


On-Chip Data Communication

Analysis, optimization and circuit design

Daniël Schinkel



Samenstelling promotiecommissie:

Voorzitter:	prof. dr. ir. A.J. Mouthaan	Universiteit Twente
Secretaris:	prof. dr. ir. A.J. Mouthaan	Universiteit Twente
Promotor:	prof. ir. A.J.M. van Tuijl	Universiteit Twente
Assistent Promotor:	dr. ing. E.A.M. Klumperink	Universiteit Twente
Deskundige:	ir. G.W. den Besten	NXP, Eindhoven
Referent:	dr. ir. M.J.M. Pelgrom	NXP, Eindhoven
Leden:	prof. dr. ir. B. Nauta	Universiteit Twente
	prof. dr. ir. G.J.M. Smit	Universiteit Twente
	prof. dr. J. Pineda de Gyvez	Technische Univ. Eindhoven
	dr. ir. N. P. van der Meijs	Technische Univ. Delft

CTIT

University of Twente
Centre for Telematics and Information Technology (CTIT)
P.O. Box 217
7500 AE Enschede
The Netherlands



This research is supported by the Dutch Technology Foundation STW, which is part of the Netherlands Organisation for Scientific Research (NWO) and partly funded by the Ministry of Economic Affairs, Agriculture and Innovation.

print: Gildeprint Drukkerijen – www.gildeprint.nl

© 2011, Daniël Schinkel, Enschede, The Netherlands

ISBN: 978-90-365-3202-0

ISSN: ISSN 1381-3617, CTIT Ph.D. thesis series No. 11-199

DOI: 10.3990/1.9789036532020

ON-CHIP DATA COMMUNICATION

ANALYSIS, OPTIMIZATION AND CIRCUIT DESIGN

PROEFSCHRIFT

ter verkrijging van
de graad van doctor aan de Universiteit Twente,
op gezag van de rector magnificus,
prof. dr. H. Brinksma,
volgens besluit van het College voor Promoties
in het openbaar te verdedigen
op vrijdag 24 juni 2011 om 14:45 uur

door

Daniël Schinkel
geboren op 5 juni 1978
te Finsterwolde

Dit proefschrift is goedgekeurd door:

de promotor prof. ir. A. J. M. van Tuijl

de assistent promotor dr. ing. E. A. M. Klumperink

Contents

ABSTRACT	11
SAMENVATTING	13
DANKWOORD	15
CHAPTER 1 INTRODUCTION.....	17
CHAPTER 2 ON-CHIP INTERCONNECTS, SCALING AND DIMENSIONING ...	19
2.1 INTRODUCTION	19
2.2 HIERARCHICAL INTERCONNECTS	19
2.3 INTERCONNECTS FOR DATA COMMUNICATION.....	21
2.3.1 <i>Interconnect length and Rent's rule</i>	21
2.3.2 <i>Global interconnects and architectures</i>	23
2.4 ELECTRICAL PARAMETERS FOR INTERCONNECTS.....	24
2.5 INTERCONNECTS AND TECHNOLOGY SCALING	25
2.6 TECHNOLOGICAL INTERCONNECT ADVANCES.....	28
2.6.1 <i>Implemented improvements</i>	29
2.6.2 <i>Future improvements</i>	29
2.6.3 <i>Reverse scaling</i>	31
2.6.4 <i>Combination with architectural and circuit improvements</i>	31
2.7 INTERCONNECT DIMENSIONING	32
2.7.1 <i>Bandwidth per cross-sectional area optimization</i>	32
2.7.2 <i>Bandwidth per pitch optimization</i>	34
2.7.3 <i>Bandwidth optimization in general</i>	35
2.8 SUMMARY AND CONCLUSIONS.....	36
CHAPTER 3 INTERCONNECT CHARACTERIZATION AND MODELING	39
3.1 INTRODUCTION	39
3.2 INTERCONNECTS IN THIS PROJECT.....	39
3.3 INTERCONNECT PARAMETER EXTRACTION	41
3.4 INTERCONNECT TRANSFER FUNCTION.....	42
3.5 INFLUENCE OF INDUCTANCE	43
3.5.1 <i>Influence of inductance on interconnect transfer</i>	44
3.5.2 <i>Influence of inductance on propagation delay</i>	45
3.6 THE SKIN-EFFECT	47

3.6.1	<i>Influence of skin-effect on the transfer function</i>	48
3.7	CONCLUSIONS ON INDUCTANCE AND SKIN-EFFECT	52
3.8	INTERCONNECT MODELING FOR CIRCUIT DESIGN	54
3.8.1	<i>Classical delay models</i>	55
3.8.2	<i>Elmore delay model</i>	55
3.8.3	<i>Multi-drop buses and their Elmore delay</i>	56
3.8.4	<i>Inductance and termination extensions to Elmore delay</i>	58
3.8.5	<i>Higher-order (transfer) models</i>	59
3.8.6	<i>Lumped models</i>	61
3.9	SUMMARY AND CONCLUSIONS.....	64
CHAPTER 4 TERMINATION, CROSSTALK AND POWER CONSUMPTION		65
4.1	INTRODUCTION	65
4.2	INTERCONNECT TERMINATION.....	65
4.2.1	<i>Classical and characteristic termination</i>	66
4.2.2	<i>Resistive RX or Capacitive TX termination and their similarities</i>	69
4.2.3	<i>Differences between a resistive receiver and a capacitive transmitter</i>	72
4.2.4	<i>RL receiver termination</i>	74
4.2.5	<i>Other types of termination</i>	76
4.3	CROSSTALK	77
4.3.1	<i>Capacitive crosstalk problem</i>	78
4.4	DIFFERENTIAL TWISTED WIRES FOR CROSSTALK REDUCTION	81
4.4.1	<i>Costs and benefits</i>	82
4.4.2	<i>Crosstalk in differential wires without twists</i>	83
4.4.3	<i>Modal analysis for crosstalk signals</i>	83
4.4.4	<i>Twist analysis and positioning</i>	84
4.4.5	<i>Quantitative results for delay and crosstalk</i>	85
4.4.6	<i>Twists to reduce common-mode crosstalk</i>	87
4.4.7	<i>Twisting patterns to reduce crosstalk in Multi-layer buses</i>	87
4.5	INTERCONNECT POWER.....	91
4.5.1	<i>Classical interconnect power consumption</i>	91
4.5.2	<i>General model for interconnect power consumption</i>	92
4.5.3	<i>Power efficiency versus signaling bandwidth</i>	96
4.6	SUMMARY AND CONCLUSIONS.....	97
CHAPTER 5 DATA COMMUNICATION ANALYSIS.....		99
5.1	INTRODUCTION	99
5.2	GENERAL VERSUS ON-CHIP DATA COMMUNICATION.....	99
5.3	DATA TRANSMISSION WITH FINITE BANDWIDTHS AND CROSSTALK	102
5.3.1	<i>Reliable data detection and eye diagrams</i>	102
5.3.2	<i>Eye diagram properties</i>	103
5.3.3	<i>Eye diagrams and crosstalk</i>	104
5.4	SYMBOL RESPONSE ANALYSIS	104
5.4.1	<i>Symbol response introduction</i>	104
5.4.2	<i>Linear models for communication systems</i>	106
5.4.3	<i>Maximum interference and eye openings</i>	108
5.4.4	<i>Complex signal analysis versus separation of I and Q</i>	113

5.4.5	<i>Statistical analysis</i>	115
5.4.6	<i>Remarks on symbol-response analysis</i>	118
5.5	SYNCHRONIZATION	120
5.6	SUMMARY AND CONCLUSIONS.....	121
CHAPTER 6 SIGNALING AND MODULATION TECHNIQUES.....		123
6.1	INTRODUCTION	123
6.2	PLAIN BINARY SIGNALING	123
6.2.1	<i>Achievable data rate with and without crosstalk</i>	123
6.2.2	<i>Achievable data rate with differential twisted wires</i>	126
6.3	ANALYSIS SIMPLIFICATIONS FOR BASEBAND SIGNALING	129
6.3.1	<i>Eye properties for PAM with first-order channel models</i>	130
6.3.2	<i>Eye properties for binary signaling with first-order channel models</i>	133
6.4	MULTI-LEVEL SIGNALING.....	133
6.4.1	<i>Eye properties for M-ary signaling with first-order channel models</i>	134
6.4.2	<i>M-ary eye properties with higher-order channel models</i>	135
6.4.3	<i>Arguments for and against M-ary signaling ($M > 2$)</i>	136
6.5	ACHIEVABLE RATES FOR BAND-PASS SIGNALS.....	137
6.5.1	<i>Single carrier PAM modulation</i>	138
6.5.2	<i>Single carrier quadrature modulation</i>	141
6.5.3	<i>Multi-Carrier and OFDM or CDMA</i>	142
6.6	SUMMARY AND CONCLUSIONS.....	144
CHAPTER 7 EQUALIZATION TECHNIQUES.....		147
7.1	INTRODUCTION	147
7.2	EQUALIZATION OVERVIEW	148
7.2.1	<i>Transmitter-side equalization</i>	149
7.2.2	<i>Receiver-side equalization</i>	149
7.2.3	<i>Transmitter and receiver equalization</i>	150
7.2.4	<i>Adaptive equalization</i>	151
7.2.5	<i>Adaptive equalization and clock recovery</i>	152
7.3	FIR-PRE-EMPHASIS	153
7.3.1	<i>FIR pre-emphasis with first-order channel models</i>	154
7.3.2	<i>Achievable data rate with FIR pre-emphasis for on-chip wires</i>	156
7.4	PULSE-WIDTH PRE-EMPHASIS	156
7.4.1	<i>PW pre-emphasis with first-order channel models</i>	157
7.4.2	<i>Achievable data rate with PW pre-emphasis for on-chip wires</i>	159
7.5	FIR VERSUS PW PRE-EMPHASIS.....	159
7.5.1	<i>Differences for on-chip and off-chip applications</i>	159
7.5.2	<i>Implementation differences</i>	160
7.5.3	<i>FIR pre-emphasis and capacitive transmitters</i>	161
7.6	DECISION FEEDBACK EQUALIZATION.....	161
7.6.1	<i>DFE with continuous-time feedback filter</i>	162
7.6.2	<i>Continuous-time DFE with first-order channel models</i>	163
7.6.3	<i>Achievable data rate with continuous-time DFE for on-chip wires</i>	164
7.7	EQUALIZATION AND PROCESS SPREAD.....	166
7.7.1	<i>Dealing with mismatch at design time</i>	166

7.7.2	<i>Adaptive equalization for on-chip transceivers</i>	167
7.8	EQUALIZATION COMBINED WITH M-PAM	168
7.9	SUMMARY AND CONCLUSIONS	170
CHAPTER 8 FIRST DEMONSTRATOR IC		171
8.1	INTRODUCTION	171
8.2	INTERCONNECT ANALYSIS AND DIMENSIONING	172
8.2.1	<i>Interconnect Model</i>	172
8.2.2	<i>Twisted differential interconnects</i>	173
8.3	PULSE-WIDTH PRE-EMPHASIS	174
8.4	TRANSCIVER IMPLEMENTATION	176
8.4.1	<i>Transmitter</i>	176
8.4.2	<i>Receiver</i>	177
8.5	COMPARISON WITH REPEATERS	178
8.5.1	<i>Receiver clocking</i>	180
8.6	DEMONSTRATOR IC TOP-LEVEL	180
8.7	MEASUREMENT SETUP	184
8.8	EXPERIMENTAL RESULTS	186
8.8.1	<i>Parameter characterizations</i>	186
8.8.2	<i>Signal measurements</i>	187
8.9	CONCLUSIONS FROM FIRST DEMONSTRATOR IC	190
CHAPTER 9 IMPROVED SENSE AMPLIFIER		193
9.1	INTRODUCTION	193
9.2	CONVENTIONAL SENSE AMPLIFIER AND ITS DRAWBACKS	194
9.3	DOUBLE-TAIL SENSE AMPLIFIER	196
9.4	SENSE AMPLIFIER SPEED, OFFSET AND NOISE ANALYSIS	197
9.4.1	<i>Double-tail sense amplifier dimensioning for low offset</i>	198
9.5	COMPARISON OF DOUBLE-TAIL WITH CONVENTIONAL	199
9.6	SENSE AMPLIFIER MEASUREMENTS	201
9.7	SENSE AMPLIFIER CONCLUSIONS	204
CHAPTER 10 TRANSCIVER ON THE SECOND DEMONSTRATOR IC		205
10.1	INTRODUCTION	205
10.2	EFFECT OF TERMINATION ON BANDWIDTH AND POWER	206
10.3	TRANSCIVER IMPLEMENTATION	208
10.3.1	<i>Capacitive pre-emphasis transmitter</i>	208
10.3.2	<i>Sense amplifier with decision feedback equalization</i>	210
10.4	DEMONSTRATOR IC TOP-LEVEL AND MEASUREMENT SETUP	211
10.5	EXPERIMENTAL RESULTS	214
10.6	CONCLUSIONS FOR TRANSCIVER ON SECOND DEMONSTRATOR IC	217
CHAPTER 11 TRANSCIVERS FOR NETWORKS ON CHIPS		219
11.1	INTRODUCTION	219
11.2	DATA COMMUNICATION ON A NOC	221
11.2.1	<i>Interconnects for Networks on a Chip</i>	221
11.2.2	<i>Conventional data transmission</i>	222

11.2.3	<i>Link improvements</i>	223
11.3	LOW-SWING TRANSMITTERS	224
11.4	RECEIVER AND OPTIMAL SWING	228
11.5	COMPLETE TRANSCIEVER	230
11.5.1	<i>Transceiver with synchronization</i>	230
11.5.2	<i>Cascaded transceivers</i>	231
11.6	CONCLUSIONS ON NOC TRANSCIEVERS	234
CHAPTER 12 CONCLUSIONS AND RECOMMENDATIONS		235
12.1	CONCLUSIONS	235
12.2	ORIGINAL CONTRIBUTIONS	238
12.3	RECOMMENDATIONS FOR FURTHER STUDY	239
12.3.1	<i>Recommendations on side-topics</i>	241
LIST OF PUBLICATIONS		245
ABOUT THE AUTHOR		247
APPENDIX A STANDARD DEVIATION ESTIMATION IN COMPARATORS ...		249
A.1	ACCURACY OF STANDARD DEVIATION ESTIMATION	249
A.2	DECISION AVERAGING VERSUS IMPEDANCE SCALING	252
APPENDIX B OVERVIEW OF ACHIEVABLE DATA RATES		255
REFERENCES		257

Abstract

On-chip data communication is an active research area, as interconnects are rapidly becoming a speed, power and reliability bottleneck for digital CMOS systems. Especially for global interconnects that have to span large parts of a chip, there is an increasing gap between transistor speed and interconnect bandwidth. To alleviate this problem, improvements in technology, architectures and circuits are needed. On the technology side, low-k dielectrics and reverse scaling can improve the interconnect behavior. On the architecture side, Network on chips (NoCs) can reduce the number of global interconnects. On the circuit side, which is the focus area of this thesis, more advanced strategies than the classical repeater insertion can be used to reduce the power consumption and increase the communication speed.

In the thesis, it is shown that the bandwidth of interconnects is either limited by their distributed RC behavior (for long interconnects), or by the skin-effect. In both cases, the bandwidth is proportional to the cross-sectional area and inversely proportional to the length squared. The aggregate bandwidth per cross-sectional area can be optimized by choosing all cross-sectional dimensions roughly equal. The bandwidth of a single interconnect can be increased by using resistive (or resistive-inductive) receiver termination or capacitive transmitter termination. The crosstalk can be mitigated with twisted differential interconnects, where the number of twists determines for how many neighbors the crosstalk can be cancelled. With the aid of a symbol response analysis method, it is shown that simple equalization schemes are very effective to boost the achievable data rate, more so than multi-level signaling or band-pass modulation.

To validate the concepts two demonstrator ICs were developed, both using 10mm long interconnects. The first chip, in a 130nm CMOS process, showed that a combination of pulse-width pre-emphasis, twisted interconnects and low-ohmic receiver termination can boost the data rate to 3Gb/s/ch (at 2pJ/bit), while a conventional transceiver reached only 0.55Gb/s/ch. The second test-chip, in 90nm CMOS, showed that a combination of a capacitive transmitter and a low-power sense-amplifier with DFE at the receiver can reduce the energy consumption to 0.28pJ/bit (at 2Gb/s), much lower than competing designs.

Circuit simulations show that a capacitive transmitter and a low-power sense amplifier can also be very effective as transceivers in a NoC, with data rates in excess of 9Gb/s (at 130fJ/transition) over 2mm interconnects. Multiple transceivers can be connected back-to-back to create a source-synchronous transceiver-chain with a wave-pipelined clock, operating with 6σ offset reliability at 5 Gb/s.

Samenvatting

Data communicatie binnen geïntegreerde elektronische schakelingen (chips) is tegenwoordig een actief onderzoeksgebied omdat de metalen verbindingen een limiterende factor aan het worden zijn wat betreft snelheid, vermogensverbruik en betrouwbaarheid van digitale CMOS systemen. Met name de lange verbindingen die grote delen van de chip moeten overbruggen worden steeds trager ten opzichte van transistors. Om dit probleem op te lossen zijn er verbeteringen nodig in zowel technologie als architectuur als circuits. Op technologieniveau kunnen isolerende materialen met lage diëlektrische constanten verbetering bieden, tezamen met meer dikkere metaallagen. Op architectuurniveau kunnen de zogenaamde ‘netwerken op chips’ (NoC’s) het aantal lange verbindingen beperken. Het onderzoek naar verbeteringen op circuitniveau is het onderwerp van dit proefschrift. Traditioneel worden er simpele repeterende versterkers gebruikt om verbindingen te versnellen, maar geavanceerdere circuits kunnen het vermogensverbruik reduceren en de communicatiesnelheid verhogen.

In dit proefschrift wordt aangetoond dat de bandbreedte van de verbindingen beperkt wordt door hetzij een gedistribueerd RC gedrag (voor lange verbindingen), hetzij het zogenaamde ‘skin-effect’. In beide gevallen is de bandbreedte evenredig met de oppervlakte van de dwarsdoorsnede van de verbinding en omgekeerd evenredig met de lengte in het kwadraat. De som van de bandbreedtes van een aantal verbindingen binnen een bepaalde dwarsdoorsnede kan worden geoptimaliseerd door alle dwars-afmetingen gelijk te kiezen. De bandbreedte van een enkele verbinding kan worden verhoogd door een resistieve (of resistieve en inductieve) afsluiting te gebruiken aan de ontvangtzijde of door een capacitieve serie afsluiting te gebruiken aan de zenzijde. De overspraak tussen verbindingen kan worden verminderd door gevlochten aderen te gebruiken. Hierbij bepaalt het aantal draaiingen in een aderpaar van hoeveel naburige aderen de overspraak kan worden onderdrukt.

In dit proefschrift wordt ook een analysemethode gepresenteerd die is ontwikkeld om de effectiviteit van verschillende data transmissie technieken te kunnen kwantificeren en die werkt op basis van de symbool respons. Met behulp van deze methode is aangetoond dat simpele egalisatie technieken zeer effectief zijn om de bandbreedte te vergroten, veel effectiever dan signalen met meer dan twee niveau’s of met banddoorlaat modulatie technieken.

Om de hierboven genoemde concepten te kunnen valideren zijn twee demonstratie chips ontwikkeld. In beide chips zijn verbindingen van 10mm lang gebruikt. De eerste chip is

gefabriceerd in 130nm CMOS technologie en had als doel om pulsbreedte egalisatie, gevlochten verbindingen en een laagohmige afsluiting aan de ontvangstzijde te kunnen testen. De combinatie van deze technieken maakte een communicatie snelheid mogelijk van 3Gb/s per kanaal (bij een vermogensverbruik van 2pJ/bit), ten opzichte van 0.55Gb/s per kanaal met conventionele circuits. De tweede chip is gefabriceerd in 90nm CMOS en met deze chip werd aangetoond dat het mogelijk is om het vermogensverbruik naar beneden te brengen tot 0.28pJ/bit (bij 2Gb/s). Dit lage vermogensverbruik – veel lager dan concurrerende circuits – werd bereikt door een combinatie van een capacitieve zender en een ontvanger op basis van een energiezuinige detectie versterker met ingebouwde egalisatie.

Met simulaties is aangetoond dat een capacitieve zender en een energiezuinige detectie versterker ook zeer geschikt zijn voor communicatiecircuits in NoC systemen. De circuits maken snelheden mogelijk van meer dan 9Gb/s (bij een verbruik van 130fJ/transitie) over 2mm lange verbindingen. Meerdere van dergelijk circuits kunnen ook achter elkaar geplaatst worden om een communicatie-keten op te bouwen, inclusief synchronisatie vanuit de bron. Het resulterende systeem werkt op een snelheid van 5Gb/s en is robuust, met een verwachte uitval door spreiding van slechts 2 op de miljard exemplaren (6σ).

Dankwoord

De laatste woorden van het proefschrift, eindelijk. Het heeft er op enkele momenten om gespannd of ik op dit punt aan zou komen, dus ik ben erg blij dat het nu zo ver is. Alhoewel er volgens mij weinig promovendi zijn die de afronding van het proefschrift eenvoudig vinden, geloof ik ook niet dat ik de makkelijkste route genomen heb. Een bedrijf starten en kinderen krijgen zijn beide grote levensveranderende projecten en niet heel makkelijk te combineren met de laatste loodjes van een promotie. Desondanks, als ik het nu kon overdoen zou ik het niet anders doen (wat betreft het bedrijf en kinderen).

Dat het zover gekomen is heb ik aan diverse mensen te danken. Een aantal hiervan wil ik hieronder in het bijzonder bedanken.

Als eerste natuurlijk mijn promotor Ed van Tuijl. Het is al weer twaalf jaar geleden dat je me voor het eerst begeleidde (met vier studenten werkten we aan een audio compressie project) en sinds die tijd heb ik met veel plezier met je gewerkt aan vele projecten, inclusief de oprichting van Axiom IC B.V.

Daarnaast ben ik ook veel dank verschuldigd aan mijn assistent-promotor Eric Klumperink, met name voor het leiden van het project en voor de gedetailleerde feedback op het manuscript. Bram Nauta, leerstoelhouder van de IC-design groep, wil ik bedanken voor de motiverende discussies om toch vooral (wat sneller) het proefschrift af te maken en natuurlijk ook voor de leuke windsurfessies. Ook dank aan de STW die dit project mogelijk maakte en aan de gebruikerscommissie voor alle discussies.

Het onderzoek in dit promotieproject deed ik gelukkig niet alleen, maar samen met Eisse Mensink. Eisse, bedankt voor de goede samenwerking. Binnenkort kan ik dan eindelijk onze afspraak van een wederzijds paranimschap nakomen.

Naast het onderzoek heb ik in de tijd dat ik bij de IC-design groep werkte een leuke tijd gehad, waarvoor ik de volgende personen met name wil danken. Mijn kamergenoten Mustafa, Eisse en Kasra voor de plezierige technische en niet-technische discussies. Natuurlijk ook Gerdien, Cor, Frederik, Gerard en Henk voor alle ondersteuning.

Ook de Universiteit Twente ben ik dankbaar. Mijn meeste vrienden heb ik hier tijdens mijn studietijd leren kennen. De UT is niet alleen een geweldige plek om kennis te vergaren en onderzoek te doen, maar de campus is met al het groen ook een heerlijke plek om te vertoeven. Dank aan Joost Kauffman en Annet Schenk voor de heerlijk ontspannende lunchwandelingen tijdens het aioschap. Ook dank aan mijn oud-huisgenoten van de WBW. Dat we na al die jaren nog steeds zoveel lol kunnen hebben sterkt mij in mijn geloof

dat ik geen betere plek had kunnen treffen om de eerste jaren van de studie door te brengen. Veel dank ook aan Kiman Velt, Steven Leussink en Wouter Groothedde. We kennen elkaar al sinds de allereerste dag dat we elektrotechniekstudent werden en we zijn nog steeds niet uitgepraat.

Inmiddels ben ik alweer vier jaar bezig met Axiom IC - een geweldig vervolg op het promotieproject. Wat een buitenkans voor iemand die altijd al iets ondernemends wilde doen, maar niet precies wist hoe hij dat aan moest pakken. Bij deze dank aan mijn vier mede-oprichters en aan alle collega's.

Ten slotte wil ik mijn familie van harte bedanken. Mijn ouders, omdat ze mij de vrijheid gaven mijzelf te ontwikkelen en me toch ook altijd bleven uitdagen verder te kijken dan mijn eigen interesses (als ik alles zelf had mogen beslissen was ik nu misschien kraanmachinist geweest). En natuurlijk mijn partner Henriët, mijn steun en toeverlaat en moeder van onze kind(eren). Bedankt dat je het na je eigen promotie nog al die tijd met een wannabee-doctor hebt uitgehouden en al die avonduren proefschrift schrijven mogelijk hebt gemaakt.

Chapter 1

Introduction

Over the last 50 years, integrated circuits have seen an immense progress, from the early developments that contained only a few transistors to the current microprocessors that can contain billions of elements. This progress has been made possible by a continued downscaling of the circuit dimensions, which goes hand in hand with an increase in transistor speed and a reduction in cost per transistor, famously known as ‘Moore’s law’.

However, not all aspects of a circuit improve with a reduction of their size. This is most notably true for the wires that interconnect the transistors (the interconnects). The resistance of interconnects increases disproportionately when their cross-sectional dimensions are reduced, which makes them slower when they are scaled down.

Already back in the 1970’s, this potential showstopper for continued scaling was brought up in a well-known paper by Dennard [1], but back then the interconnects were still far away from becoming a bottleneck. But over the last decade, the interconnects indeed have become a real limiting factor for large digital integrated circuits – which are nowadays made almost exclusively in CMOS technologies. This is especially the case for those interconnects that are used for data communication from one block on the chip to another and hence need to bridge ‘large’ distances.

The problem of interconnect scaling thus received renewed attention, and a number of improvements have been suggested. Some technological improvements have already been implemented, such as the introduction of copper interconnects which have lower resistivity than their aluminum predecessors. Other improvements are in progress, such as the move towards insulators with lower capacitance (the low-k dielectrics). In the more distant future other opportunities for improvements might become viable, such as 3D integration (multiple chips on top of each other) or optical interconnects, but it remains to be seen whether these technologies will really leave the research phase.

Next to these technological advancements, there is also room for much improvement in the circuits that are used for on-chip data communication. The existing approach is to use simple repeater circuits that are placed along the wire to boost the signal. However, repeaters already cost quite some chip area and power consumption, and their number is projected to rise rapidly in future IC technology generations.

The central theme in this PhD project - of which this thesis is one of the results - is how circuit techniques can be used to improve on-chip data communication. This project was carried out by two PhD students, Eisse Mensink [2] and this author. Over the course of the project, four major topics were investigated. The first is how the wires themselves can be optimized for high-speed data communication, within the boundaries set by the technology. The second is how the effect of crosstalk between the wires can be reduced. The third is what type of signaling methods are most suitable for on-chip communication and the fourth is how these signaling methods can be implemented with power and area-efficient circuits. The two main criteria that were used in these investigations are, one: how can the speed of the communication be improved, and two: how can the power consumption of the communication be reduced.

In this thesis, it will be shown that it is possible to optimize the interconnects for data transmission by choosing their width and height approximately equal. It will also be shown how twisted differential wires can reduce crosstalk and how equalization and wire termination can be used to optimize the speed of the communication. A number of circuit improvements will be presented, of which a capacitive transmitter with an optimized sense amplifier is the best candidate for low power high-speed communication.

This thesis is roughly divided into three parts. In the first part, the interconnects themselves are discussed. This part starts in the next chapter with an introduction to on-chip interconnects, how they scale over technologies and how their physical properties can be optimized for data communication. It is followed by an analysis of interconnect transfer functions in Chapter 3. That chapter also discusses models of different degrees of complexity to capture the interconnect behavior. Chapter 4 discusses other interconnect topics important for data communication, namely interconnect termination, crosstalk and power consumption.

In the second part of the thesis, data communication is discussed and how it can be best applied to on-chip communication. Chapter 5 presents techniques for the analysis of the achievable speeds for data communication over bandlimited channels (such as on-chip interconnects). In the next two chapters, these techniques are applied, first to modulation methods in Chapter 6 and then to equalization techniques in Chapter 7. A lot of quantitative data is generated in this part of the thesis, which is summarized in Appendix B.

In the third part of the thesis, practical circuits for on-chip data communication are discussed, applying the results of the first two parts. Two demonstrator IC's were made in the course of this project to validate the proposed methods and circuits. The first demonstrator IC is discussed in Chapter 8. As part of the second demonstrator IC, a more widely usable building block – a clocked comparator – was optimized, which is discussed separately in Chapter 9, with some background analysis in Appendix A. The transceiver on the second demonstrator IC is discussed in Chapter 10. The third part of the thesis concludes with Chapter 11, where it is discussed how the circuits from the second demonstrator IC can be adapted and optimized further for application in 'Networks on a Chip' (NoCs), an emerging strategy for on-chip communication. The last chapter of the thesis summarizes the results and conclusions from the earlier parts and presents recommendations for further study.

Chapter 2

On-chip Interconnects, scaling and dimensioning

2.1 *Introduction*

This chapter presents a background for on-chip interconnects. It is discussed how they are used, what their basic properties are, how they scale over technology generations, and how the resulting scaling problem manifests itself. This scaling problem is most severe for global, chip-wide interconnects. This is because the interconnect dimensions – both their length and their cross-sectional parameters – play a vital role in determining the interconnect bandwidth. For the highest bandwidth, the interconnect length should be kept short and the cross-sectional dimensions large. This chapter briefly discusses a number of architectural and technological advancements that aim to do this, including a short discussion on methods that try to tackle the problem in a whole different way. As circuit designers also have some control over the interconnect cross-sectional dimensions, an analysis is presented that predicts the desired dimensions for the highest data rate.

The chapter starts in the next section with a general interconnect overview. Section 2.3 discusses the use of interconnects for data communication. Section 2.4 discusses interconnect parameters and section 2.5 shows how these parameters affect interconnect performance and how they scale over technology. Section 2.6 discusses advancements in interconnect technology. Section 2.7 presents the analysis on optimal cross-sectional dimensions.

2.2 *Hierarchical interconnects*

On-chip interconnects are of course vital components of any chip in any technology. Without them, the various devices on a chip could not be connected and integrated circuits would not exist. Large-scale digital integrated circuits nowadays usually use hierarchical designs and interconnection styles. At the lowest level – the local level – metallic wires or wires from semiconducting materials (such as polysilicon) interconnect the various devices in a small circuit, for example a digital gate or flip-flop. At the next level – the intermediate level – metallic wires interconnect the different sub-circuits (gates) to form larger




Figure 2.1: Cross-sectional view of hierarchical wiring approach. Source:[3].

functional blocks such as an ALU-unit, a multiplier, or a memory bank. At the highest level – the global level – interconnects are used to create communication fabrics such as busses or even on-chip networks to link all the functional blocks together.

The hierarchical multiple tier interconnect structure is reflected in current CMOS IC processing technologies (which is the standard technology for almost every large-scale digital circuit), with small pitched wires at the lowest metal layers and large, thick wires at the top metal layers, as visible in Figure 2.1.

Next to the interconnects that are used for data signals, a large number of interconnects are used for the distribution of power. These power (V_{DD}) and ground wires can span the entire chip and are often organized in mesh-type grids with thick and wide wires at the highest metal levels [4, 5]. An example of such a mesh configuration is shown in Figure 2.2. This configuration makes low impedance power connections available throughout the area of the chip. Especially in the top metal layers, quite a large percentage of the wires can be reserved for power distribution.

Clock distribution also occupies a significant part of the interconnect fabric, apart from circuits that use asynchronous design styles. Most often, a tree structure is used for the distribution of the clock [6] with wide wires (and large buffers) for the chip-wide top-level distribution and finer wires at lower levels. An example of a common clock tree with hierarchical wiring is shown in Figure 2.3. This H-tree has the nice property that the clock (ideally) has equal delay at every end-branch.

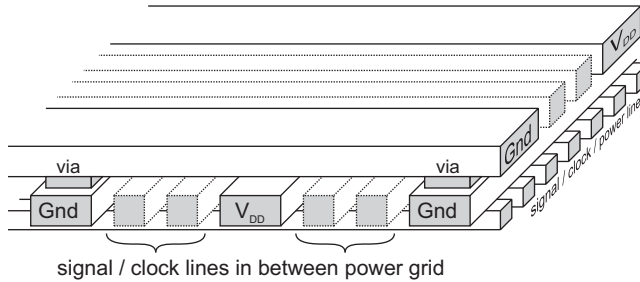


Figure 2.2: Multi-layer Power grid, with a mesh of Gnd and V_{DD} wires, possibly with signal or clock wires in between.




Figure 2.3: H-tree for low skew clock distribution.

So the three purposes of interconnects, signal transportation, clock distribution and power distribution all use hierarchical wiring structures and compete for the same wiring resources. At the top-metal layers, power and clock distribution are the dominant purpose for the interconnects. In [5] it is argued that the percentage of the top metal layers that is occupied by the power grid increases as the process technologies scales down, leaving little room there for other types of interconnects in future CMOS processes. Fortunately, the number of metal layers available also increases over process generations, to facilitate the ever increasing demand in wiring resources.

2.3 Interconnects for data communication

In this thesis we focus on interconnects for data communication (digital signal transportation), as was discussed in the introduction. We will especially focus on data communication over long wires (global data communication), as that is the type of interconnect that poses the highest limitations in current and especially future CMOS processes [7, 8].

Of course, there are far fewer global interconnects that span large portions of the chip than there are short, local interconnects, but the global wires still play a vital role in integrated circuits. They are for example used for on-chip buses to connect the different parts of a microprocessor or a system on a chip (SoC) [6]. They can also be found in memories, as global address or data-lines, or to interconnect the different levels of caches.

2.3.1 Interconnect length and Rent's rule

To get a general indication of how many (signal or data) interconnects of a given length can be found on a chip, Rent's rule is often used [6, 7, 9, 10]. Rent's rule gives a simple




Figure 2.4: On-chip network in a mesh configuration.

empirical relationship between the number of wires (K) that cross the boundary of a circuit-block, as a function of the number of transistors or nodes within the block (N) and the number of interconnections inside the block (k):

$$K = k \cdot N^p \quad (2.1)$$

With p being the Rent exponent, which usually varies between 0.55 for regular circuits such as a memory up to 0.85 for highly irregular circuits such as random logic (automatically synthesized logic). Rent's rule was originally used to predict the number of I/O pins for a module as a function of the number of gates inside that module [11], but it also proved valuable as a basis for the prediction of wire length distributions [9]. By simplifying the analytical results in [9] and removing some of the more higher-order modeling, we can make a simple estimate for the wire density as a function of wire length l :

$$i(l) \approx C \cdot (2L - l)^2 \cdot l^{2p-3} \quad (2.2)$$

Where p is the Rent exponent, L is the size of the chip and C is a constant that depends on a number of factors, most notably on the number of transistors on the chip. The formula starts to become valid for lengths exceeding the transistor size. For small lengths, the distribution is roughly proportion to l^{2p-3} . With p being about 0.8 for microprocessors investigated in [9], this amounts to $i(l) \propto l^{-1.4}$. For large lengths, the distribution decreases more rapidly because the distribution is naturally cut off at lengths exceeding path-lengths ($2L$) between opposite corners of the chip. Graphical overviews for wire-density distributions of actual chips are given in [6] (page 41) and [9], which indeed have shapes corresponding to (2.2). The graph in [6] also shows that interconnects on processors from the Intel Pentium series have lengths of up to about 20mm.

When we integrate the wire density, starting at the gate size l_0 , then we get the cumulative distribution $I(l)$, giving the total wire-count up to a certain length. Assume for example that L is 1cm and that l_0 is 10000 times smaller than L (wire length starting at $1\mu\text{m}$). Then, evaluation of $I(l)$ predicts that only 3% of the wires are longer than one tenth of the chip-size (1mm) and only 700ppm of the total wires are longer than the chip size (10mm). When we assume a smaller l_0 of for example $L/100000$, then these percentages drop further to 1% and 300ppm respectively.

2.3.2 Global interconnects and architectures

Of course, the actual number of the long, global interconnects differ widely for different IC's and is also influenced by the fact that global interconnects are becoming a significant performance bottleneck. Global interconnects for high-speed data communication are for example often broken down into smaller segments with inverters as signal amplifiers (repeaters) in between [7]. These repeaters prevent signal deterioration due to e.g. bandwidth bottlenecks, just as they do in off-chip communication over for example long intercontinental data cables. To give an indication of some practical numbers, consider for example the Cell processor [12]. This processor contains about 234M transistors, connected by 1.4M nets (probably excluding the nets inside the gates). It also contains a total of 580k repeaters, of which 32k are used to ensure signal integrity for global nets.

Repeaters however are also not ideal, as will be seen later on, so perhaps future interconnection styles will become more locally oriented, either because of advances in CAD tools [8] or because of changes in chip architecture. It is often argued that new chip architectures are needed, not only because on-chip interconnects are becoming a performance bottleneck, but also because systems on chips are becoming so complex that they require new interconnection approaches [13, 14].

Networks on chips (NoCs) have emerged as such a new approach, they should be suitable to connect the many functional elements on present and future SoCs [13-18]. In these NoCs, global communication is carried out over a network, with routers as network nodes that interconnect with each other and with the functional elements on the chip. These functional elements are usually called processing elements, but they can be any circuit that generates or requires data, including input/output circuitry. An often used NoC topology is a mesh network, as shown in Figure 2.4. A mesh topology has the advantage that global wires for data communication are omitted altogether.

Still, also in these developing architectures, the availability of fast global interconnects will be desirable. A NoC for example can benefit from circular network topologies, such as torus or folded torus configurations [14], which require longer interconnects than the standard mesh topology. Wherever the trend in architectures leads to, one thing remains certain and that is that global communication is a vital aspect of digital chips. How this global communication is arranged in the future remains to be seen and will depend on the specific communication requirements. The options are: either directly over long (un-interrupted) interconnects or with interruptions along its path, whether these interruptions be in the form of simple repeaters or in the form of more advanced network routers. Examples of these different arrangements will resurface in various parts of this thesis, accompanied by discussions of their advantages and disadvantages.




Figure 2.5: Interconnects and their distributed resistance, inductance and capacitance.

2.4 Electrical parameters for interconnects

As far as data communication is concerned, on-chip interconnects have three important parameters, as shown in Figure 2.5. First, the distributed capacitance (C' in F/m), consisting of a number of contributing parts to the different conductors in the surrounding environment. Second, the distributed resistance (R' in Ω/m), as defined by the cross-section and conductivity of the interconnect. Third, the distributed inductance (L' in H/m), which complements the capacitance and together they create the well-known transmission-line behavior.

Sometimes, a fourth parameter, the (frequency-dependent) shunt conductance (G' in $1/\Omega m$) is used in interconnect models, in analogy with standard transmission line parameters where it is used to describe losses in the dielectrics. However, dielectric losses in on-chip interconnects are insignificant compared to other losses [19]. Although G' can also be used to model losses in for example return paths [19], it is not a meaningful physical parameter in that sense, nor is it a necessary parameter (return-paths can be modeled in other ways). In [2], values for G are obtained, but for the interconnects in this project they were not needed for accurate interconnect modeling. We will therefore not further use G' in this thesis, and use the more common RLC or RC models instead.

Capturing the resistance, capacitance and inductance in single valued parameters is sometimes quite difficult. The effective capacitance to ground is for example quite a complex property, influenced not only by the distances and dimensions of the neighboring interconnects, but also by the size, structure and termination impedances of these neighbors. Even the signals on the neighboring interconnects affect the capacitance when the signals are correlated. We will return to this topic in section 4.3. For now, we use the common assumption that the interconnect capacitance is simply referenced to ground, which is usually reasonably accurate for practical interconnect configurations.

Actual interconnects are also not infinitely small and processes like electrical conduction are not necessarily uniformly distributed inside the interconnect (for example due to skin-effect). In this sense, the three parameters are also a simplification of the actual properties of the interconnect. In general transmission-line theory, the parameters are often specified as a function of frequency to improve the correspondence between the models and the actual behavior. However, on-chip wires are so small that constant values usually suffice to describe their dominant behavior. The RLC parameters for the global interconnects that we analyzed in this project vary for example by less than 3% over a frequency range of 10GHz

[2]. At really high frequencies, or for very wide and thick interconnects, the skin-effect – the confinement of conduction in the outer part of a conductor at high frequencies – becomes an issue. Skin-effect does add a frequency dependency to the R and L parameter, but it turns out that the effect can still be described in terms of the original frequency independent RLC parameters, as discussed in section 3.6.

The significance of the RLC parameters has changed over time and differs per application. The inductance of on-chip interconnects for example has only recently begun to receive attention and it can still be neglected in many cases. Only in some applications are inductive effects clearly present, either by intention [20, 21] or as parasitic effect [19, 22], but for most interconnects for data communication, it is an irrelevant parameter, as will be discussed in section 3.5 and in section 3.8.4. The wire resistance is another parameter that is often disregarded (for short wires) but which effect is becoming increasingly important, as is discussed next.

2.5 Interconnects and technology scaling

Traditionally, when interconnects for CMOS data communication were concerned, IC-designers were only interested in the capacitance of the interconnect. This is because in CMOS processes digital gates usually have no static currents (apart from leakage currents) and energy costs are primarily caused by signal switching actions. The capacitance determines this energy-cost and also determines the required size of the driver to get suitably low switching times.

However, as technology feature sizes scaled down, the resistance of the interconnect also became important, because the wire-resistance increases with smaller cross-sectional dimensions. The distributed interconnect resistance and capacitance together create RC-delay and bandwidth limitations. The resistance and capacitance not only limit the bandwidth, but also create crosstalk between wires. Switching voltages on an interconnect will also pull at the voltage levels of the surrounding interconnects, mainly through capacitive coupling (see Figure 2.5). This crosstalk effect would not be present when the entire interconnect would be tied to a low-impedance driver, but the interconnect resistance weakens this link with the driver.

Parameter	scaling factor
feature size	s
operating frequency f	$1/s$
Devices	
$t_{ox}, W_{min}, L_{min}, 1/N_a, V_{DD}$	s
Delay time, $V_{dd}C_{MOST}/I_d$ (s)	s
Energy/transition $C_{MOST}V_{DD}^2$ (J)	s^3
Power density $f \cdot E/A$ (W/m ²)	1
Device density (1/m ²)	$1/s^2$
Interconnects	
Cross dimensions w, h (m)	s
Length l (m)	s
Distributed R' (Ω/m)	$1/s^2$
Distributed C' (F/m)	1
Energy/transition $C' l \cdot V_{DD}^2$ (J)	s^3
Power density $f \cdot E/A$ (W/m ²)	1
Drive delay $V_{dd} \cdot C' l / I_d$ (s)	s
Interconnect RC delay ($R' C' l^2$)	1

Table 2.1: Technology scaling and the impact on devices and interconnects, assuming Dennard scaling rules [1].

Already back in 1974, Dennard showed in his seminal paper about (constant-field) technology scaling [1], that transistors get faster with technology-scaling, but interconnects do not, as shown in Table 2.1. The table shows that, if we could neglect wire resistance, then the delay and power in interconnects scale at the same pace as the delay and power of transistors and we would have ideal scaling. But unfortunately, as interconnects get smaller cross-sections, their R' increases while the C' stays roughly equal because decreasing plate surfaces are canceled by decreasing spacings to neighboring conductors. This results in interconnect RC delays that do not track scaling parameters, with delays that stay equal over scaling and even increase when the interconnect is kept at a certain length.

For many years, this scaling discrepancy was not a problem, as the inherent time constant of the interconnects were much shorter than the time constant of the drivers. But in the past decade, after many years of successful scaling, technology feature sizes have become so small that the interconnect resistance and the associated interconnect RC time constant become a significant speed bottleneck.

Of course, actual technology scaling has deviated quite a bit from the idealized Dennard scaling. Many other hurdles have been faced along the way (such as the increasing

problems with for example leakage power consumption). But thanks to the huge efforts of many engineers, scaling still continues. Unfortunately, so does the discrepancy between transistor and interconnect delay.

In the public literature, the interconnect problem has not gone by unnoticed. Already in 1990, Bakoglu presented a comprehensive overview of the subject [7]. A number of influential papers also started to appear from the mid nineties onwards. In 1995, Bohr [23] fueled the interest in interconnect delay, by mentioning that standard techniques to keep interconnect delay within bounds – such as the addition of metal layers and the increase in aspect ratio (height over width) – were reaching their limits. Later on, in 2001, Davis et. al. [24] made a more general overview and formulated a number of limitations for interconnects, ranging from fundamental (Information theory) limitations, to material, device, circuit and system limitations. Regarding interconnect literature, 2001 was quite a productive year. Next to the paper from Davis et. al., a number of other invited papers also appeared in the proceedings of the IEEE, including the often cited papers from Deutsch et. al. [25] and Ho et. al. [8].

An interesting nuance in the discussion is the distinction between local interconnects that scale together with transistors and global ones that span large portions of the chip. This distinction was discussed in some early work [7, 26] and revitalized and applied to modern processes by Ho [8]. It is argued that the biggest problems are found in the global interconnects, which span the entire chip and are used for example for chip-wide buses. These global interconnects do not scale down in length as the perimeter of large-scale digital IC's has remained roughly constant over different technologies. Even when repeaters are added to break up these long interconnects, they will still pose a delay and bandwidth problem. Local wires on the other hand connect gates inside a functional block and the length of these wires scales down together with the gates. Ho argued that 'the relative change in speed of local wires to the speed of gates is modest', so local wires should not be the first cause of concern.

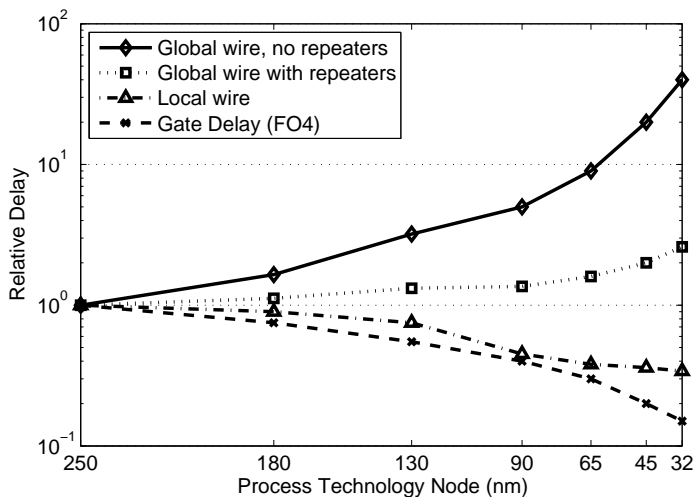


Figure 2.6: Normalized delay of gates and wires versus technology feature size.
Source:[27].

This distinction between local and global wires is also found in the 2001 ITRS roadmap [27] and in its successors. A graph from this roadmap that was often used in the interconnect delay discussion is shown in Figure 2.6. The graph clearly shows the significance of the delay problems for global interconnects, whether repeated or not. As mentioned before, this is one the prime reason why we focus primarily on global interconnects in this thesis.

The graph also shows a lowering of the predicted delay for local wires, in line with Ho’s argument about local wires. The Dennard scaling rules predict no decrease of this delay (Table 2.1) and the reason that the actual delay is estimated to decrease is due to technology improvements, such as projected changes in interconnection dielectrics. Still, the delay of these local wires is predicted to decrease at a slower pace than the gate delay. That means that local interconnects still pose some issues. This is explained clearly in [10], where it is stated that a re-design of a circuit in a newer technology is no longer essentially just a matter of downsizing all dimensions: When all dimensions are downsized according to Dennard scaling, then some fraction of the interconnects, which had acceptable RC delays before scaling, will no longer satisfy the timing constraints of the scaled circuit, given that the operating frequency is also scaled. These local wires will have to be moved to higher metal layers to get larger diameters and decrease their RC delay. This is one of the reasons why a truly hierarchical wiring scheme as shown in Figure 2.1 has become a real necessity, not only to enable proper low-impedance power grids, but also for data communication. In fact, a good hierarchical interconnect stack with a layer count that increases over technology generations is part of the solution to the interconnect problem, as will be discussed in the next section.

2.6 Technological interconnect advances

To postpone the difficulties with interconnects and continue successful scaling (‘Moore’s law’), the semiconductor technology industry has devised a number of workarounds. A

number of these options have been used in the past, some are entering mainstream technologies and some are planned for the near or far future.

2.6.1 Implemented improvements

Aspect ratio increase

As mentioned earlier, one of the first techniques that was used to avoid a scaling-disparity between transistor speed and wire speed, was to raise the wire aspect ratio [7]. As the capacitance of an interconnect consists partly of fringe capacitances that do not scale with perimeter size, one can raise the height of interconnects and benefit from a resistance that initially decreases more rapidly than the capacitance increases. But, already in 1995, with an average aspect ratio that had risen from 0.4 to 1.3, Bohr [23] predicted that this option would soon reach its limits as the RC delay benefits from increasing aspect ratio diminish above ratios of about 2. Also, patterning and etching become more difficult at higher aspect ratios and intra-layer crosstalk is worsened. And indeed the latest ITRS [3] predicts little increase in future aspect ratios, which are currently ranging from 1.8 for local and intermediate wires to 2.3 for global ones.

Copper interconnects

Bohr [23] also mentioned the search for new conductor and dielectric materials, to meet 'future ULSI interconnect requirements'. Around 1998, the industry indeed shifted from the use of aluminum interconnects to copper interconnects [28], as copper has 40% lower resistivity. A remaining problem is the increase in copper resistivity at small dimensions (i.e. <100nm line-width) due to grain boundaries and interfaces [24, 29]. Continued research for e.g. other barrier materials that simultaneously suppress electromigration and provide a smoother interface might postpone the increase in resistivity, and more experimental options such as carbon nanotubes might become available in the future, but a true solution has not yet been found [3]. Fortunately, this is not (yet) a major problem for global interconnects, as these usually reside in the higher, larger metal layers.

Low-k dielectrics

Regarding the dielectric materials, a lot of research has also been carried out in the past decade, with the goal to replace (or mix) the traditional silicon oxide with other dielectrics to get a lower dielectric constant (the so-called low- ϵ , or low-k dielectrics) and less capacitance as a consequence. Some of the initial steps, such as the use of fluorine doped silicon dioxide ($\epsilon=3.7$) were quite successful and at present other reliable insulators with $\epsilon=2.7-3.0$ are used. However, further reduction of the dielectric constant with the use of porous materials was hampered by reliability and yield issues [29] and a reduction below $\epsilon=2$ is deemed extremely difficult [3]. An alternative that is considered is the use of air gaps to lower the dielectric constant [3].

2.6.2 Future improvements

Even when the industry succeeds to find better low-k dielectrics, material changes alone can not provide interconnect improvements forever. There are not many practical metals with a lower resistivity than (bulk) copper, and the relative dielectric constant has a natural

lower limit of one (vacuum). To still be able to improve interconnect performance, especially for global interconnects, a number of more radical changes have been proposed and are under active investigation. These include research to replace (global) electrical interconnects by RF/Wireless or optical interconnects or to move to 3D integration [3]. Each option is shortly discussed below.

Wireless data transmission

Although not strictly a technological advancement, wireless data transmission is mentioned in the technology roadmaps as a candidate for global on-chip communication [3, 27]. However, wireless (or in a more general term: unguided) data transmission [30], faces the problem that there will only be one communication channel available, at least as long as the wave-length is larger than the antenna-size. Directional beam-forming is only a real option when the antenna-size sufficiently exceeds the wave-length. With an antenna of for example 1mm in diameter, the RF frequency needs to be higher than $f > (c_0/\epsilon_r)/1\text{mm} \approx 150\text{GHz}$ before a somewhat directional beam becomes feasible. When we assume that such frequencies become feasible in the near future and we also assume that the link would be so wideband that e.g. 100Gb/s could be transmitted in the direction perpendicular to the 1mm wide antenna, then this option would still not be competitive with interconnects. With a pitch of e.g. $1\mu\text{m}$, thousand interconnects would fit in the same cross-section as the antenna. As will be shown in this thesis, each of these interconnects could easily transmit at data rates exceeding 1Gb/s, creating an aggregate data rate of more than 1Tb/s, ten times larger than the antenna.

So, wireless data transmission is no good alternative for interconnects in on-chip data communication. Other application areas where it can perhaps be beneficial in the future is in clock distribution and intra-chip communication [3].

Optical data transmission

Data transmission over optical interconnects might become a viable option in the far future, but still requires a huge number of technology advances [31]. This includes the implementation of dielectric materials with sufficiently different refractive indices to confine the beams into small guided channels (and avoid crosstalk). It also requires the integration of very high-speed optical transmitters (laser-diodes or light-modulators) and receivers (photo-diodes). These are challenging issues, as photo-diodes for example suffer from finite bandwidth problems, when integrated in standard CMOS [32]. Quantifying this finite bandwidth leads to a prognosis that optical interconnects could only compete with copper interconnects when wavelength-division-multiplexing (WDM) would be used [31]. WDM would complicate the technology integration issues for the optical elements even further.

3D integration

The stacking of multiple IC's on top of each other or the integration of multiple active Si layers in one IC is the basis of 3D integration [3]. Using the 3rd dimension more effectively can significantly reduce the footprint of an IC, thereby alleviating the problems with long interconnects. It is a whole research area on its own as it promises increased integration, but one of the problems that is faced is how to remove heat from the chip. The power increases with higher integration, while the surface area over which heat can be transferred decreases.

Still, it is believed to be a potential solution for the interconnects limits associated with ‘gigascale integration’ [3, 24]. When analyzing interconnects, 3D integration is however not a really radical change and can be regarded as being similar to the use of more interconnect layers as discussed next.

2.6.3 Reverse scaling

One of the most practical solutions seems to continue to increase the number of metal layers on a chip, in a true hierarchical fashion. At the bottom of the interconnect stack, smaller metal layers are added to interconnect the transistors over short distances. Because these local wires are very short, it should still take many generations before they become a real bandwidth bottleneck, even with the increasing resistivity of copper for small cross-sections [24]. Longer wires that do become a bandwidth bottleneck in scaled designs can move up in the interconnect hierarchy to thicker metal layers, with lower resistance per length. The consequence is that metal layers at the top of the stack have to become increasingly large in future processes, the so-called ‘reverse scaling’. So the metal stack increases at both sides with smaller layers at the bottom and thick layers at the top.

In [10], the required number of metal layers for future process generations is discussed in detail and estimations are given for three different scaling scenarios. Rents rule (see section 2.3.1) is used to estimate the required number of interconnects and their required metal layer (depending on their length) in each scenario. In the first scenario, the number of transistors doubles with each process generation while the perimeter of the chip stays constant. The resulting estimation for this scenario shows a quick explosion in the number of required metal layers, which have to increase by a factor of 1.7 per generation, much faster than the linear increase of about 0.5 metal layer per generation as predicted by the ITRS [3, 29]. Even when predicted material changes such as low-k dielectrics are taken into account, then the number of metal levels still needs to increase with a factor of 8 over 5 generations. This is clearly not a practical situation and reconfirms the difficulties with global interconnects that do not scale down in length.

The other two scenarios describe interconnects that scale down in length, either at a slower pace than the scaling of device size (proportional to the square root) or at the same pace as the scaling of devices. In this last situation, the number of metal layers still needs to increase to keep up with the increases in transistor speeds, but now at a manageable rate, reachable with the ITRS prediction of about 0.5 layers/generation.

This last situation implies either architectural changes to keep the interconnects ‘local’ (e.g. networks on a chip), or implies other solutions with the same effect such as 3D integration. Without these changes, a continued scaling of clock-speeds and integration densities simply does not seem feasible. On a positive note, when these changes are incorporated, then the more radical technology changes might not be necessary and interconnects do not necessarily have to be a showstopper for future CMOS technologies.

2.6.4 Combination with architectural and circuit improvements

To summarize the above results, continued scaling in CMOS is possible, but requires design strategies that are very ‘interconnect aware’. Any technique that can improve the

performance of interconnects should be embraced, as interconnect bandwidth will become scarce and a critical cost factor. So, next to the technology and architectural improvements, data communication improvements at the circuit-level can also be a very beneficial approach.

Circuit techniques can not only alleviate the interconnect bandwidth problem, but potentially also the power problem. The ITRS predicts for example that the average power per GHz per cm² per metallization layer will increase from the current 1.3W/GHz/cm²/layer to about 2W/GHz/cm²/layer in 2020 [3]. But as the number of metal layers also increases (with the mentioned 0.5/layer per process node), the total power consumption will increase even more, which adversely contributes to the already big problem of chip heat removal. Circuit techniques that reduce the power consumption for on-chip communication, as will be presented in this thesis, help to tackle this problem.

2.7 Interconnect dimensioning

From a circuit design perspective, given that we operate in standard CMOS and are operating with band-limited copper wires as interconnects, we can still try to optimize these wires for data communication. When global data transport is concerned, what is usually the most important factor is throughput (a.k.a. aggregate data rate, or sometimes also called ‘bandwidth’), or how to transport as many bits per second from A to B. Whether this data transport occurs with many bits in parallel or with all bits in series is often only of secondary importance. To maximize the data throughput over a link, it intuitively makes sense to use wide data paths [14] with many densely packed interconnects. It will be shown in this section that this intuitive notion is only partly true and that it is actually not advantageous to make the width and spacing of the wires smaller than their vertical dimensions.

2.7.1 Bandwidth per cross-sectional area optimization

To maximize the throughput for a certain bus, we will optimize the ‘bandwidth per cross-sectional area’ (BW/Area). A bus with these optimized interconnects will have the highest achievable throughput for a certain bus area.

The bandwidth of a single interconnect is inversely related to its RC time constant, and consequently depends on its dimensions. The length of the interconnect is determined by the application and we therefore do not include it in the optimization, but use the normalized R’C’ instead. We are free to choose the width (w) and spacing (s), as shown in Figure 2.7. When we assume that we have control over the technology (or perhaps less radical: control over the metal layer) then we can also choose the vertical dimensions (h and t).

In [2, 33] we discussed how these cross-sectional dimensions should be chosen to optimize the bandwidth per cross-sectional area (BW/Area). A first-order analysis predicts that the BW/Area peaks when all the wire and spacing dimensions (w , h , s and t in Figure 2.7) are about equal, which is illustrated with the equations below (neglecting fringe-capacitance):

$$C' = 2C'_{side} + 2C'_{topbottom} \rightarrow C' \propto \left(\frac{h}{s} + \frac{w}{t} \right), \quad R' \propto \frac{1}{wh} \quad (2.3)$$




Figure 2.7: Interconnect cross-sectional dimensions.

$$BW \propto \frac{1}{R'C'} \quad , \quad Area = (w+s)(h+t) \quad , \quad \frac{BW}{Area} \propto \frac{1}{R'C'} \frac{1}{(w+s)(h+t)} \quad (2.4)$$

$$\frac{BW}{Area} \propto \frac{wh}{\left(\frac{h}{s} + \frac{w}{t}\right)(w+s)(h+t)} = \frac{1}{\left(\frac{h+t}{s} + \frac{h+t}{w} + \frac{w+s}{h} + \frac{w+s}{t}\right)} \quad (2.5)$$

The partial derivatives of (2.5) are all zero if $w = h = s = t$. Usually, the h and t are fixed by the process and choice of metal layer (or at least h/t is), but w and s can be varied independently. Taking the partial derivative of (2.5) to w and s and solving it for zero yields:

$$w_{opt} = s_{opt} = \sqrt{\frac{h+t}{\frac{1}{h} + \frac{1}{t}}} = \sqrt{ht} \quad (2.6)$$

In most technologies h and t are not very different, which means that we can approach the real optimum - where all dimensions are equal - quite well.

However, second-order effects such as fringe capacitance; different dielectric constants for inter and intra-layer dielectrics; barrier layers; or the use of the top metal layer without any top-plate capacitance, all give an alteration of the optimum. Differential signaling also changes the optimum as the capacitance between the two differential halves is doubled as a result of the Miller effect (section 4.4). To include these effect and fine-tune the optimum dimensions for w and s , more elaborate calculations were carried out, in combination with EM-field simulations [2]. The analytical results above will perhaps not yield the most accurate value for the optimum, but they do provide a first-order estimate and aid to establish a few general conclusions.

As discussed earlier, most new technologies use a hierarchical wiring system with increasing wire thickness for higher metal layers. This is beneficial for the data rate per interconnect, as the use of a thicker metal layer with larger inter-layer dielectrics will give a lower resistance (2.3), (2.4). The data rate per cross-sectional area is however not changed because, with optimal dimensions ($w=h=s=t = d_{opt}$), the $BW/Area$ is independent on d_{opt} :

$$\left. \frac{BW}{Area} \right|_{w=h=s=t=d_{opt}} = \text{technology constant} \quad (2.7)$$

So the required data rate per single interconnect can determine the choice of metal layer, with little impact on aggregate data rate per cross-area.

The downside of this independence of the BW/Area on d_{opt} is that we are apparently not able to improve the throughput for a given length through a certain cross-area beyond a certain limit. For global buses that do not scale in length over technology, this means that the only way to increase the throughput is to increase their cross-sectional area, for example by adding metal layers.

2.7.2 Bandwidth per pitch optimization

Instead of focusing on the BW/area as criterion, we could also have optimized for the highest BW/Pitch ($pitch=w+s$), and not regard vertical size as a cost-factor. In fact, optimization of the BW/Pitch is more frequently used and discussed [8, 34-36]. Many of these papers discuss optimization of interconnects and repeaters simultaneously [34-36]. The reason that we started with an optimization of the BW/Area is because vertical dimensions can certainly be cost factors, both for design and technology. From a design-perspective, it is for example possible to leave a metal layer empty to reduce the capacitances of the layers around it, but the equations above predict that this is not beneficial for the total throughput. From a technology perspective, the BW/area equations predict that the thickness of the metal does not have to be much larger than the minimum allowable width and spacing, at least not for optimal throughput.

The optimization of the BW/pitch is actually not really different from the optimization of the BW/Area, at least not for the simple model as used in (2.3)-(2.5), which gives a similar equation as $BW/pitch=(h+t) \cdot BW/area$. The resulting optimum is the same, with $w = h = s = t = d_{opt}$. The optimum for w and s , given a certain h and t also is the same as (2.6).

An aspect that is different however, is the fact that the BW/pitch does increase when we increase d_{opt} :

$$\left. \frac{BW}{pitch} \right|_{w=h=s=t=d_{opt}} \propto d_{opt} \quad (2.8)$$

In other words, higher, larger metal layers have more BW/pitch than the smaller, lower layers, which gives a clear motivation for the reverse scaling of wires as discussed in section 2.6.




Figure 2.8: *BW/Area (a) and BW/pitch (b) as a function of vertical interconnect spacing and height (with $w=s=$ square root of ht).*

To illustrate the effect of metal layer spacing and height on the BW/pitch and the BW/Area, 3D plots are shown in Figure 2.8, as a function of the cross-sectional dimensions h and t and using (2.6) to define s and w .

2.7.3 Bandwidth optimization in general

So, from a BW/Area point of view there is no compelling reason to increase the sizes of metal layers, but from a BW/pitch perspective there is. One could thus argue that aggressive reverse scaling will be beneficial as reverse scaling of a metal layer increases its available BW/pitch. An alternative would be to stack multiple thin metal layers with many small parallel wires, which would have the same BW/area, but a lower BW/pitch and higher manufacturing costs (due to the additional masks and processing steps).

When the data that has to be transported is high-frequency and serial in nature, then the argument is straightforward and thick reversely scaled wires are clearly the wires of choice. But, it might well be that the source-data is present in a parallel form, as is often the case in large-scale digital circuits for example for register or memory data. Transporting such data over a few very thick wires to obtain the highest BW/pitch will require high-speed serializing and de-serializing (serdes) circuitry, which require power and area overhead. So in this case, it might be favorable to use many small wires for data transport. These many small wires will still fit in the same area as the few large wires as their BW/area is the same as for the reversely scaled wires. When multiple layers are used for the small wires, then the BW/pitch is also the same as for the large wire. This is illustrated in Figure 2.9.

An observation that can be made for both BW/pitch and BW/area bandwidth optimization is the fact that it is not beneficial for throughput to use wires with high aspect ratios. At first sight, this seems to be contradictory to the fact that high aspect ratios are very common in current CMOS technologies, as mentioned in section 2.6. High aspect ratios can be beneficial for other reasons, of which a few are discussed next. First, as mentioned in

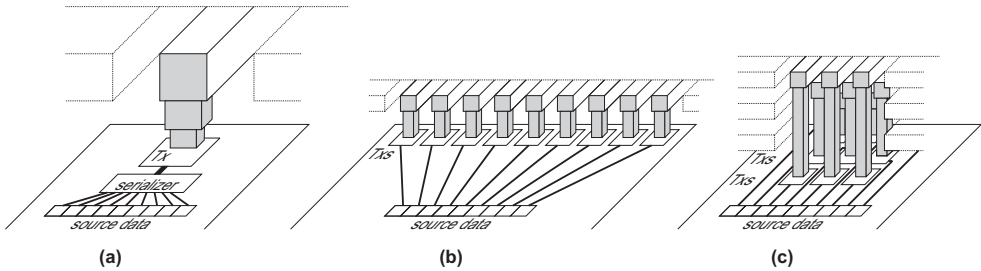


Figure 2.9: Bus configurations with (a) one large wire, (b) nine small wires in one metal layer, (c) nine small wires in three metal layers. All configurations have the same BW/area, with area defined as in Figure 2.7. (a) and (c) have the same BW/pitch.

section 2.6, historically it seemed favorable to increase the aspect ratio to enable a rapid porting of a design and avoid problems with timing-closure without the need for a complete re-design (or a new architecture). Second, the RC time constants of most interconnects on a chip are still small enough to not be a bandwidth bottleneck, but their RC time does create some delay. For these interconnects it can be favorable to increase the height of the interconnects and lower their delay, without a penalty of an increase in foot-print for the interconnect. Or argued in a different way: for these (local) interconnects, it is desirable to decrease their lateral dimensions to be able to keep up with the diminishing dimensions of the transistors and avoid routing congestion.

It does seem that these arguments are not convincing to keep a high aspect ratio for the intermediate and global metal layers, not in the least because it costs significant manufacturing effort to enable these high aspect ratio's [29]. It hence might be a good strategy to (slowly) migrate back to layers with lower aspect ratios (down to about unity), at least for those layers where the intermediate and global wires for data communication reside.

Other optimization criteria

In this section, we focused on throughput as criterion for wire optimization, but there exist of course other criteria. We could for example dimension the wires for minimal crosstalk or for minimal power consumption. However, for criteria such as minimum power, there does not really exist a practical optimum, as optimization would lead to an isolated interconnect with minimum dimensions and maximum spacings [2], to minimize the capacitance. Still, for some applications it might be good to deviate from the BW/area optimum in favor of other criteria. But in an IC environment where the bandwidth of the wires itself becomes an increasing concern, especially for long wires, optimization of the BW/area (throughput) is a solid choice for most applications.

2.8 Summary and conclusions

The list below shortly summarizes the results and conclusions from this chapter:

- Modern CMOS processes have a hierarchical wiring style. The bottom wire layers are for local interconnects and the top layers are mostly used for power and clock routing, with only limited room for data communication. In this thesis, interconnects in the intermediate metal layers are therefore used for communication analysis.

- According to Rent's rule, the number of data-wires that span a 10mm perimeter of a 10x10mm chip constitute less than 0.1% of the total wire count. Modern architectures such as NoCs can further reduce this percentage and break the wires into smaller segments (a task currently done by repeaters). But although low in numbers, wires in between 1mm and 10mm are very important for effective global on-chip communication.
- Technological material advancements such as copper interconnects and low-k (or air-gap) dielectrics give only a finite increase in interconnect speed. Other technology advancements such as reverse scaling are also needed, in combination with a reduction of interconnect lengths by using e.g. NoCs or 3D integration.
- Circuit techniques are also needed to reduce the power and increase the speeds over the interconnects, to avoid that interconnects become a dominant factor in the already complicated topic of on-chip heat removal. This only becomes more complicated with more metal layers (and even more so with 3D integration).
- The cross-sectional dimensions of an interconnect and the spacings to other interconnects should all be chosen roughly equal (assuming one uniform dielectric parameter) to optimize the interconnects for aggregate data rate.
- The aggregate data rate per cross-sectional area is in first order not dependent of the thickness of the metal layer, but the bandwidth for a single wire is. It is hence beneficial to continue to expand the interconnect stack with additional thick metal layers (reverse scaling). And for data communication, these thick layers do not need high aspect ratios.

Chapter 3

Interconnect characterization and modeling

3.1 Introduction

This chapter discusses how the interconnect properties that are important for on-chip communication can be characterized and modeled. In the thesis of Eisse Mensink [2], the interconnect characterization and parameter extraction is discussed in greater detail. In this chapter we will summarize some of that material and elaborate on how the parameters affect the interconnect behavior and which parameter is important under what conditions. The last part of this chapter discusses how the interconnect models can be adapted and simplified for behavioral modeling and for circuit simulations.

The next section defines in more detail which interconnects were investigated in this project. Section 3.3 discusses how their parameters were characterized. Section 3.4 discusses the transfer function of interconnects. Section 3.5 and 3.6 discuss the influence of inductance and skin-effect on this transfer, for which conclusions are drawn in section 3.7. Section 3.8 discusses interconnect models suitable for circuit design and for behavioral modeling. Section 3.9 gives a summary and conclusions.

3.2 Interconnects in this project

In this project, we focus on global or semi-global interconnects, which means that we are looking at the higher metal layers. We however assume that the topmost metal layer is completely reserved for power and clock routing (see the discussion in section 2.2), so the interconnects will be completely surrounded by other interconnects.




Figure 3.1: Interconnect configuration as assumed in this thesis.

The interconnect configuration is shown in Figure 3.1. A Manhattan routing style is common in on-chip interconnect routing, so the surrounding interconnects will usually be oriented perpendicular when a bus is located in one metal layer as in Figure 3.1.

To make the discussion more practical, we normally assume that such an interconnect has a length of 10mm, to represent a typical global interconnect and allow for easy comparison with prior work. But when theoretical findings are discussed, we will try to present them as general as possible and abstract from actual interconnect lengths. We also assume that the communication structure consists of buses with all signals traveling in the same direction. Routing of wires with signals traveling in opposite directions would greatly increase the problems with crosstalk and can be regarded as bad design-practice for high-speed data communication. We furthermore only regard point-to-point buses and no types of multi-drop buses as point-to-point connections are much more efficient in terms of link bandwidth and power, as will be discussed in section 3.8.2.

In a number of cases, we will also discuss shorter wires with lengths in the order of 1mm to 2mm. This is because the long wires can be broken up into smaller segments with repeaters in between. These repeaters can be simple inverters or perhaps even complete routers that are used for Network on Chip communication, as mentioned earlier in section 2.3. In essence, one could regard such a router as a more advanced type of repeater. In this project, the primary goal is in all cases how to go cross a global, chip-wide distance (10mm) as fast and as power efficient as possible, regardless of whether the interconnects are broken up into smaller sections or not.

We assume throughout this thesis that we are operating in standard CMOS processes, as CMOS processes are the dominant choice for large-scale digital integrated circuits, both at present and in the foreseeable future. During the project, we used CMOS processes from two contemporary technology nodes. The first is a 0.13 μm , 1.2V, 6-metal copper CMOS process and the second is a 90nm, 1.2V, 7-metal copper CMOS process (with two additional metal layers available as options).

For our demonstrator IC's and for most of our analysis, we used the intermediate metal layers for the data buses. M5 and M4 were used in the 0.13 μm process and M4 was used in the 90nm process. These metal layers satisfy the constraint of a reserved top layer, but they also have only a moderate thickness. For the width and spacing, we use dimensions optimized for highest BW/area, as was discussed in section 2.7. The resulting thin and narrow wires have a higher distributed resistance than the top metal layers. This gives, in combination with the length of 10mm, an interconnect bandwidth that is significantly lower

than the transistor frequency limits. In this way, the limited interconnect bandwidth problem manifests itself in full force. As technology progresses and transistors get faster, these results will also become representative for interconnects that are thicker (and wider) or shorter than the ones that were used in this project.

3.3 Interconnect parameter extraction

To be able to accurately model the behavior of on-chip interconnects, it is necessary to know their parameters, most notably the values for the distributed capacitance, resistance and to some degree also the inductance. There are a number of ways to acquire these parameters for a given interconnect configuration. The design-manual of a process usually gives some data about the metal-to-metal capacitances and sheet-resistances for all the metal layers, but especially the capacitance is broken up into many different (fringe and sheet) terms and it is quite configuration dependent. One could also use the general interconnect data as published in the roadmaps from the ITRS [3], but that data is not necessarily accurate for the specific process. To get analytical expressions for the distributed parameters as function of material properties and dimensions, one can also assume some simplifications in the configuration and use the Maxwell equations themselves [2].

In our project, we obtained material properties such as the sheet-resistance and dielectric permittivities from the design manual and used this as a basis to construct an accurate model in a 3D EM-field solver [2]. Simulations with this solver were used to analyze the behavior of the interconnects and extract s-parameters and distributed RLC parameters.

In the EM-field solver model, metal plates were used to approximate the surrounding perpendicular interconnects (assuming a Manhattan routing style), as a large-scale IC usually has a high wire density in all layers. In the actual demonstrator IC, ground- and Vdd-connected metal stripes were used to fill these surrounding metal layers. Simulations and measurements showed that the capacitance between the interconnect and a set of (Gnd or Vdd-connected) metal stripes with about 50% fill density is quite similar to the capacitance between an interconnect and a metal plate, justifying this approximation. The cross-sectional dimensions that are used for the interconnects in the EM-field solver and on the demonstrator IC's are optimized to get the highest bandwidth per area, as discussed in section 2.7 and in more detail in [2].

Table 3.1 on the next page shows the distributed parameters for these interconnects, as obtained with the field solver. The parameters are obtained for interconnects when used either single-ended or differentially. For differential interconnects, the parameters are slightly different due to the Miller-multiplication of the mutual capacitance and the cancellation of the mutual inductance. For the single-ended interconnects, the return paths for the current are the surrounding (Gnd- or Vdd-connected) plates. These plates (or in actual circuits: the power grid) have such low impedances that their influence can be neglected.

CMOS Technology node	0.13μm	90nm
Interconnect length	10mm	
Metal (copper) layer	M5	M4
Width, Spacing	0.4 μ m, 0.4 μ m	0.54 μ m, 0.32 μ m
R'	150 Ω /mm	135 Ω /mm
C'		
single-ended wire:	0.23pF/mm	0.24pF/mm
differential wire:	0.27pF/mm	0.28pF/mm
L'		
single-ended wire:	0.41nH/mm	0.35nH/mm
differential wire:	0.25nH/mm	0.24nH/mm

Table 3.1: Dimensions and distributed parameters for interconnects with optimized bandwidth per cross-sectional area. Source: [2].

Note that the extracted capacitance agrees quite well with capacitance estimations from the ITRS 2006, which predicted about 1.8-2pF/cm for interconnects in intermediate metal layers in 90nm CMOS technology [37]. The 20% difference with our extracted values can be attributed to the smaller interconnect size that is assumed by the ITRS. The capacitance does not scale much over technologies, as predicted by Dennard scaling (Table 2.1 on page 26), and observed by [8]. The 2006 ITRS also predicted only a very small decrease to 1.5-1.8pF/cm for the 20nm technology node. The 2010 ITRS is slightly more optimistic and predicts that the capacitance will drop a bit, to 1.3-1.6pF/cm in 2020 [3].

The obtained distributed parameters can be used in the well-known telegraphers equations, which have known solutions for the frequency-domain transfer function [2, 38]. However, to obtain solutions for practical interconnects, which are not infinitely long and have certain termination impedances, s-parameter models are more suitable. S-parameter equations are a useful tool to calculate transfer-functions in more complex situations and in the presence of reflections at the terminations of the interconnects, as is carried out and explained in [2]. The resulting equations can be used in numerical tools such as Matlab, to inspect the transfer functions for interconnects or use it for post-processing, for example to obtain the time-domain impulse or step-response through (numerically evaluated) Fourier transforms.

3.4 Interconnect transfer function

Figure 3.2 shows a numerically evaluated transfer function, using the RLC parameters from Table 3.1, in this case for a single-ended interconnect in the 0.13 μ m CMOS technology and with the assumption of an idealized transmitter with zero ohm impedance and infinite impedance as load. The transfer function shows three different regions, where region one and two are caused by the RC behavior and region three by the LC behavior and the skin-effect.

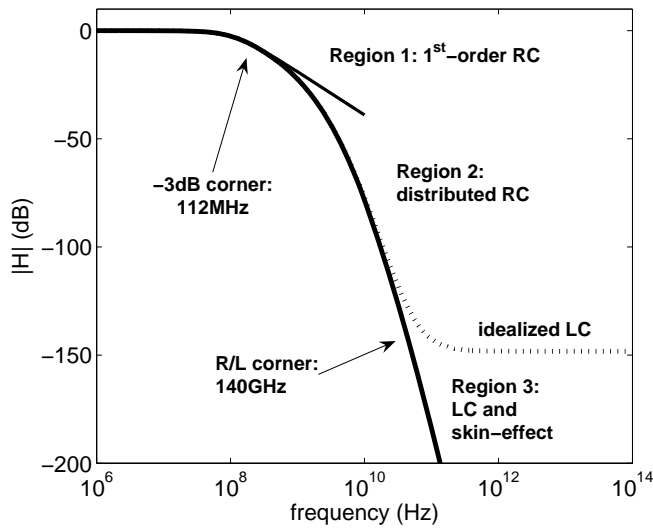


Figure 3.2: Transfer function for an interconnect-model with the parameters from Table 3.1, in this case for a 10mm single-ended interconnect in 0.13 μ m CMOS.

An interesting aspect of these RC-limited interconnects is that they have a single dominant pole, giving a high resemblance to a first-order roll-off in region one. This is an advantageous aspect that not only enables simple single-pole modeling, but also enables simple equalization circuits, as will be discussed in later sections. The dominant pole behavior is a nice characteristic of terminated RC-limited interconnects of finite length, and is not found in infinitely long RC-limited interconnects or in wireline channel models [39] as will be shortly discussed in section 4.2.1.

The higher-order part of the distributed RC-line transfer starts to dominate in region two. Only in the third region, for frequencies where $\omega L \geq R$, does the inductance begin to play a role. However, for these thin and long wires, this frequency region starts around 140GHz and is useless for data transmission as the attenuation is more than 148 dB. Furthermore, if skin-effect is taken into account, then there will not be a flat region in the transmission, as will be shown in section 3.6. But before skin-effect is analyzed, the next section first discusses in which cases the inductance of interconnects has to be taken into account.

3.5 Influence of inductance

The interconnect is RC-limited and the inductance can be neglected as long as the RC time constant ($\text{length}^2 R' C'$) is much larger than the L/R time constant (L'/R'). Only for short wires does this no longer hold, as can be seen in Figure 3.3 on the next page. The figure shows that at an interconnect length of 0.5mm, typical transmission-line effects start to appear, such as a rippling transfer function (or ringing in the time domain) due to impedance mismatch. Note that the driver impedance is set at zero Ohm, to show the idealized situation. In actual applications, short wires are likely to get small drivers, creating a relatively high driver impedance and dampening the transmission-line effects to some degree [40].




Figure 3.3: Idealized transfer functions of interconnects of different lengths (single-ended, $0.13\mu\text{m}$ CMOS), assuming a voltage source as driver ($Z_s=0$), no load ($Z_L=\infty$) and frequency-independent RLC parameters.

Because inductance can cause ringing and other detrimental effects, its presence is sometimes even seen as a disadvantage [24, 40]. The fact that inductance, in combination with capacitance, also creates propagation delay does not help to make it a favorable property. But this propagation delay, as defined by the LC product, is ultimately related to the speed of light in the medium, for which there are not many workarounds. So in those cases where inductance is clearly present one could simply terminate the interconnect with its characteristic impedance and count the blessings of very high bandwidths.

However, infinite bandwidth is not a realistic condition and an effect that is not yet accounted for in this section (nor in Figure 3.3) is the skin-effect, which also causes bandwidth limitations as will be discussed in section 3.6. For now, we concentrate on the more basic aspects of inductance and assume uniform conduction throughout the conductor

3.5.1 Influence of inductance on interconnect transfer

Given the RC and L/R time constants, it is possible to quantify at which lengths the inductance can start to play a role (assuming that driver impedances and load impedances do not mask the inductance effects). Because of the distributed nature of the R,C and L, we need to correct the time constants by a certain scaling factor (also see section 3.8). Based on the data from Table 3.1 and Figure 3.2 we can derive the following empirical correction factor for the RC time constant:

$$\tau_{RC} = \frac{1}{\omega_{-3dB}} \approx \frac{1}{2\pi \cdot 112 \cdot 10^6} = 0.41 \cdot R' C' l^2 \quad (3.1)$$

The L/R time constant should be independent of length, but as visible in Figure 3.3, the +3dB R/L corner has a little dependency on the length, also because a 3dB corner is only a

good measure for the actual time constant in first-order systems. More detailed inspection of the transfer function shows that the +3dB R/L corner varies between about 50GHz and 150GHz. By taking the average of 100GHz, we can derive the following correction factor:

$$\tau_{L/R} = \frac{1}{\omega_{+3dB}} \approx \frac{1}{2\pi \cdot 100 \cdot 10^9} = 0.58 \cdot \frac{L'}{R'} \quad (3.2)$$

The L/R time constant will start to get a significant influence (will start to cancel the RC time constant) when it approaches or exceeds the value of the RC time constant which happens for small lengths:

$$\tau_{RC} < \tau_{L/R} \rightarrow 0.41 \cdot R' C' l^2 < 0.58 \cdot \frac{L'}{R'} \rightarrow l < 1.2 \cdot \sqrt{\frac{L'}{C'}} \frac{1}{R'} \quad (3.3)$$

In our case, for the single-ended wire in the 0.13 μ m CMOS process, the length at which the two time constants are equal is only 0.3mm. Note that the square root of L/C normally represents the characteristic impedance Z_0 of lossless transmission lines [38, 41] so another way to formulate (3.3) is to say that the length should be so short that the characteristic impedance (multiplied by a correction factor that is almost unity) still dominates over the resistance ($R'l$). A similar statement is found in [42], but with a slightly different correction factor of 2 ($R'l < 2Z_0$), which is based on analytical models instead of the empirical factor of 1.2 that is found here. Other conditions given in [42] that should roughly determine whether transmission-line (inductive) effects are relevant are: 1) a load capacitance C_L that is smaller than the wire capacitance ($C'l$) and 2) a driver impedance that is smaller than 0.5-1 times the characteristic impedance Z_0 . But note that these conditions are very rough guidelines and are not always true. A driver impedance that exceeds the characteristic impedance can for example still give transmission-line effects when the transmission line itself has low losses.

3.5.2 Influence of inductance on propagation delay

Another parameter on which inductance has influence is on the delay of the interconnects. The square root of the product of L' and C' determines the propagation velocity and delay in lossless transmission lines [38, 41]:

$$v|_{R=0} = \frac{1}{\sqrt{L'C'}} = \frac{c}{\sqrt{\mu_r \epsilon_r}} \quad , \quad t_d|_{R=0} = \frac{l}{v} = \sqrt{L'C'} l^2 \quad (3.4)$$

The propagation velocity can not exceed the speed of light in vacuum ($c=300\text{Mm/s}$) and is in a medium with $\mu_r \approx 1$ and $\epsilon_r \approx 4$, such as silicon dioxide, limited to about half this speed, 150Mm/s. Note that when the velocity is calculated with the L' and C' values in Table 3.1 (on page 42), then it amounts to only 100Mm/s. This can (at least partly) be attributed to the fact that the tabulated L' is the inductance at low frequencies where current-densities are still uniform inside the interconnect. At very high frequencies, the skin-effect limits conduction to a thin outer shell of the interconnect, which gives a lower value for the effective inductance [43, 44], as is also discussed in the next sub-section. The inductance as characterized with the 3D EM models indeed starts to decrease in the multi-GHz range [2], while the capacitance changes much less.

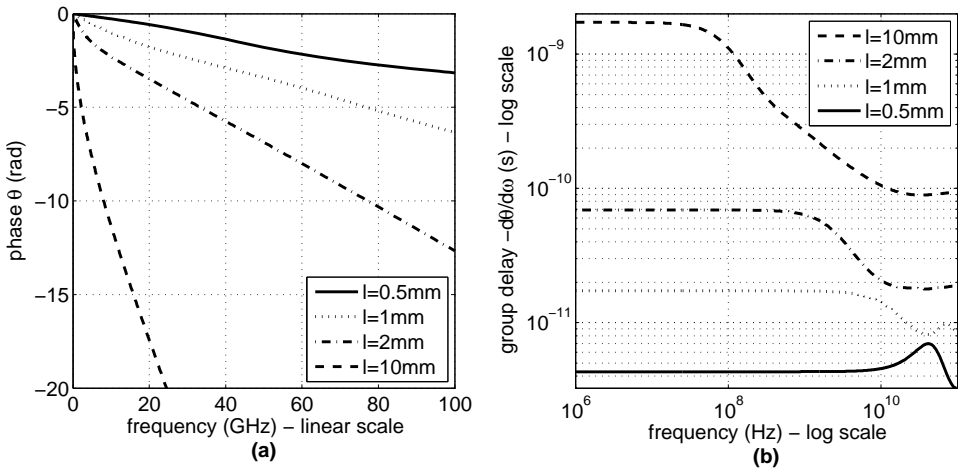


Figure 3.4: Phase shift (a) and group delay (b) for the same interconnect models as used in Figure 3.3.

The actual delay is usually limited by the RC delay, as is visible in Figure 3.4, where the phase and group delay ($=-d\theta/d\omega$) of the same wires as in Figure 3.3 are shown. For the 10mm interconnect the RC time constant dominates most clearly, creating a delay for the low-frequencies of over 1.7ns. In the frequency region where the LC regime takes over, the delay is only 0.1ns for the 10mm interconnect. As the line-length goes down, the RC delay decreases quadratically while the LC delay decreases only linearly. This is also visible in Figure 3.4b (with a few aberrations in the transition region for the smallest lengths). Below a certain length, roughly the same length as in (3.3), the RC delay becomes lower than the LC delay, meaning that the LC delay will also take over at low frequencies (as the speed of light dictates the lower bound on delay).

The fact that, for long wires, the LC delay is so much lower than the RC delay has prompted some researchers to investigate band-pass data transmission in the LC frequency region [20, 21]. But to get the attenuation in this region to suitably low levels, very large interconnects are required. In [21], a microstrip configuration is used with a 6 μm wide and 2 μm thick top-layer copper wire above a ground-plane and without any other metal in the immediate surroundings. In [20], a coplanar waveguide is used with a Ground-Signal-Ground configuration where each wire is 4 μm wide and 0.53 μm thick, with a spacing that is also 4 μm . The resulting interconnects might be fast by themselves, but they are very slow in terms of BW/area, as explained in section 2.7, and when compared to other solutions [2].

Of similar origin is the high importance that is sometimes attributed to inductance in the analysis of long on-chip interconnects [19, 24]. Those papers usually discuss very wide wires, sometimes without other metal in the immediate surroundings, which might not be very representative for typical data communication situations. When we would have metal in the surrounding metal layers, then neither a very large width nor a large spacing would change the dominance of the RC time constant for the long 10mm wires [2].

The papers that attribute a high importance to the inductance in future technologies make a stronger case when we take reverse scaling into account, as discussed in section 2.6.3. Wires will become thicker with reverse scaling, so the R' will drop, creating higher RC bandwidths and lower R'/L' corner frequencies. However, as will be shown below, for these larger wires, the skin-effect will have to be taken into account, which will still limit the bandwidth to finite values.

To illustrate when the inductance (and skin-effect) becomes important for reversely scaled global wires, the following first-order estimation can be used. When both the horizontal and vertical dimensions of the interconnects from Table 3.1 are enlarged by a factor s , then the resistance decreases by a factor of s^2 , while the L' and C' should stay roughly constant. This means that the ratio of the τ_{RC} over the $\tau_{L/R}$ decreases by s^4 . For the 10mm interconnect from Figure 3.2 with its ratio of $\tau_{RC}/\tau_{L/R} \approx 1500$, we have to increase the dimensions by a factor of $s=6.2$ to let the τ_{RC} equal the $\tau_{L/R}$ (at an equivalent frequency of $f=1/2\pi\tau=4\text{GHz}$). This will give an interconnect width and a height of $6.2 \cdot 0.4\mu\text{m}=2.5\mu\text{m}$.

3.6 The skin-effect

As mentioned, the aspect that is left out in the analysis above is the skin-effect [19, 38, 43, 45, 46]. Skin-effect used to only be a concern for off-chip communication, where much larger conductor cross-sections are used. But with a skin-depth in copper of only $2.1\mu\text{m}$ at 1GHz [38], it will become a limiting factor for the high-frequency performance of the thick on-chip wires in reverse-scaled metal stacks.

In a sense, skin-effect can be attributed to inductance [46], or in a more general sense, to electrical flow (as opposed to electrical potential). A change in current in one (part of the) wire induces a voltage difference in another (part of the) wire and this induced voltage will ‘try to counteract its origin’. This means that alternating currents flowing in the same direction are repelled from each other while currents in opposite direction are attracted to each other. The repulsion effect forces conduction at high-frequency to the outer shell of a conductor. In a bundle of interconnects, conduction is forced to the outer conductors [45] when they transport current in the same direction. In that last respect, the skin-effect is similar to inductive crosstalk and its effect can be reduced by intertwining the bundle of interconnects with return paths that transport current in the opposite direction (see section 4.3).

Due to the skin-effect, the effective resistance of an interconnect increases proportional to the root of the frequency above the corner-frequency f_s where the skin-depth becomes equal to the size of the interconnect. For a round conductor with radius a , conductance σ and magnetic permeability μ , this corner frequency is given by [41, 44]:

$$f_s = \frac{1}{\pi\mu\sigma r^2} = \frac{R'_{DC}}{\mu} \rightarrow \omega_s = \frac{2}{\mu\sigma r^2} = \frac{2\pi}{\mu} R'_{DC} \quad (3.5)$$

As the resistance at low-frequencies (R'_{DC}) is also related to the radius squared, the corner-frequency is proportional to the DC resistance. Above the corner-frequency, the effective resistance increases due to the skin-effect R_{AC} :

$$R'_{AC} = \lambda\sqrt{\omega} \quad (3.6)$$

Skin-effect thus creates bandwidth limitations. In fact, the skin-effect behavior can be described by the diffusion equation [43, 46], which is the same equation as the one that describes the behavior of an RC-wire of infinite length, as will be discussed in more detail in section 4.2.1. So, for thick interconnects, the original RC bandwidth limitation is replaced by a skin-effect bandwidth limitation which actually has the same type of attenuation curve (at least in the ‘distributed RC region’).

For off-chip wires, there exist well-known models that describe the impact of skin-effect on the transfer function quantitatively (as discussed in [44]). With some adaptations, these equations can also be used for on-chip wires, to quantitatively support the claims made above.

3.6.1 Influence of skin-effect on the transfer function

As a starting-point, note that the transfer-function of a characteristically terminated transmission-line is given by the following well known solution to the ‘telegraphers equation’ [2, 38, 41, 44]:

$$H(j\omega) = e^{-\gamma l} \quad (3.7)$$

With l being the length of the transmission-line and γ is expressed in terms of the four general transmission-line parameters R' , C' , L' , G' :

$$\gamma = \sqrt{(R' + j\omega L')(G' + j\omega C')} \quad (3.8)$$

When we assume that the transmission-line is RC-dominated, that the R and C parameters are constant and that we can neglect the L' and G' parameter then (3.7) simplifies to:

$$H(j\omega) = e^{-\sqrt{j\omega R' C'} l} = e^{-\left(\frac{1}{2}\sqrt{2} + j\frac{1}{2}\sqrt{2}\right)\sqrt{\omega R' C'} l} \quad (3.9)$$

When we look back to Figure 3.2, this characteristic equation describes the interconnect transfer-function in region two. The transfer in region one is slightly different because in region one, reflections due to non-characteristic termination have a (positive) influence on the transfer, as discussed in [2] and in section 4.2.

However, when we do incorporate the inductance L' into the equations and incorporate the skin-effect, then the equations become a bit more complex. To model the frequency-dependence of the parameters, we assume that the on-chip interconnect resembles a coaxial cable configuration, such that we can reuse the standard equations. This assumption is less strange than it seems, especially for the interconnects in a configuration as shown in Figure 3.1, which ideally has equal dimensions in both directions and also equal spacing in both directions to the surrounding metal. It is not a totally round configuration and there are gaps in the surrounding metal, unlike the shield in a coaxial cable, but these are tolerable modeling inaccuracies as we are trying to compute the trends and not a very exact transfer. What is most important is that we have to assume that the surrounding metal acts as a return-path, similarly to the shield in a coaxial cable, which is true when the surrounding wires are connected to low-impedance sources.

For coaxial cables, the parameter λ that describes the magnitude of R'_{AC} in (3.6) is given by [44]:

$$\lambda_{\text{coax}} = \frac{1}{2\pi} \left(\frac{1}{a} + \frac{1}{b} \right) \sqrt{\frac{\mu}{2\sigma}} \quad (3.10)$$

With a is the radius of the inner conductor and b is the distance from the center to the shield. For on-chip wires, we can approximate these dimensions with $a = 1/2 d_{\text{opt}}$ and $b = 1/2 d_{\text{opt}}$ when we assume that the wires are sized for optimal BW/area, with all dimensions and spacings equal to d_{opt} , as was discussed in section 2.7:

$$\lambda \approx \frac{1}{2\pi} \left(\frac{2}{d_{\text{opt}}} + \frac{2}{3d_{\text{opt}}} \right) \sqrt{\frac{\mu}{2\sigma}} = \frac{4}{3\pi} \frac{1}{d_{\text{opt}}} \sqrt{\frac{\mu}{2\sigma}} \quad (3.11)$$

This parameter λ can also be expressed in terms of the DC resistance. If we assume $R'_{\text{dc}} = 1/(\pi r^2 \sigma)$, as in equation (3.5), and substitute $r = 1/2 d_{\text{opt}}$ then λ can be rewritten as:

$$\lambda \approx \sqrt{\frac{4}{18} \frac{\mu}{\pi} \frac{4}{\pi d_{\text{opt}}^2 \sigma}} = \sqrt{\frac{2}{9} \frac{\mu}{\pi} R'_{\text{dc}}} \quad (3.12)$$

Apart from its role in the definition of the AC resistance in (3.6), λ also returns in the definition of the inductance. The inductance consists of two components, an internal component (L_i) and an external one (L_e):

$$L' = L'_i + L'_e \quad (3.13)$$

The internal component defines the relation between the flux inside the conductor and the current and is frequency dependent above the skin-effect corner-frequency, just as the resistance [44]:

$$L'_{i,AC} = \frac{\lambda}{\sqrt{\omega}} \quad (3.14)$$

The external component is frequency independent as long as the dielectric is frequency independent:

$$L'_{e,\text{coax}} = \frac{\mu}{2\pi} \ln\left(\frac{b}{a}\right) \rightarrow L'_e \approx \frac{\mu}{2\pi} \ln(3) \quad (3.15)$$

As with the external inductance, the capacitance of an interconnect is relatively frequency independent, at least in the frequency region where the dielectric permittivity ϵ is frequency independent:

$$C'_{\text{coax}} = \frac{2\pi\epsilon}{\ln\left(\frac{b}{a}\right)} \rightarrow C' \approx \frac{2\pi\epsilon}{\ln(3)} \quad (3.16)$$

The fourth transmission-line parameter is the shunt conductance G' , but this parameter can be disregarded as was mentioned in section 2.4 because on-chip insulators such as silicon-dioxide have low losses [19]. So, with the assumption that G' is zero, equation (3.8) can be rewritten to:

$$\gamma = \sqrt{j\omega R' C' - \omega^2 L' C'} \quad (3.17)$$

When the frequency dependence from (3.6), (3.13) and (3.14) are substituted into (3.17) then γ becomes (for frequencies larger than the skin-effect corner-frequency f_s):

$$\gamma_{ac} = \sqrt{j\omega^{\frac{3}{2}}\lambda C' - \omega^{\frac{3}{2}}\lambda C' - \omega^2 L'_e C'} \quad (3.18)$$

Which can be rewritten as:

$$\gamma_{ac} = j\omega \sqrt{(1-j)\lambda \frac{C'}{\sqrt{\omega}} + L'_e C'} = j\omega \sqrt{L'_e C' \left(\frac{(1-j)\lambda}{L'_e \sqrt{\omega}} + 1 \right)} \quad (3.19)$$

As discussed in [44], this equation can be approximated by using the first terms of the Taylor approximation:

$$\text{For small } x: \sqrt{x+1} \approx \frac{x}{2} + 1 \rightarrow \quad (3.20)$$

$$\text{with } x = \frac{(1-j)\lambda}{L'_e \sqrt{\omega}}, \quad \gamma_{ac} \approx j\omega \sqrt{L'_e C' \left(\frac{(1-j)\lambda}{2L'_e \sqrt{\omega}} + 1 \right)} \Leftrightarrow \quad (3.21)$$

$$\gamma_{ac} \approx j\omega \sqrt{L'_e C'} + (j+1)\sqrt{\omega} \frac{\lambda}{2} \sqrt{\frac{C'}{L'_e}} \quad (3.22)$$

The first term in (3.22) is the familiar wave propagation delay. The second term in the equation is caused by the skin-effect and consists of two components: an imaginary part (γ_{skin-i}) that creates additional delay and a real part (γ_{skin-r}) that creates attenuation in the transfer function from (3.7). This real part is examined further, by substituting (3.12) and (3.15) into the equation:

$$\gamma_{skin-r} \approx \sqrt{\omega} \frac{\lambda}{2} \sqrt{\frac{C'}{L'_e}} = \sqrt{\omega} \frac{1}{2} \sqrt{\frac{2}{9} \frac{\mu}{\pi} R'_{dc}} \sqrt{\frac{2\pi C'}{\mu \ln(3)}} \quad (3.23)$$

Simplifying this equation gives:

$$\gamma_{skin-r} \approx \frac{1}{3\sqrt{\ln(3)}} \sqrt{\omega R'_{dc} C'} \quad (3.24)$$

So the loss term can be described in terms of the DC resistance and capacitance. This loss term is in fact the same as the RC-dominated loss in (3.9), except for the scaling factor that is now $1/(3\sqrt{\ln(3)})$ instead of $\frac{1}{2}\sqrt{2}$, which is about a factor 2.3 lower. Note that the same holds for the imaginary part (γ_{skin-i}) so the phase-shift due to the skin-effect is also very similar to (3.9).

Note that these approximations for γ_{skin} are only valid when x from (3.20) is much smaller than one. With this condition for x , with x as defined in (3.21) and with the substitution of (3.12) and (3.15), this results in the following condition for the parameters:

$$\frac{\lambda}{2L_e'\sqrt{\omega}} \ll 1 \rightarrow \frac{\sqrt{\frac{2}{9}\frac{\mu}{\pi}R_{dc}'}}{\sqrt{4L_e'L_e'\omega}} \ll 1 \rightarrow \frac{\sqrt{\frac{2}{9}\frac{\mu}{\pi}R_{dc}'}}{\sqrt{\frac{4\mu}{2\pi}\ln(3)L_e'\omega}} \ll 1 \rightarrow \frac{1}{\sqrt{9\ln(3)}}\sqrt{\frac{R_{dc}'}{L_e'\omega}} \ll 1 \rightarrow \quad (3.25)$$

$$R_{dc}' \ll 9\ln(3)L_e'\omega \quad (3.26)$$

So the skin-effect equations hold for that frequency region where the inductive impedance dominates over the resistive impedance. In other words: the skin-effect formula's become valid for the frequency region where $\omega \gg 1/\tau_{skin}$ with τ_{skin} being approximately:

$$\tau_{skin} \approx 9\ln(3)\frac{L_e'}{R_{dc}'} \quad (3.27)$$

We can relate this time constant to the $\tau_{L/R}$ that was defined earlier in (3.2), by assuming that the external inductance L_e' is about 1.5^2 smaller than the inductance at DC, to account for the factor 1.5 difference in propagation velocity that was discussed in section 3.5:

$$L_e' \approx \frac{1}{1.5^2}L_{dc}' \rightarrow \tau_{skin} \approx 4.4\frac{L_{dc}'}{R_{dc}'} \quad (3.28)$$

So this skin-effect time constant is larger than the transmission-line time constant $\tau_{L/R}$ from (3.2), which means that the transmission-line region where the transfer function would have become flat when there would be no skin-effect (region 3 in Figure 3.2) is overshadowed by the skin-effect behavior.

The region where $\omega > 1/\tau_{RC}$ but not $\omega > 1/\tau_{skin}$ is a transition region where the skin-effect gradually takes over the role from the DC resistance (unless $\tau_{skin} > \tau_{RC}$, then the skin-effect would be the dominant effect for all frequencies). This transition region is difficult to describe analytically (if at all possible) and is not covered by the classical skin-effect models for R' and L_i' from (3.6) and (3.14). To still be able to evaluate transfer functions, we defined the following simple transitional behavior models:

$$R' \approx \sqrt{R_{dc}'^2 + R_{ac}'^2} \quad , \quad L_i' \approx \frac{1}{\sqrt{\left(\frac{1}{L_{i-dc}'}\right)^2 + \left(\frac{1}{L_{i-ac}'}\right)^2}} \quad (3.29)$$




Figure 3.5: Transfer functions for the same interconnects as used in Figure 3.3 but now including skin-effect models.

These models for R' and L' were used to draw the transfer-functions shown in Figure 3.5, with interconnects with the same dimensions as in Figure 3.3 but now including the skin-effect formulas. The figure shows a few interesting aspects of the equations.

First of all, the figure confirms that the skin-effect indeed limits the bandwidth, also for the shorter wires. There is no longer a flat part in the transfer, as was present in the idealized situation of Figure 3.3. The transition-region that was mentioned above is also visible in the figure. Because of the earlier mentioned difference in scaling factor between the loss from distributed RC from (3.9) and the skin-effect loss from (3.24), the transfer is not a continuous smooth line, but shows different regions, especially visible for $l=1$ and $l=2$ mm. The scaling factor for the skin-effect loss in (3.24) does have a big tolerance, as it was the result of an approximation of the interconnect by a coaxial cable. Experiments showed that when the actual scaling factor would be a factor three larger, then the transition-region between RC and skin-effect losses would become less distinct and the transfers in Figure 3.5 would become a smooth line.

What is also visible in the figure, especially for $l=0.5$ mm is that termination is important in the region where skin-effect takes over from RC-behavior. Experiments showed that, when instead of $Z_1=0$, Z_1 is chosen equal to $Z_0 \approx \sqrt{L_0/C}$, then the ringing in the transfer almost vanishes (not completely, as Z_0 is not perfectly constant).

3.7 Conclusions on inductance and skin-effect

The previous two sections discussed the effect of inductance on the transfer, first with a one-dimensional model for the wire (constant inductance) and then refining it with the influence of skin-effect. The quantitative skin-effect derivations in the section above do contain some approximations, and the model might be further improved with a better than coaxial approximation of the on-chip wire (incorporating e.g. experimental data, such as

those found in [19] although there, other wire configurations are used). However, the general trends sketched in the equations appear to be plausible, based on which the following conclusions can be drawn:

- A low-pass behavior is always present in an interconnect transfer function, even when the τ_{RC} is smaller than $\tau_{L/R}$. When $\tau_{RC} < \tau_{L/R}$, then the inductance overshadows the DC resistance and conventional transfer models without skin-effect (as those used in section 3.5.1) predict infinite bandwidth. However, the skin-effect will still ensure a low-pass behavior, with the same shape in the transfer as a distributed RC-line model and the losses can still be expressed proportionally to R_{dc} and C' . This implies that the intrinsic bandwidth of a wire is always proportional to the cross-section squared and inversely proportional to the length squared.
- Although the simplified treatment of inductance in section 3.5 thus not correctly predicts the behavior of the transfer, it does still predict correctly whether the inductance has to be taken into account, depending on τ_{RC} versus $\tau_{L/R}$. This can be used for example to determine the ideal termination impedance and the delay. For small and long wires, only the $R_{dc}C$ is important, but for short or thick wires the inductance creates a characteristic impedance and unwanted reflections can occur if the termination is not matched with this impedance.
- The linear relation between length and propagation delay at high frequencies, as discussed in section 3.5.2 is still largely correct. This is because the frequency-dependent impact of skin-effect on the group-delay is much smaller than the wave-propagation delay itself: the $j\omega\sqrt{L_e C}$ term in equation (3.22) dominates at high frequencies. Do note that the predicted value of the propagation delay based on $L_e C$ is less than when $L_{dc} C$ would be used, as was already discussed in section 3.5.2. On the side, note that the simplified model from section 3.5 gave some dips in the group-delay around the $\tau_{L/R}$ corner as were visible in Figure 3.4. With the skin-effect models, the group-delay behavior is smoother and lowers asymptotically to the limit set by the $L_e C$ propagation-velocity. This observation strengthens the confidence in the skin-effect model.
- The derivations above restate what has been common knowledge in the microwave community, namely that guided transmission via transverse-electromagnetic (TEM) waves is inherently band-limited as high-frequency TEM-waves are attenuated by skin-effect. To obtain really high bandwidths while maintaining small cross-sections, one has to switch to other modes of propagation such as TM or TE waves [38]. These modes can even exist without any conducting material (optical transmission), eliminating conduction losses altogether.

For the wires in this thesis, which have lengths larger than 1 mm with small cross-sections (optimized for BW/area), RC line models that disregard the subtle effects from the inductance and skin-depth are still adequate. This is also reflected in the results that are obtained with for example lumped-element RC models (100 lumps), which are nearly indistinguishable from the EM-field solver results for those wires that were used as actual test-cases. The results from these last sections are more important for future, reversely scaled, interconnects with larger cross sections or interconnects with shorter lengths.

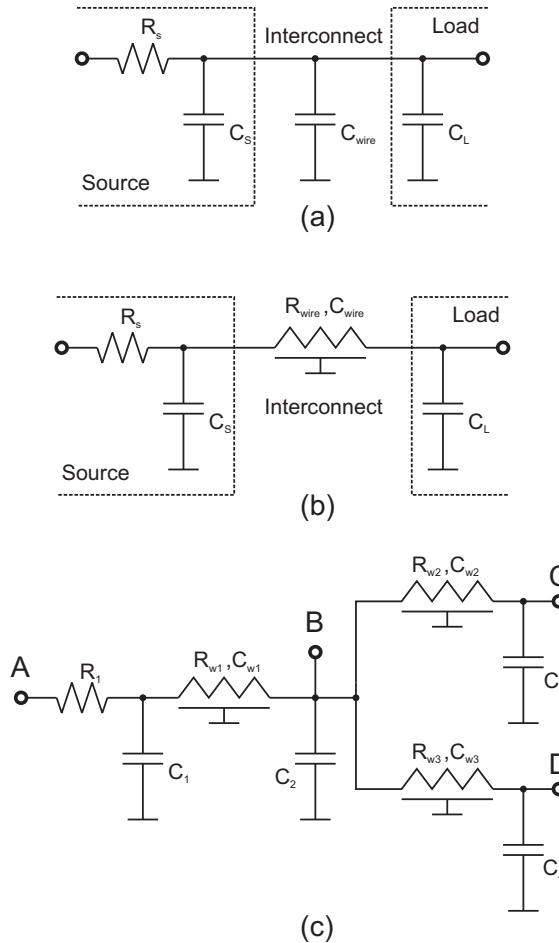


Figure 3.6: *Different levels of interconnect models: (a) simple model ignoring wire resistance, (b) distributed model, (c) model of branched interconnects.*

3.8 Interconnect modeling for circuit design

The characterization of the interconnect and the models that were discussed so far are accurate, but also quite complex and difficult to apply in a circuit design environment. The models are also a bit limited in their use in the sense that it is quite difficult (or cumbersome) to extend them to e.g. branched interconnects such as shown in Figure 3.6. This section therefore discusses a number of models that capture the various aspects of the behavior of the interconnect that are of interest for circuit design and enable simpler interconnect analysis. The section starts of with the simplest classical models and ends with the more complex lumped models.

3.8.1 Classical delay models

Traditionally, IC-designers were mainly interested in the capacitance of the interconnect, because the capacitance determines the power that is consumed ($\sim CV^2$ per switching action) and also determines how large the driver should be to get a suitably low driving resistance and enable slow enough delay (fast enough rise and fall-times). Such an interconnect model is shown in Figure 3.6a. The total capacitance of the circuit is $C_t = C_s + C_{\text{wire}} + C_l$. Together with the driving resistance this creates a time constant $R_s C_t$ with a step response:

$$V(t)_{t>0} = \hat{V}_{\text{drive}} \left(1 - e^{-\frac{t}{R_s C_t}} \right) \quad (3.30)$$

To quantify the speed of such a transceiver, we could use several different measures, for example the time it takes to go from 0 to 90% of the final value, or the 10%-90% time (rise-time), or the time it takes to reach 50% of the final value [7]. This last figure is usually taken as the measure of delay. For a first-order transfer it is easily solvable:

$$\frac{V(t_{\text{delay}})}{\hat{V}_{\text{drive}}} = 0.5 \rightarrow t_{\text{delay}} = \ln(2) \cdot R_s C_t \approx 0.7 \cdot R_s C_t \quad (3.31)$$

This measure for delay is just an approximation, as it for example assumes that the driving voltage is a step with infinite rise-time. It is still a very usable approximation, also in more realistic situations where the driving voltage V_{drive} will not switch instantaneously. In an inverter for example, the rise- and fall-times of the drain currents are much faster than the rise- and fall-times of the output, and V_{drive} can be approximated to be a step, provided that the rise-time of the input signal of the inverter is sufficiently high. The effect of finite rise-times can also be accounted for in the value for the effective driver resistance R_s , or capacitance C_s , as these are already approximations of the actual non-linear behavior of an inverter

The step-approximation is often used to analyze the delay of drivers and interconnects [7]. It also makes the analysis of cascaded stages much simpler (such as optimal repeater insertion, as discussed later in section 8.5, as the delay of one driver plus interconnect stage is assumed independent of the waveform of the previous stage. In those cases where the approximation is not accurate enough, one can use more complex models that take finite rise-times into account [47-50].

3.8.2 Elmore delay model

In the last two decades, as transistors became faster and interconnects got smaller cross-sections, it proved necessary to also take the resistance of the interconnect itself into account for accurate modeling of the delay as shown in Figure 3.6b. As the interconnect resistance is distributed over the interconnect, it can not simply be taken as part of the driver resistance, so the model from the previous sub-section is insufficient. This is especially the case when the interconnect is not simply a single wire but also contains branches, as visible in Figure 3.6c.

Fortunately, there exist a simple method to approximate the delay of linear networks, originally developed by W.C. Elmore over fifty years ago [51] and later applied to RC trees

with both lumped as well as distributed RC's [52, 53]. The Elmore delay model approximates the (50%) delay time to a step response and is valid for responses that are monotonic. For RC trees, the Elmore delay can be computed by a summation of the partial RC products found at each node k when traversing through the network from the input-node to the desired output-node. A partial RC product is found by multiplying the resistance between node k and the next node with the total load capacitance after node k . When the node in question is followed by a distributed R and C then that distributed C is weighted with a factor half [7, 53]. The equation below shows how this delay is calculated for the distributed driver-interconnect system from Figure 3.6b:

$$T_{Elmore} = R_s(C_S + C_{wire} + C_L) + R_{wire}\left(\frac{1}{2}C_{wire} + C_L\right) \quad (3.32)$$

For the branched interconnect tree from Figure 3.6c, the Elmore delay from A to B is [7]:

$$T_{AB} = R_1(C_1 + C_{w1} + C_2 + C_{w2} + C_3 + C_{w3} + C_4) + R_{w1}\left(\frac{1}{2}C_{w1} + C_2 + C_{w2} + C_3 + C_{w3} + C_4\right) \quad (3.33)$$

And the Elmore delay from B to C is:

$$T_{BC} = R_{w2}\left(\frac{1}{2}C_{w2} + C_4\right) \quad (3.34)$$

The total Elmore delay from A to C is simply the sum of the two: $T_{AC} = T_{AB} + T_{BC}$.

It should be noted that the Elmore delay is not a direct measure of the 50% crossing delay, but equals the first moment of the derivative of the step response [47, 53]. In other words, it represents the mean of the impulse response when the impulse response is regarded as a probability density function. In many cases the Elmore delay accurately matches the actual delay, but sometimes it is a bit off. As an elaboration to the Elmore delay, upper and lower bounds for the delay are given in [52, 53]. In [53] it is also shown that the Elmore delay is actually a good approximation for the dominant time constant of the network. In its capacity as time constant, the Elmore delay is widely used for interconnect analysis and analysis of repeater insertion [7, 54], with $0.7 \cdot T_{elmore}$ representing the 50% delay, as in (3.31).

3.8.3 Multi-drop buses and their Elmore delay

An interesting use of the Elmore delay model is to provide a simple quantitative motivation as to why data channels with more than one Tx and Rx, such as multi-drop buses [41] are not very efficient in terms of data rate, nor are they efficient in terms of power. Every transmitter in a multi-drop bus sees the capacitance of the entire interconnect with all its branches as load, which not only results in high energy consumption per data transition, but also affects the delay.

The delay and power penalty of a multi-drop bus is most severe when the target of a transmission is in the vicinity of the source, as the capacitance of all the 'unused' parts of the interconnect still add to the delay and power. As an example, assume that a bus connects four equal (tristate) transceivers as in Figure 3.6c and each interconnect has equal length. We can use equation (3.33) to calculate the Elmore delay for communication from transceiver A to B, assuming that each transceiver has the same capacitance that consists of both driver and receiver capacitance ($C_{TransC} = C_S + C_L$):

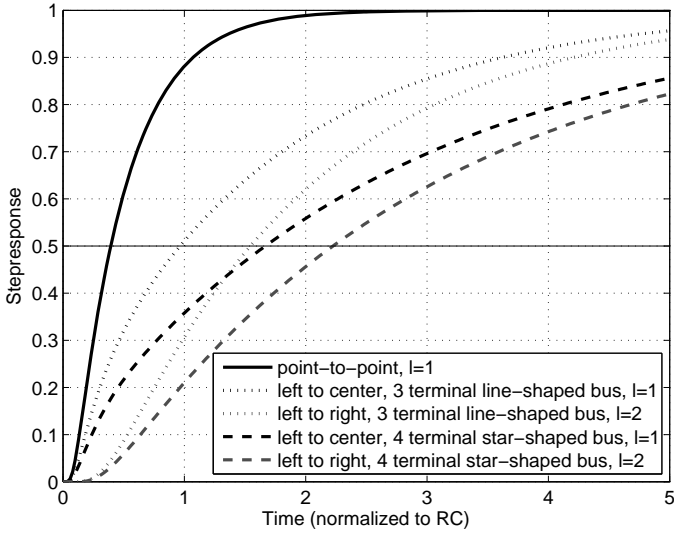


Figure 3.7: Step responses for different bus topologies, normalized ($R_{\text{wire}}C_{\text{wire}}/l^2 = 1$) and with ideal transceivers ($C_S=C_L=0$).

$$T_{AB} = R_S(4C_{\text{TransC}} + 3C_{\text{wire}}) + R_{\text{wire}}(3C_{\text{TransC}} + 2\frac{1}{2}C_{\text{wire}}) \quad (3.35)$$

This is much more than the delay of a single point-to-point interconnect expressed in (3.32), not in the least because the $R_{\text{wire}}C_{\text{wire}}$ term becomes five times as large. The terms caused by the source and load capacitances also increase by more than a factor four.

In a more general sense, when we assume that a bus connects N transceivers and we assume for simplicity that the $R_{\text{wire}}C_{\text{wire}}$ term is dominant, then we can easily compare the delay to the delay for a point to point connection. Assume for simplicity a line-shaped bus that uses one path with $N-1$ wire segments of equal length for the interconnections. The delay to go from one end of the bus to the other equals the delay of a point-to-point connection of the same length as we take only the $R_{\text{wire}}C_{\text{wire}}$ term into account. But, the Elmore delay between two neighboring transceivers will be $R_{\text{wire}}(\frac{1}{2}C_{\text{wire}} + (N-2) \cdot C_{\text{wire}})$, which is $2N-3$ times as large as when a point-to-point connection with $T_{\text{Elmore}} = \frac{1}{2}R_{\text{wire}}C_{\text{wire}}$ would be used.

In other bus configurations such as a star-shaped bus, the equations will be slightly different but the general trend is the same: a bus increases the power and delay for the interconnects. That is why we focus in this project on point to point connections, where all the interconnect capacitance is actually used to reach the desired destination. This is in line with the general architectural trend to move away from buses and towards (point-to-point) network-based communication, as was also mentioned in 2.3.2

In Figure 3.7, simulated step responses are shown for a few different interconnect configurations, to get a more visual idea of the different responses of different bus topologies. Two factors influence the speed of the response in this figure: First of course the length of the interconnect, which influences the time-scale but not the shape of the response. Second, whether or not the output node is loaded with other interconnects (that

Bus type	length	t_d	t_d/length^2	T_{Elmore}	$0.7 \cdot T_{\text{Elmore}}$
point-to-point (Figure 3.6b)	1	0.39	0.39	$\frac{1}{2}$	0.35
Three-point line-shaped bus (Figure 3.6c, without branch D)	1	0.97	0.97	$1\frac{1}{2}$	1.05
	2	1.56	0.39	2	1.40
Four-point star-shaped bus (Figure 3.6c)	1	1.67	1.67	$2\frac{1}{2}$	1.75
	2	2.23	0.56	3	2.10

Table 3.2: Simulated t_d and Elmore delay for different bus topologies, normalized ($R_{\text{wire}}C_{\text{wire}}/l^2 = 1$) and with ideal transceivers ($C_S=C_L=0$).

continue towards other destinations) impacts both the delay and the shape of the response. These different shapes are also a reason why the Elmore delay only gives an approximation of the actual delay. In Table 3.2, the delay as extracted from the simulations and the Elmore delay are tabulated. When we regard the Elmore delay as a time constant and multiply it by 0.7 to get the 50% delay as discussed above, then it matches quite well to the actual delay, with a margin of about 10%. The difference becomes 20% when we regard the Elmore delay of $\frac{1}{2}R_{\text{wire}}C_{\text{wire}}$ as a measure of the intrinsic time constant (τ_{wire}) of the interconnect, which was empirically determined as $0.41 \cdot R_{\text{wire}}C_{\text{wire}}$ in (3.1). This can be explained by the fact that higher-order parts of the transfer add delay, but do not change the dominant time constant, as is modeled in more detail in section 3.8.5.

With regard to the multi-drop buses, both the Elmore delay and actual simulations clearly show a large delay penalty compared to point-to-point buses. As mentioned before, this is why point-to-point busses are used in this project.

3.8.4 Inductance and termination extensions to Elmore delay

In recent years, a lot of work has been published with models that improve the accuracy over the original Elmore delay model or make it more widely applicable. One often used improvement incorporates the inductance into the model for delay [40, 55, 56]. As was discussed in section 3.5.2, the inductance does limit the propagation velocity, so in some cases it can be beneficial to include it, as RC models can predict physically impossible delays that would require propagation velocities larger than the speed of light.

However, as discussed earlier in section 3.5 and 3.6, for the long and thin wires that we use in this project, the inductive effects are negligible and delay approximations that only take the R and C into account are sufficient. The t_d values from Table 3.2 for example differed by less than 1% when simulated with a 1cm RC line or a 1cm RLC line, when we use the parameters from Table 3.1.

Ramp delay model

Another method to predict the dominant time constant is to take a ramp-function and compute what the delay between the ramp at the output and the ramp at the input will be [2, 57]. This ramp-model is very useful to predict the delay when a resistance is used as line

termination at the receiver, instead of the capacitive load from a gate in the conventional situation.

For this model, one assumes that the input voltage is a ramp with a certain slope p_v : $V_{drive}=p_v t$. For an RC-line, the steady-state output will then also be a ramp, with a certain delay τ and possibly with a different slope p_o : $V_{out}=p_o(t-\tau)$. In [2, 57], the solution for the delay is found by regarding the RC-line as an infinite number of RC lumps, with a ramp current through all the resistances and a constant current through all the capacitors (as the capacitor current is the derivative of a ramp). The resulting infinite series that relates the output to the input ramp has an algebraic solution, from which the delay figure τ can be extracted.

That this τ is a good measure for the dominant time constant of an RC-line can be readily shown for a first-order system. A ramp function that starts at $t=0$ has a derivative that is the step response for which we now the solution from (3.30):

$$\left. \frac{dV_{drive}(t)}{dt} \right|_{t>0} = p_v, \quad \left. \frac{dV_{out}(t)}{dt} \right|_{t>0} = p_o \left(1 - e^{-\frac{t}{\tau}} \right) \quad (3.36)$$

By integrating (3.36) from 0 to t , we can find the corresponding ramp-response:

$$V_{drive}(t)|_{t>0} = p_v t, \quad V_{out}(t)|_{t>0} = p_o \left(t + \tau e^{-\frac{t}{\tau}} - \tau \right) \quad (3.37)$$

So, after an initial exponentially decreasing transient, the output indeed becomes a ramp with a delay equal to the time constant τ .

With resistive line termination the dominant time constant τ_c of the RC-limited interconnect depends on the resistance of the transmitter (R_S) and receiver (R_L) and can be approximated by applying the ramp-model, as was done in [57]:

$$\tau_c = \frac{1}{2} R_{wire} C_{wire} \cdot \frac{R_S + \frac{R_{wire}}{3} + R_L + \frac{2R_S R_L}{R_{wire}}}{R_S + R_{wire} + R_L} \quad (3.38)$$

When we neglect the transmitter resistance and assume an infinite receiver resistance, then the time constant corresponds to the Elmore approximation $\tau_c = \frac{1}{2} R_{wire} C_{wire}$. If, instead, a resistor is used as receiver termination with a value sufficiently lower than the wire resistance then the time constant decreases (and the bandwidth increases), up to a factor of three according to equation (3.38) in the limit of zero Ohm termination ($R_S=R_L=0$).

3.8.5 Higher-order (transfer) models

So far, first order models that predict the delay or dominant time constant have been discussed. These models capture the dominant behavior of the interconnect, but disregard the more subtle effects. It is for example quite clear from Figure 3.7 that line responses differ from a standard first-order response, especially in the initial time-phase. It is also visible that different type of interconnects can have a comparable delay, but with quite a different shape of the step response.




Figure 3.8: Comparison of low-order models to actual response for a 10mm single-ended interconnect in 0.13 μ m CMOS. (a) Zoomed-in step response for time-domain comparison, and (b) transfer function for frequency domain comparison.

To capture these effects and better model the shape of the response, higher-order models can be used. One model that is often referenced, is presented in [47]. There it is proposed to use ‘asymptotic waveform evaluation’ to estimate waveforms for linear circuit models, not only including R’s and C’s but also including inductances. The model resembles the Elmore delay model in the sense that it determines moments of the response, but now extracts the first $2q-1$ moments instead of only the first moment, such that a q -pole dynamic model can be formed.

Parametric modeling

However, compared to the asymptotic waveform model, much simpler alternatives exist in case the step response of the system is already known (for example through field-solver or lumped-element simulations). In such cases, one can use a variety of methods to fit the coefficients of a low-order dynamic model to the step response. A method that is very convenient in our case, as simple Matlab functions are available, is the use of parametric modeling algorithms that were originally developed for statistical signal processing and (adaptive) filter design [58].

We used the covariance method (Matlab function ‘arcov’) to estimate the parameters of the transfer function fractional polynomial, assuming that it is an autoregressive (AR, or all-pole) system. The covariance method is a method that finds a linear predictor (a filter) with minimum forward prediction error [58]. When we use the sampled impulse response as input and remove the first few samples (to remove the delay-effect of higher-order poles), then the covariance method is able to find parameters of a low-order transfer function that match well to the actual transfer function. Results from this low-order estimation are shown in Figure 3.8, together with actual response, both in time-domain and in frequency domain.

The time-domain graph of the step-response is zoomed-in on the first part, where the differences between the different estimates are largest (the first part of the time-span in the time-domain is related to the high-frequency part in the frequency domain). On a larger

Time constant / model order	τ (s)	delay (s)	$\tau/(R'C'l^2)$	delay/ $(R'C'l^2)$
1 st	$1.42 \cdot 10^{-9}$	3.5e-10	0.41	0.10
2 nd	$1.62 \cdot 10^{-10}$	1.9e-10	0.047	0.055
3 rd	$6.3 \cdot 10^{-11}$	1.34e-10	0.018	0.041

Table 3.3: Time constants and delays which together match the low order models to the actual transfer function, for a 10mm single-ended interconnect in 0.13 μ m CMOS.

time-scale the step responses from the low-order models nicely coincide with the actual step response, except for a delay term which is added separately (to model the delay effect of the higher-order poles). On that larger time-scale, all curves resemble the step response of the point-to-point interconnect in Figure 3.7.

The frequency domain graph shows that the 3rd-order model is able to capture the behavior of the actual interconnect up to 50dB of attenuation, which is more than enough for our practical purposes.

In Table 3.3, the time constants are given for the first, second and third-order model, as well as the additional delay term (which depends on the model order). Values normalized to the RC-product of the interconnect are also given. Note that the most dominant time constant is $0.41 \cdot RC$, which corresponds well with the frequency-domain analysis from equation (3.1). If we use (3.31) to convert the first-order time constant from Table 3.3 to the 50% delay and add the additional delay then we get a total normalized delay of $0.39 \cdot RC$, which also matches well with the simulated delay (t_d) of the point-to-point interconnect in Table 3.2.

So the first-order parametric model plus delay accurately predicts the dominant behavior in both the time and frequency domain. When we would have used only a time constant without an additional delay (as is done in the Elmore delay model), then we would get a slower slope for the step response and underestimate the bandwidth of the actual interconnect.

The higher-order models further improve the modeling accuracy. Especially the 3rd-order model captures the actual response very well. A further increase of the model order can be done, but the added value reduces as a whole series of ever more closely spaced time constants is necessary to capture the higher-order part of distributed RC-behavior.

3.8.6 Lumped models

The models discussed so far in this section are very usable for high-level modeling and analysis, but they do have their limitations, as they for example do not model the impedances of the interconnect systems. They are furthermore not very usable in a circuit simulator as they are not specified in terms of circuit values but in terms of time constants or transfer coefficients.

With some adaptations, it is quite possible to incorporate these models into a circuit-simulator, by either using dedicated circuit-elements that represent low-order transfer functions or by creating simple linear circuits to generate the desired transfer functions. It should furthermore not be very difficult to incorporate the correct impedance levels into

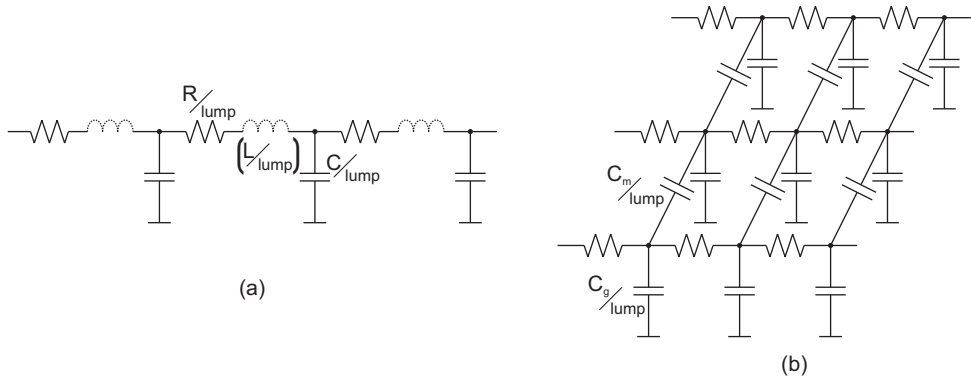


Figure 3.9: Lumped models for (a) single interconnects and (b) buses.

these models, for example with a multiple-input multiple-output (MIMO) transfer that relates the voltage and the current at input and output to each other. However, such methods are not very straightforward and experimentation will be hindered by the need for different models for different interconnect configurations.

A much simpler and widely used alternative is to use lumped element models to represent the actual distributed system. Examples of such lumped element models for both a simple single interconnect and a bus with capacitive cross-coupling are shown in Figure 3.9. The parameters in these models have direct relations to the physical distributed parameters and no conversion procedures are necessary.

Low-order lumped element models with only one or a few lumps have often been used in the past, with names corresponding to their schematic shapes, such as the L, π , T, double π , double T or triple π models. However, for a low number of sections, the direct correspondence between the lumped parameters and the distributed parameters is lost, although with the use of a correction factor, one can still get a good correspondence in transfer function [2]. With eight lumps or higher, the dominant part of the transfer functions of a lumped model corresponds to within a few tens of dB's to the actual transfer, and a correction factor can be omitted [2].

Nowadays, with the availability of fast circuit simulators, it is no longer necessary to use only a low number of lumps, the more so because a circuit simulator can solve the linear differential equations from these lumped models much faster than the many non-linear equations that describe transistor behavior. In circuit simulations we therefore use models with 100 lumps. With this high lump count and with the inclusion of inductive elements in the model, the circuit simulator is able to model the actual behavior with good accuracy over the entire frequency range.

Lumped signal models

Apart from circuit-simulators, we also wanted to simulate the interconnect behavior in more high-level tools such as Matlab-Simulink, for example to test behavioral models of different transceiver concepts. To this end, we converted the circuit-lumps to block-diagrams with separated signals for voltages and currents, as visible in Figure 3.10. A single lump is shown both for an RC and an RLC model. Cascading these block-level lumps and

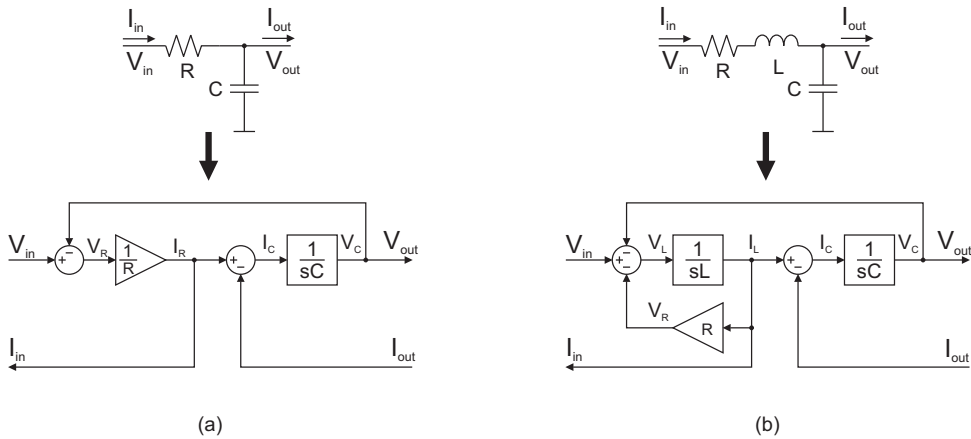


Figure 3.10: Lumped models and their corresponding block diagrams for (a) a single RC lump and (b) an RLC lump.

terminating the ends with proper signal sources or sinks, creates a dynamical model that is functionally equivalent to the circuit-level lumped models.

The last type of lumped model that we used is very similar to the block-level model from Figure 3.10, but then written in code and with a substitution of the continuous time integrators ($1/s$) by discrete time counter-parts.

Lumped discrete-time models

The dynamical models from Figure 3.10 still require an ordinary differential equation (ODE) solver, as they contain continuous-time (s -domain) integrators. When we substitute these integrators by discrete-time (z -domain) equivalents, then we can omit an ODE-solver completely. Calculating the response then translates to a simple iterative evaluation of the difference equations. When we vectorize these difference equations, with the vector indices representing the lumps, then we can use numerical software tools such as Matlab to generate sampled responses at high simulation speeds.

To convert the integrator in the RC-model to the discrete-time domain, we used the forward Euler method ($1/s = T_s/(z-1)$), also known as the forward difference approximation). The forward Euler conversion gives simple and directly usable difference equations without algebraic loops in Figure 3.10a. It is however known to produce difference equations that may be un-stable [59] when the step-size (or sample-period T_s) is not chosen adequately short. For the RC-model, experimentation showed that a step-size shorter than $0.5R_{lump}C_{lump}$ prevents instability and ensures accurate modeling.

For the conversion of the RLC-model in Figure 3.10b, we can use backward Euler conversion ($1/s = T_s z/(z-1)$) for the right integrator, without generating algebraic loops. We do this to improve stability [59]. But, as with other numerical methods to solve differential equations, stability is not guaranteed unless suitably small time steps have to be chosen [60]. For certain combinations of the RLC parameters (for example a very small R when LC is high) we can use very large step-sizes, but a more general rule that always led to stable results was to keep the step-size at a fraction of the shortest lump time constant:

$$T_s = A \cdot \min\left(R_{lump} C_{lump}, \frac{L_{lump}}{R_{lump}}\right) \quad (3.39)$$

An A that is 0.5 or smaller gave stable results in every tested situation. In practice, we normally used an A of 0.2. An A of 0.2 gives results indistinguishable from the continuous-time lumped model results, except for a slight ($\sim 1\%$) error for those cases where a large L causes under damped ringing parts in the response.

We used these time-discrete lumped models extensively for fast time-response evaluation in Matlab. We also used these models to simulate the response along RC-lines and (with some model extensions) busses and visualize the results in 2D or 3D animations to get a more direct and intuitive understanding of the behavior of the interconnects.

3.9 Summary and conclusions

The list below shortly summarizes the results and conclusions from this chapter:

- In this project, 10mm long interconnects in metal layers just below the top layers are assumed to represent typical interconnects for global data communication. These interconnects are bandlimited limited by distributed RC behavior with a low-pass corner frequency in the order of 100MHz (for M5 wires in 130nm CMOS technology)
- For the global communication, point-to-point buses with all signals traveling in the same direction are considered. Point to point interconnects have higher bandwidths than multi-drop busses.
- For 10mm interconnects, inductance and skin-effect start to play a role when the cross-sectional dimensions width and height exceed roughly $2.5\mu\text{m}$.
- For thick or short wires where skin-effect does become important, the magnitude transfer is still very similar to distributed RC behavior and the bandwidth remains proportional to the cross-sectional dimensions. Only the phase characteristic and the influence of termination impedances changes significantly.
- For behavioral modeling, simple wire models can be used, ranging from Elmore delay models (with $t_d = 0.5R_w C_w$), to third-order parametric models that, together with a delay term, capture both the time and frequency domain behavior very accurately. For more detailed analysis, lumped wire models can be used, which also capture other effects than only the transfer, such as the wire impedance.

Chapter 4

Termination, crosstalk and power consumption

4.1 Introduction

This chapter discusses how termination at the start and end of the interconnect affects the bandwidth of the interconnect. The chapter also discusses the subject of crosstalk between different interconnects in a bus and how its effects can be minimized with twisted differential interconnects. Power consumption in the interconnects and in the termination circuits is the third topic of this chapter. The reason that these three subjects are treated here together, is because they interact with each other. Termination not only influences the interconnect bandwidth, but also influences crosstalk and power consumption, as will become clear in this chapter.

The next section discusses the termination. Section 4.3 subsequently discusses the crosstalk problem, with section 4.4 discussing how twisted differential wires can mitigate crosstalk. Section 4.5 discusses power consumption and section 4.6 closes the chapter with a summary and conclusions.

4.2 Interconnect termination

It has already been briefly mentioned earlier in section 3.8.4 that the type of termination of the interconnect has a profound impact on its behavior. This is of course not only true for on-chip interconnects but also for off-chip ones (perhaps even more so).

In [2] it was already discussed in great detail which termination impedances can be used to boost the bandwidth (and lower the power) of on-chip interconnects. In this section, some additional background is given. First, in section 4.2.1, the classical approach to off-chip wire termination is compared to the classical method for on-chip termination and it is discussed why characteristic termination is not a good idea for on-chip interconnects. Next, in section 4.2.2, resistive receiver termination and capacitive transmitter termination are discussed, which are the two practical termination improvements that were used in this project. For a detailed separate treatment of these two termination schemes, see [2]. Here, we will focus more on the similarity between the two schemes in terms of transfer function, and on the creation of a simple low-order transfer model. In section 4.2.4, the discussion moves on to a special RL receiver termination scheme that was investigated in this project

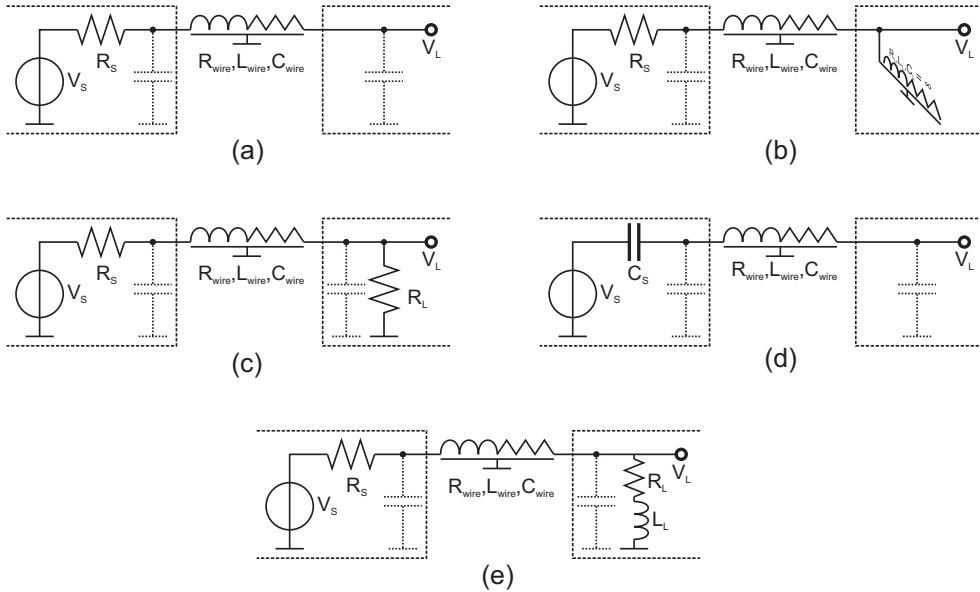


Figure 4.1: Different terminations for on-chip interconnects: (a) conventional, (b) characteristic, (c) resistive receiver, (d) capacitive transmitter, (e) resistive-inductive. Some common parasitic capacitances are also shown (dotted).

to further extend the interconnect bandwidth, but was eventually never used in a demonstrator IC. The discussion is concluded in section 4.2.5 with a short discussion of other termination concepts.

An overview of the different types of termination that will be discussed in the coming sections is shown in Figure 4.1, with classical (Figure 4.1a) versus characteristic (b) termination being discussed in section 4.2.1, (low) ohmic receiver (c) and capacitive transmitter (d) termination in section 4.2.2, and RL receiver termination (e) in section 4.2.4.

4.2.1 Classical and characteristic termination

For off-chip wires, it is general practice to terminate them with their characteristic impedance Z_0 , to avoid reflections and the associated ripples, or sometimes even deep nulls in the transfer function.

However, it is actually not favorable to terminate RC-limited on-chip interconnects with their characteristic impedance. For RC-limited channels, the characteristic impedance is a frequency-dependent imaginary term [2] which does not resemble any circuit element. Terminating the interconnect with such an impedance would be equivalent to the loading of this interconnect with another interconnect of infinite length, as shown in Figure 4.1b. In section 3.8.3 it was already discussed (and shown in Figure 3.7) that the loading of an interconnect by another interconnect is not favorable for its performance. Here, we will examine the behavior of infinite length interconnects in more detail.

Charge transport over an RC-line is governed by the same differential equations as the standard diffusion equation [41]:

$$\frac{\partial^2 V}{\partial^2 x} = RC \frac{\partial V}{\partial t} \quad (4.1)$$

A frequency-domain solution for the signal (is charge) propagation over such an RC-line is easily derived, for example from the solution of the standard telegraphers equation [2, 41, 44], as was mentioned earlier in section 3.6, equation (3.9). When we omit the inductance (L) and admittance (G), then the solution represents a true RC-line and becomes:

$$H(j\omega) = e^{-\sqrt{j\omega RC}} \quad (4.2)$$

Where $RC = R'C'l^2$. The analytical function for the accompanying step response is given by [7, 44]:

$$h_{step}(t) = \text{erfc}\left(\sqrt{\frac{RC}{4t}}\right) \quad (4.3)$$

The step response from equation (4.3) has a very long tail because signals have to diffuse over the RC-line, which is a slow process, especially for long diffusion lengths. For a characteristically terminated interconnect, or an interconnect of infinite length, it will therefore take an infinite time before the charge is fully distributed over the wire.

For the analogy with diffusion, one can speed up the process by containing the diffusion within an enclosed volume with a source at one end. For the RC line this is equivalent to a voltage source at one end - as a voltage difference is in this case the source for diffusion, as reflected in equation (4.1) - and an open other end (to stop the charge from diffusing further).

One can also argue from a signal transmission theory point of view that it is better to leave an RC line open at the receiving end instead of using characteristic termination. This was done in [2], where it was argued that the resulting reflections at the end of the wire are actually beneficial for the speed of RC-lines.




Figure 4.2: Transfer functions (a) and step responses (b) for a 10mm single-ended interconnect in 0.13 μ m CMOS with either conventional or characteristic termination at the receiver side (where a 2x scaled copy of the latter curves are also drawn).

Figure 4.2 illustrate this point and shows the transfer function and step response, both for an interconnect with conventional termination (idealized, $R_S=0$, $R_L=\infty$) and with characteristic termination ($R_S=0$, $Z_L=Z_c$). In Figure 4.2b, the long tail for the step response with characteristic termination is clearly visible. Initially, the response is similar to the response of a conventional wire, but at the 50% crossing, the delay is already significantly larger; with $t_d = 1.1 \cdot RC$ instead of $0.39 \cdot RC$ for the conventional wire (from Table 3.2). After the 50% crossing, the differences become even larger, with a tail that only very slowly settles to the final value of unity, instead of the exponential settling of the conventional wire (with $R_L=\infty$).

An interesting side-aspect that is also shown in the figure is that the high-frequency part of the transfer – and the first part of the step response – of a characteristically terminated wire actually corresponds much better to the response of a conventional wire when the former is multiplied by two. This can probably be attributed to the reflection that is present in the conventional wire [2], which doubles voltage at the end of the wire in the initial phase of the response.

When we compare the transfer function of the conventional and characteristically terminated wire then a factor 10 difference can be observed in the -3dB corner frequency (112MHz and 11MHz respectively). The transfer of the characteristically terminated wire already starts to go down at very low frequencies, albeit with a very gentle slope (slow roll-off).

Sometimes, the characteristic transmission from (4.2) is taken as the basis for evaluating interconnects, without taking the reflections into account. In [61] for example, the effects of placing active negative impedance elements along the line are analyzed in comparison to this characteristic transmission. But although the characteristic transmission is a simple, compact equation with easy to use properties such as a constant attenuation (in dB) per unit length at a certain frequency, it does give a very pessimistic view for the bandwidth. Signal

reflections from simple terminations at the end of the wire can already account for a large portion of the bandwidth improvement, as opposed to more complex techniques such as the negative impedances from [61].

Correspondence between characteristic termination and skin-effect

Interestingly enough, the diffusion equation (4.1) that describes the characteristic transmission along an RC-line also governs the behavior of skin-effect [43]. This similarity was also discussed in section 3.6. It was also shown that skin-effect is not yet a limiting factor for most on-chip wires. For off-chip interconnects, it is however one of the main bandwidth limiting effects, especially in wireline communication channels [39, 44]. Another band-limiting effect found in off-chip interconnects is dielectric absorption, but for many types of cables, the attenuation caused by the skin-effect is more dominant.

So, Figure 4.2 basically also shows (at least part of) the difference between on-chip and off-chip interconnects. When an on-chip interconnect is terminated with $R_L = \infty$, then the reflections not only speed up the wire, but also create a transfer with a dominant first-order behavior. This makes on-chip wires very simple to equalize, as will be discussed in Chapter 7.

Off chip wires do not have this dominance in the first order pole, as visible in Figure 4.2. And indeed, simple first-order pre-emphasis do not work well for off-chip wires [39, 44], except for special schemes with additional high-frequency boost, as will be discussed in section 7.5.1.

4.2.2 Resistive RX or Capacitive TX termination and their similarities

Resistive receiver termination

In the previous sub-section, the correspondence of RC-line transmission to diffusion processes was discussed. This correspondence can also be used to explain why it is beneficial to use a resistive termination not only at the source, but also at the receiving end of the wire, as was shown earlier in Figure 4.1c (and for an idealized case, also in Figure 4.3a on the next page). A resistor at the receiver acts as a sink for the charge, and the diffusion process reaches its steady state faster when there is a sink at both ends of the wire. This speed-up was already discussed quantitatively in section 3.8.4 and indicated by equation (3.38).

A sink at both ends also makes the behavior of the wire more symmetric. If R_S and R_L are chosen equal, then the behavior of the wire is entirely symmetric and the transmit and receive side become interchangeable. Such symmetric behavior is good for crosstalk minimization, as discussed in [2] and in section 4.4 but from a bandwidth perspective, it does not really matter if R_S and R_L are equal, as long as both are at least much smaller than the wire resistance. For the source-resistance, this is quite intuitive, and one would naturally be inclined to make the source-resistance as low as (economically) possible. For the load-resistance however, the side effect is that the received voltage also decreases, as it is proportional to $R_L / (R_{\text{wire}} + R_L + R_S)$.




Figure 4.3: Resistive receiver termination (a) and Capacitive transmitter termination (b), showing idealized cases without any parasitics.




Figure 4.4: Step responses for resistive receiver termination (a) and capacitive transmitter termination (b) for a 10mm single-ended interconnect in $0.13\mu\text{m}$ CMOS.

Choosing R_L can be viewed as a sort of gain-bandwidth trade-off with a fixed gain-bandwidth product. In a first-order model where the interconnect would just be a single R and C , the trade-off is completely valid, as both the RC -product and the attenuation from in to output will be proportional to $R_L/(R_{\text{wire}} + R_L)$. Actual interconnects however consists of distributed R 's and C 's and the increase in bandwidth is finite and limited to a factor three, as was also discussed in section 3.4 of [2].

To illustrate the speed-up from resistive receiver termination, step responses for different R_L values are shown in Figure 4.6a, for the idealized transceiver (with no parasitics) from Figure 4.5a.

Capacitive transmitter termination

During the project, it turned out that there is another elegant method of termination that trades of bandwidth for attenuation, which is capacitive transmitter termination. For an idealized capacitive transmitter and wire, as shown in Figure 4.5b, the response at the end of the wire is actually the same as for the resistive transmitter when a reciprocal design strategy is used: placing a capacitor C_S at the transmitter in series with the interconnect creates a capacitor divider. When C_S is smaller than the wire capacitance, then the voltage swing of the wire is reduced and the bandwidth is increased. When the divider ratio's are




Figure 4.5: Comparison of low-order models to actual response, valid for resistive receiver termination ($R_L/R_{\text{wire}}=1/10$) as well as capacitive transmitter termination ($C_s/C_{\text{wire}}=1/10$). Compare with Figure 3.8

Time constant / model order	τ (s)	delay (s)	$\tau/(R'C^2)$	delay/($R'C^2$)
1 st	$4.27 \cdot 10^{-10}$	$2.8e-10$	0.12	0.08
2 nd	$1.1 \cdot 10^{-10}$	$1.7e-10$	0.032	0.049
3 rd	$5.3 \cdot 10^{-11}$	$1.31e-10$	0.015	0.037

Table 4.1: Time constants and delays for a 10mm single-ended interconnect in 0.13 μ m CMOS, equal for both resistive receiver termination ($R_L/R_{\text{wire}}=1/10$) and capacitive transmitter termination ($C_s/C_{\text{wire}}=1/10$). Compare with Table 3.3

chosen the same as for the resistive receiver (meaning $C_s/C_{\text{wire}} = R_L/R_{\text{wire}}$), then the step responses are exactly equal, as can be seen by comparing Figure 4.6b with Figure 4.6a.

Simulations showed that this equality in response not only holds for the idealized case where R_s and C_L are zero, but also for the case where the source and load impedances are symmetric, meaning $R_s=R_L$ for the resistive receiver and $C_L=C_s$ for the capacitive transmitter.

Transfer and Low-order models for R_L or C_s termination

To extract the most important properties of the wire transfer the parametric modeling from section 3.8.5 is repeated but now with R_L and C_s termination, with $R_L/R_{\text{wire}}=1/10$ and $C_s/C_{\text{wire}}=1/10$ respectively. Time constants for the low-order models are given in Table 4.1, with the corresponding time and frequency-domain graphs of the interconnect transfer shown in Figure 4.5. Some observations can be made from this analysis:

- The first-order time constant is less than $1/3^{\text{rd}}$ of the value obtained for conventional termination (Table 3.3 on page 61). This means that, in terms of bandwidth

improvement, the results are better than the factor 2.5 that is predicted by (3.38) for $R_L/R_{\text{wire}}=1/10$.

- The higher-order components change much less than the first-order time constant. Compared to Table 3.3, the 2nd order component only decreases with 32% and the 3rd order with only 16%. So the dominance of the first-order component is reduced with termination, which will decrease the effectiveness of simple first-order equalization schemes when they are used in conjunction with C_S or R_L termination. Equalization does still have a positive influence on the achievable data rate, but less pronounced than with the conventional termination, as will be discussed in Chapter 7 (and with the quantitative data also repeated in Appendix B).
- Related is the fact that the additional delay term in Table 4.1 does not decrease as much as the first-order time constant when compared to Table 3.3, because this delay compensates for the omission of the higher-order components. When the delay term and the time constant are combined to get the 50% crossing delay ($t_{50\%} = \ln(2)*RC+\text{delay}$), then the improvement with respect to conventional termination amounts to a factor 2.3. So the improvements in terms of 50% delay are slightly smaller than predicted by (3.38).

4.2.3 Differences between a resistive receiver and a capacitive transmitter

So, as we saw above, either R_L and C_S termination can be used to increase the bandwidth, at the cost of a reduced transfer magnitude.

The cause of the bandwidth increase with C_S termination is a bit different than with R_L termination. Where resistive termination gets its bandwidth gain from the additional ‘charge-sink’ at the receiver, capacitive termination achieves a bandwidth gain because it emphasizes the high-frequency signals at the transmitter-side. This difference is visible in the signal plots in Figure 4.6, shown next to possible circuit implementations, which are discussed further below. With capacitive termination, the overshoot after each transition (the high-pass emphasis) is due to the fact that the bulk of the wire capacitance is shielded by the wire resistance. Note that this is the reason why we termed the circuit in Figure 4.6b a ‘capacitive pre-emphasis transmitter’ [62].

The big advantage of C_S termination is that it reduces the power consumption in the wire because of the lower signal swing on the wires (this relation with power consumption is discussed in more detail in section 4.5.1). The signals in Figure 4.6 show that with R_L termination, the source of the wire has a swing equal to the driver’s supply voltage, whereas with C_S termination, the swing along the whole interconnect is reduced. The increase in power efficiency from C_S termination was the prime reason to implement this concept on the second demonstrator IC, which is discussed in more detail in section 10.2

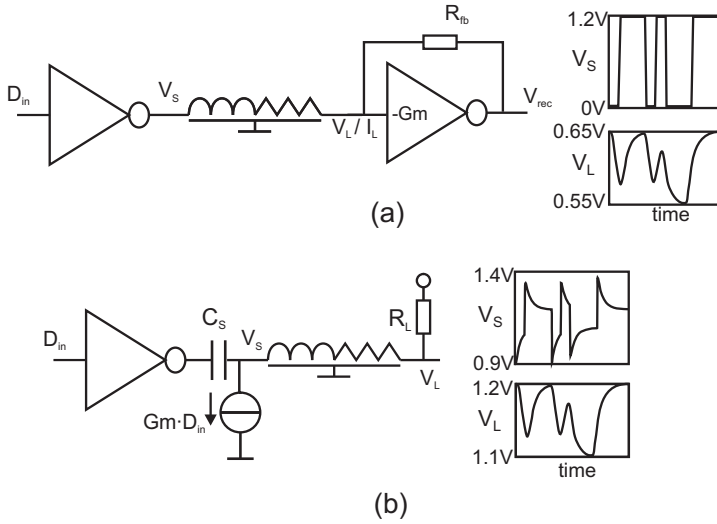


Figure 4.6: Circuit implementations and signals for resistive receiver termination (a) and capacitive transmitter termination (b).

Circuit-level differences between resistive and capacitive termination

There is a small price to pay for the power advantage of C_s termination and that is the increased sensitivity to circuit offset due to the low signal swing at the receiver. When the receiver is implemented with a real resistor, as in [63, 64] then this disadvantage also holds for R_L termination.

However, in many implementations [65-72], current-sensing amplifiers with low input impedance are used to create the resistive receiver impedance. A current-sensing receiver based on a trans-impedance amplifier, as shown in Figure 4.6a, was also used in this project [33, 73]. Its implementation will be discussed in more detail in section 8.4.2. Here we just mention its offset advantage: In a current-sensing scheme, the magnitude of the voltage transfer over the interconnect is no longer relevant and it is the voltage to current-transfer that counts. In Figure 4.6a for example, with idealized circuits, the gain of V_s to the output V_{rec} is R_{fb}/R_{wire} . When this gain is chosen to be e.g. 0.5 then the gain of the amplifier's offset would only be 1.5 as $V_{rec}/V_{offset} = (R_{wire} + R_{FB})/R_{wire}$. This offset gain is independent of the swing at V_L . With capacitive transmitter termination, the magnitude of the voltage transfer to V_L is important and one is limited in the range of transmitter capacitances as the receiver needs some voltage swing to operate reliably.

With regard to the implementation of C_s termination, some additional measures have to be taken, because with a only series capacitor (AC-coupling), the DC voltage on the interconnect is not well defined as there is no (well behaved) DC path to one of the supplies. One option would be to use DC-balanced codes and a high-ohmic load to the desired receiver potential, as was proposed in [74], but that would complicate the source. Instead, in this project, a load resistor R_L and a transconductance G_m were added to defined the DC swing, as shown in Figure 4.6b. When $C_s/(C_s+C_{wire})$ equals $G_m \cdot R_L$ then the magnitude-transfer of the high-frequency path through the circuit equals the magnitude-transfer of




Figure 4.7: Transfer functions for a 10mm single-ended interconnect in 0.13 μ m CMOS, (see Table 3.1) with different receiver termination impedances ($Z_S=0$).

low-frequency path and the transfer-function remains similar to the original. If a small G_m and a large R_L are chosen, then the static current can be kept small which avoids a significant power increase. The implementation is discussed in more detail in section 10.3.1.

4.2.4 RL receiver termination

Another form of termination that was investigated in this project is resistive-inductive receiver termination, as was shown earlier in Figure 4.1e. This form of termination emerged after experimentation with various receiver equalization circuits, inspired by the bandwidth-boosting effects that inductors have in general high-frequency circuits [75]. It turned out that when the inductance was directly coupled to the line, instead of being used in an internal node of an equalizing receiver circuit, then its bandwidth enhancement effects increased. It was also observed that the theoretically ideal receiver termination can have an inductive component (see [2] and next sub-section).

In Figure 4.7, the magnitude of the voltage transfer of the interconnect is shown for conventional, resistive and RL receiver termination. As visible, the RL-loaded wire transfer is similar to the transfer with only a resistor as load but with a significant bandwidth extension. The -3dB bandwidth of the interconnect with either conventional, resistive or RL receiver termination is 112MHz, 360MHz or 1.1GHz respectively. The inductance value L_L was determined by hand-optimization for a maximally-flat transfer. The time constant L_L/R_L is 0.333ns, which is not surprisingly in the same range as the dominant time constant of a resistively terminated interconnect (about 22% smaller, see Table 4.1), acting as a sort of pole-zero cancellation.

Circuits for RL receiver termination

So inductive termination can give an interesting bandwidth boost, but the required inductance (50nH in the example in Figure 4.7) is too high to make with a practical, small-sized, on-chip inductor. To circumvent this problem, active circuits that behave like RL combinations were investigated. Such circuits often employ a ‘gyrator’ element which




Figure 4.8: A resistive-inductive impedance (a) and a common ‘gyrator’ circuit to create this impedance (b). Source: [76].




Figure 4.9: Circuit implementation for the resistive-inductive receiver termination as tested in this project.

transforms a capacitive impedance into an inductive one. A simple example is shown in Figure 4.8.

The circuit in Figure 4.8 however also poses a few implementation challenges. It requires for example a differential amplifier with an output that can both source and sink current and ideally has a large voltage range. Normally, this is not a big problem, but this amplifier also requires a very high gain-bandwidth product, preferably with low power consumption. The most simple circuits that meet these criteria are plain inverters. A circuit with only inverters was therefore investigated in this project. Its schematic is shown in Figure 4.9.

This circuit uses inverters as transconductor elements (a.k.a. voltage-controlled current-sources or G_m 's), similar to the resistive termination circuit in Figure 4.6a. Three cascaded G_m stages are used for the feedback in this circuit, enabling a high feedback-gain and low termination resistance. However, the gains of the first two stages does have to be limited, to increase their bandwidth (and create a wideband V_{rec}/V_L transfer) and to avoid instability. The gain can be controlled by using either a ground-connected load resistance or a feedback resistor over the second stage as shown with the dotted resistors in Figure 4.9. In the latter case, the second G_m also participates in the definition of the time constant and the separate R' can be omitted. Both circuit variants were simulated in 0.13 μ m CMOS technology.

For the first variant, one can actually use V_{dd} -connected resistors and use only NMOSTS as G_m elements, which should give a bandwidth compared to PMOS/CMOS circuits. But this setup also proved to limit the gain V_{rec}/V_L , due to the low voltage available for the first resistor (as $V_{rec}/V_L = G_m \cdot R \approx G_m \cdot V_R / I_d$). The second variant with the resistor across the second G_m proved better to dimension, as the bias currents for the G_m can be decoupled

from the resistor value. With this circuit, simulations showed that a data rate of 3Gb/s could be achieved over a wire $R_{\text{wire}}=1.9\text{k}\Omega$ and $C_{\text{wire}}=0.25\text{pF}$, with the circuit having an effective R_L of 100Ohm and with 400mV swing at V_{rec} .

So significant data rates can be achieved with this circuit, but it also has some considerable drawbacks. It is not very straightforward to design and re-use for different situations, as both the stability and the effectiveness of the circuit depend on the interconnect properties (notably its impedance). Stabilizing the circuit for all possible situations (process-corners, temperatures, voltages) can be problematic. With respect to stability, note that the circuit from Figure 4.8 might be better, but that was not further investigated.

In the project, attention shifted towards capacitive termination, which has lower implementation costs and, more importantly, lower power consumption. The RL receiver termination was therefore not tested on actual silicon. In [68] however, a ‘current-mode’ receiver circuit is presented that quite resembles the three-Gm circuit from Figure 4.9. In that paper, RL-termination is not mentioned as a concept and the time constants of the circuit do not seem to be specially matched to the line, but a significant speed-up is still observed compared to conventional voltage-mode termination. However, similar to the findings in this project, it was observed that the circuit is not ideal with respect to power consumption.

4.2.5 Other types of termination

Besides the different methods for termination that were discussed above, many other possibilities exist. In [77] for example, it is proposed to combine capacitive transmitter termination with resistive receiver termination and furthermore also include a series capacitance at the receiver (also done in [78]) and a grounded resistor at the transmitter. It is claimed that such a scheme leads to additional bandwidth improvements (from 212MHz with a capacitive transmitter to 1.26GHz with the additional elements). But there are also drawbacks. For one, the low-frequency transfer is blocked, so a simple two-state receiver will no longer suffice. Also, the attenuation is high, requiring a very low-offset receiver. Although a test-chip was designed, unfortunately no measured data was published.

Other termination methods than those discussed in the previous sections were also investigated in this project, with analytically optimal source or load impedances as a basis [2]. Interestingly enough, when allowing certain attenuation in the magnitude transfer (e.g. $V_I/V_S=0.1$), then capacitive transmitter termination or RL-receiver termination actually come quite close to ideal termination. When attenuation is not allowed, then the ideal Z_S or Z_L have more resemblance to negative impedances. Active circuits that create negative impedance levels were shortly investigated but without much promising results. It should be noted that in [61, 79] good results are claimed with negative resistance circuits along the wire, but the resulting transceivers become quite complex, area intensive and not very power efficient (as is also shown quantitatively in the conclusions in Table 12.1 and Figure 12.1

4.3 Crosstalk

Crosstalk is the effect that the signal level on a (victim) interconnect is influenced by signal changes on other (aggressor or attacker) interconnects. It is a result of electrical and magnetic coupling between these different interconnects. The electrical or capacitive coupling is usually regarded as the most dominant effect between direct neighbors, while the magnetic or inductive coupling can stretch across larger distances, especially when there are no return paths for the current in the direct vicinity [8]. They are two effects that generate crosstalk in opposite directions. The capacitance between two wires pulls the victim signal in the same direction as the aggressor signal. On the other hand, a changing current in one wire induces a voltage and a current flow in the other wire, were the latter is oriented in the opposite direction as the causing current.

In this respect, it is interesting to discuss what happens when we exert a bundle of wires together, for example with a step signal. When we regard only capacitive crosstalk, then the wires in the center of the bundle will have the fastest response, as all the neighboring wires are also charged in the same direction and effectively help to charge this central wire. When we regard inductive crosstalk, then the wires in the center of the bundle will have the slowest response, as the magnetic fields of all the surrounding wires induce currents counteracting the charging.

At very high frequencies, this latter effect is dominant, as induction becomes stronger with higher frequencies. So, at high frequencies, conduction will be mostly confined to the outer wires in the bundle [45], similar to skin-effect in a single wire (also see section 3.5). At lower frequencies however, the capacitive crosstalk is dominant, especially when return paths are also added inside or close to the bundle. Return-paths reduce the inductance because they guide current in the opposite direction and hence cancel part of the magnetic fields.

The interconnects in this thesis are usually RC-limited, with negligible inductive effects inside the pass-band (section 3.5). Inductive crosstalk effects are hence also very small. This is even more so as we use differential wires to mitigate crosstalk, as will be discussed below. These differential wires provide very local return-paths, which reduces inductance even further (while it increases capacitance, see Table 3.1). In the remainder of this section, we will therefore focus on capacitive crosstalk and ignore inductive crosstalk.

Of the many types of capacitive crosstalk, crosstalk between neighboring wires in a bus is most severe. Crosstalk from unrelated perpendicular wires in other metal layers can also have some effect, as discussed in [2], but this effect is smaller and the signals are difficult to predict and can be better modeled as noise. In this section we concentrate on the dominant deterministic crosstalk between neighbors.

In this section we also assume that wire dimensions are optimized for highest BW/area (see section 2.7) such that the distributed capacitance of the wire is hence evenly distributed in all directions ($C_{\text{top}}=C_{\text{bottom}}=C_{\text{side}}$ in Figure 2.5). See section 3.2 for a more detailed discussion of the assumed routing style and wire dimensions.

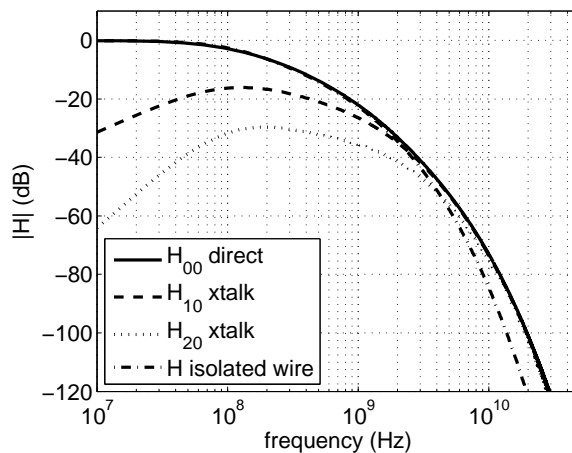


Figure 4.10: Direct and crosstalk transfer functions for 10mm single-ended interconnects in a bus in 0.13 μ m CMOS.

4.3.1 Capacitive crosstalk problem

To illustrate capacitive crosstalk, transfer functions of wires in a bus are shown in Figure 4.10. The figure shows the transfer of an aggressor wire both to its own receiving end (H_{00}), to the receiving end of the neighboring victim wire (H_{10}) and to the receiving end of the second (non-direct) neighbor victim (H_{20}). As reference, a graph for the transfer of an isolated wire (with all its capacitance to ground) is also shown in the figure. These graphs were obtained from lumped element simulations as described in section 3.8.6 (with the discrete-time models with 100 RC lumps per line, using a Fourier transform on the differentiated step response to obtain the transfer function).

The figure shows that crosstalk between wires in a bus has an impact on the achievable data rate, as the ratio of received signal power to crosstalk-interference power decreases with increasing frequency. At really high frequencies, the transfer magnitude to the wire's receiving end and to its neighbors even becomes equal. The figure also shows that the attenuation at high frequencies is lower than for a single isolated wire. This is due to the fact that in the bus, not all the interconnect capacitance connects to ground, but some part connects to the neighboring interconnects which are also partly charged. When we reduce or eliminate the charging of the neighboring wires – for example with grounded shields or with twisted interconnects as discussed below – then there is more correspondence with the transfer of a single interconnect.

A time-domain figure for the crosstalk in the bus is shown in Figure 4.11a (obtained with the same model as used for Figure 4.10). The time-axis is normalized to the RC product of the wire to enable easier comparison with other wires (e.g. those from Figure 3.7), independent of wire length or technology. The figure shows that the direct neighbors of the aggressor do indeed receive most of the crosstalk. As was discussed above, the response of the aggressor wire itself is slightly faster than the response of a single isolated wire because




Figure 4.11: Normalized step response (a), with zoom-in (b) for a bus where only the center wire is excited (with $R_S=0$, $C_L=0$ and $C_{side}=C_{top/bottom}=0.25 \cdot C_{wire}$).

not all its capacitance connects to ground. Inspection showed that the 50% delay is only $t_d=0.36 \cdot RC$, while the delay of an isolated wire was $t_d=0.39 \cdot RC$ (from Table 3.2 on page 58).

Note that it was tried to create low-order models for the transfers from Figure 4.10 (and the step-response from Figure 4.11), as was done earlier for single wires in section 3.8.5 and 4.2.2. However, the results were not very accurate, as the wires with crosstalk have zeros in their transfer-function, while the parametric model that was discussed in section 3.8.5 was developed for all-pole systems. The results were thus not further used in signal simulations. This is not really problematic, as for most simulations it is assumed that crosstalk is canceled with twisted wires (as discussed below), for which the models from section 3.8.5 and 4.2.2 apply.

For first-order estimates for the effect of crosstalk, the peak crosstalk voltage is often used. A simple estimate for this peak voltage is given in [8]:

$$V_{xstalk-peak} = V_{swing} \frac{C_{coupling}}{C_{total}} \frac{1}{1 + \frac{\tau_{att}}{\tau_{vic}}} \quad (4.4)$$

Where τ_{att} and τ_{vic} are the time constants for the aggressor and victim drivers respectively. In Figure 4.11, those two are equal and indeed the peak noise of ~ 0.12 corresponds to the equation, which predicts 0.125 given that $C_{coupling}$ is 25% of C_{total} .

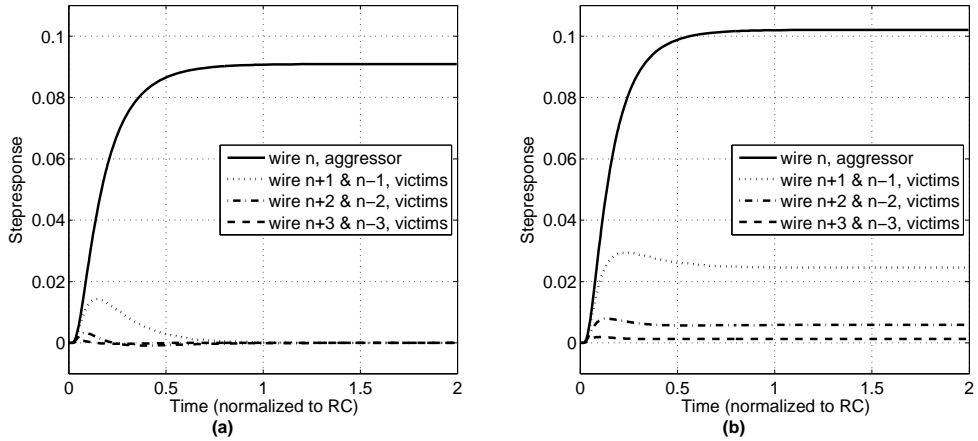


Figure 4.12: Normalized step responses for (a) resistive receiver termination with $R_L=R_{wire}/10$ and (b) capacitive transmitter termination with $C_S=C_{wire}/10$. ($R_S=0$, $C_L=0$).

Crosstalk and wire termination

The peak crosstalk equation (4.4) is only valid for conventional termination. The crosstalk with resistive receiver or capacitive transmitter termination (as discussed in section 4.2.2) is different, as can be seen in Figure 4.12.

In the step response in Figure 4.12a, it is visible that a resistive receiver not only lowers the wire delay, but also increases the peak crosstalk compared to Figure 4.11 (The quantitative data is also listed in Table 4.2 on page 86). For the capacitive transmitter in Figure 4.12b, the crosstalk is even higher and it also does not return to zero when time progresses. This is because the capacitive transmitter not only charges its own wire, but also the neighboring wires, as all transfers are defined by capacitive ratios.

For the capacitive transmitter, the charging of the neighboring wires also reduces the effective ground-capacitance of the aggressor, which is the reason why the swing in Figure 4.12b is higher than the swing in Figure 4.12a (while the R_L and C_S were dimensioned for equal DC transfer when a single wire is considered, as was done in section 4.2.2). The lower effective capacitance to ground also results in a 10% lower delay (see Table 4.2) for the capacitive transmitter termination compared to the resistive receiver termination.

The additional G_m/R_L path that was discussed in section 4.2.3 to better define the DC transfer of the capacitive transmitter will slightly change the low-frequency part of the step response (and give a swing defined by G_m/R_L). It will also reduce the crosstalk at low frequencies and avoids that the neighboring wires become permanently charged by the aggressor.

Note that the first part of the crosstalk in Figure 4.12 is still quite similar to the zoomed-in response from Figure 4.11b, because the high-frequency behavior of the wires are similar for all types of termination. The high crosstalk in this first part of the step-response is not desirable, because to obtain high data rates, we generally use higher symbol rates, which means that we are mostly interested in the first part of the step-response as shown in Figure 4.11b. On this short time-scale, the more distant wires are also clearly affected by crosstalk.




Figure 4.13: Twisted differential interconnects to mitigate crosstalk.

In a sense, the high frequency part of the transmission spreads out over the entire bus (as was also observed from the frequency domain plot in Figure 4.10). The lower the frequency, the more transmission is confined to the region of the wire itself.

Crosstalk is thus especially problematic when we want to increase the achievable data rate by using the higher-frequency parts of the transfer (with termination or other signaling techniques). Quantitatively, crosstalk between RC-limited neighboring wires in a bus in one metal layer already reduces the achievable data rate with 42% when we use plain binary signaling and conventional termination, as will be discussed later in more detail in section 6.2.1. Crosstalk problems become even worse when the surrounding metal layers are also used as data paths in the same bus.

A standard method to reduce crosstalk is to increase the spacing between the wires or insert shield-wires and shield-planes [35, 80], where the latter option also helps to define a return path and reduce inductive crosstalk. To enable the highest data rates for each channel, one would need to place a shield wire between every signal wire, but at the cost of increased wiring resources and a lower BW/area [35]. To also minimize crosstalk at high frequencies, the shields need to be connected to low-impedance ground nodes at regular position along their path. Otherwise, high-frequency crosstalk still spreads out to more distant wires as sketched above.

4.4 Differential twisted wires for crosstalk reduction

A method to mitigate crosstalk that is much more robust than shielding is the use of twisted differential interconnects, as shown in Figure 4.13. The twists invert the sign of the crosstalk for successive wire segments, such that (most of) the crosstalk is cancelled at the end of the wire, provided that the twist positions are chosen correctly [81, 82]. The (minimum) number of twists to be used and their optimal positions depend on the type of wire termination and bus configuration, as will be discussed in more detail in section 4.4.4 and section 4.4.7.

Application of twisted differential wires to reduce crosstalk is quite common in a number of other fields, such as wireline communication over twisted pairs. On-chip, twisted differential interconnects are also widely used in CMOS memory cells, to cancel crosstalk between bitlines [83].

When using twists for on-chip communication, it is favorable to not use too many of them, as each twist occupies a section in another metal layer (blocking routing space) and the contacts can add quite some resistance. In [84] for example, twisted interconnects were used for communication over global interconnects, but it was mentioned that the wire resistance increased because the via resistance due to the eight twists was overlooked. In this project, on our demonstrator IC's, we used only a single twist in the even channels and two twists in the odd channels, as shown in Figure 4.13. The size of the twists in the figure is greatly exaggerated for visual purposes. In our actual layout, one of the two wires ducked under the other via the metal layer below, with the whole twist occupying an area of about four square micron.

4.4.1 Costs and benefits

At first sight, the use of differential interconnects would seem to incur a doubling of the area. The bandwidth of the wires also decrease due to the Miller multiplication of the capacitance ($C_{se-wire} = C_{top} + C_{bottom} + 2C_{side}$, $C_{diff/wire} = C_{top} + C_{bottom} + 3C_{side}$) and to top it off, power increases by more than a factor of two due to the doubling of the number of active wires and the increase in capacitance per wire. However, in reality these drawbacks are alleviated and are counterbalanced by a number of advantages that differential wires provide.

The area overhead of differential wires is smaller than a factor two because single-ended channels also need overhead to reach high data rates, in the form of shield wires as discussed above. Even with proper shields, single-ended wires and their transceivers are less robust than (twisted) differential channels [85]. This is because differential interconnects are also immune towards noise-sources such as V_{dd} disturbances or crosstalk from perpendicular wires because they translate to common mode noise. The improved immunity to common-mode disturbance is a very general reason to use differential circuits. For this project, it for example enabled the use of a differential sense amplifier with a low offset and a high power-supply rejection as receiver [86], which can operate reliably at much lower noise margins than a single-ended latch or logic cell. This advantage is shared by other alternatives that use single-ended data wires and a shared reference, such as the pseudo-differential interconnect from [85]. The ability to cancel crosstalk is however not present in pseudo-differential interconnection schemes.

The fact that twisted differential wires can reduce crosstalk improves their performance per area cost. This is discussed quantitatively in section 6.2.2, where it is shown that there is a break-even data rate above which differential wires obtain higher aggregate data rates per cross-sectional area than the single-ended alternatives.

Despite the doubling of the number of active wires, differential transceivers can even be more power efficient than their single-ended counterparts because a differential transceiver can operate with lower signal swings. This is discussed theoretically in section 6.2.2 and is demonstrated with circuit-simulations of practical transceiver circuits in section 11.3.

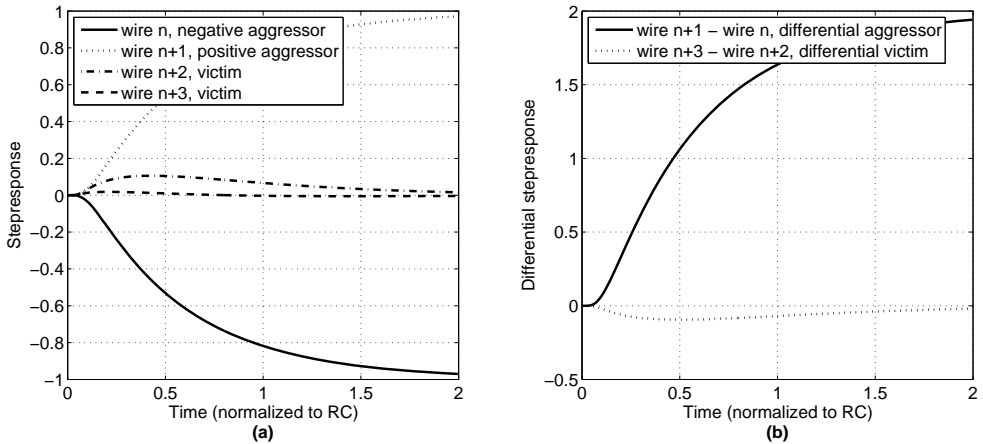


Figure 4.14: Step response with conventional termination as in Figure 4.11, but now with a differential aggressor (still without twists), showing the signals for each wire (a) and the differential signals (b).

4.4.2 Crosstalk in differential wires without twists

The use of differential wires without twists already helps to reduce the effect of crosstalk, as the effective signal swing is doubled. This is clearly visible in Figure 4.14a, where aggressor and victim signals are plotted for the case that two aggressor wires switch in opposite direction to model differential operation. Examination of the differential signals, as plotted in Figure 4.14b, reveals that the crosstalk is attenuated compared to the single-ended signals in Figure 4.11 (also see Table 4.2 on page 86), especially the high-frequency part (close to $t=0$). The earlier mentioned effect that high-frequency crosstalk spreads out to more distant wires than only the direct neighbors is thus easily reduced with differential signaling.

4.4.3 Modal analysis for crosstalk signals

We could also have reached the conclusion that differential signaling reduces crosstalk by examining the single-ended step-response from Figure 4.11, by looking at the difference between the response of victim wire $n+1$ and $n+2$. This difference is the same signal as the crosstalk response that a differential aggressor would cause on a single-wire. In that sense, the step-responses from Figure 4.11 can be used as a basis to construct more complex signals. Signals for situations where multiple wires are active can easily be found by superimposing shifted scaled copies of this basis response.

The method described above is one variant of the so-called ‘modal analysis’ [87]. Modal analysis is used to find solutions of dynamic systems with multiple outputs when excited by an input. Above, we described how a single-input multi-output (SIMO) response could be used as a basis to solve responses in more complex multi-input multi-output (MIMO) situations. It is however also possible to solve the response with a multi-input excitation, which can be easier to solve analytically for certain well-chosen input signals. In [81] for example (and in [2, 82]), Eisse Mensink used modal analysis to obtain solutions for the

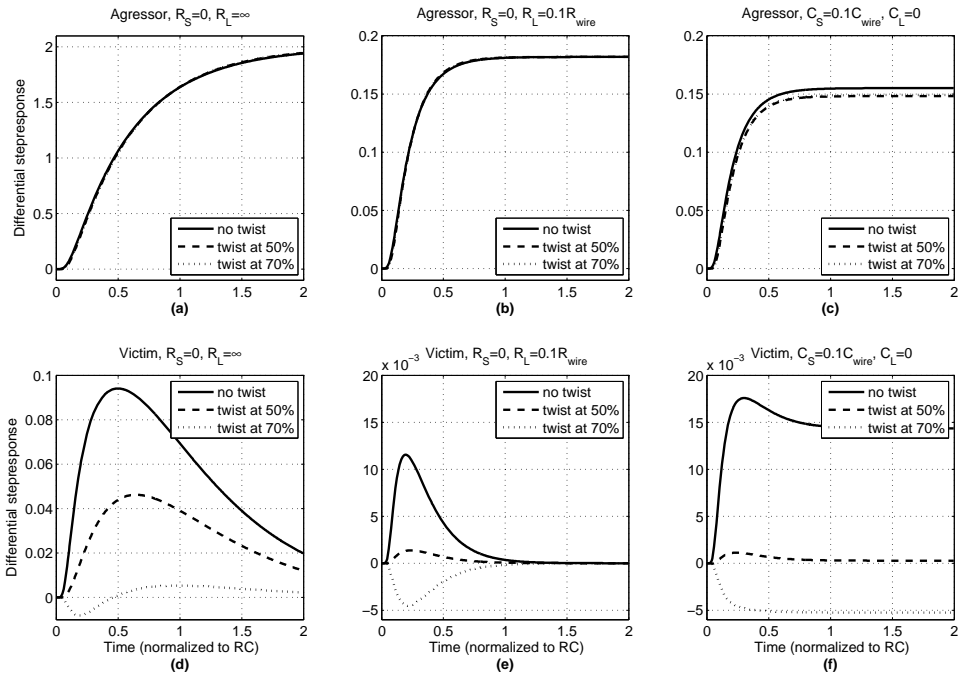


Figure 4.15: Step responses of aggressors and victims for a differential bus with either no twist or a single twist at 50% or 70% of the wire length, with conventional termination (a, d), resistive receiver termination (b,e) or capacitive transmitter termination (c,f).

(crosstalk) transfer functions of twisted differential interconnects. The response was evaluated for two situations: In the first, the interconnect pairs were all excited with the same sign (even mode) and in the second with opposite sign (odd mode).

4.4.4 Twist analysis and positioning

Intuitively it is clear that twists in the interconnects, as in Figure 4.13, can reduce crosstalk when they are spaced such that neighboring pairs get an equal amount of positive and negative crosstalk along their length. But with the above mentioned modal analysis, it can be shown that positioning of these twists is not trivial. The optimal positions and the effectiveness of twists depend for example on the type of termination (in [82], this is analyzed in detail for conventional and resistive termination).

Figure 4.15 illustrates this for a single twist in the uneven channels, which is sufficient to reduce the crosstalk from a differential aggressor to the differential-mode signal of the neighboring victims. When we use conventional termination as in Figure 4.15a,d (with an ideal zero-ohm driver and no load capacitance) and place a twist at 50% of the length, then the high-frequency part of the crosstalk is almost gone, but there is still a sizable residue of low-frequency crosstalk. To cancel the dominant part of the crosstalk, a position at 70% of the length is more favorable [81, 82], but a bit of crosstalk remains.

With resistive receiver termination, the suppression of crosstalk is more efficient and the optimal position of the twist is at 50% of the length for all frequencies [81, 82]. The same holds for capacitive termination.

Note that the swing for the capacitive transmitter is now lower than for the resistive receiver instead of the higher swing that was found for the single-ended wires in section 4.3.1. This is because the capacitance between the two differential halves is increased by the Miller multiplication. It is also because the charging of the neighboring wires is now partly canceled, depending on the twist position, which increases their effective capacitance to ground and Figure 4.15c,f show that the better the crosstalk cancellation, the lower the swing. For the case of the 50% twist position, the crosstalk is nearly perfectly canceled and the swing for the capacitive transmitter is $1/1.25$ of the swing with resistive termination due to the Miller multiplication. When the C_S is also increased by a factor 1.25, then the results match those of the resistive termination, as is also shown quantitatively in Table 4.2, except that the capacitive transmitter has a slightly lower delay. Detailed examination of the step-responses showed that the difference with resistive termination increases in the tail of the response, with a delay difference up to $0.016/RC$. This difference seems to be caused by the fact that the capacitive transmitter has increased crosstalk at the transmitting end of the wire (that is later cancelled by the twist), which slightly speeds up the charging of the wire, as was also the case for the single-ended wires from Figure 4.12. This difference in step response speed explains why the capacitive transmitter is able to reach higher data rates than the resistive receiver, as will be discussed in section 6.2.2.

4.4.5 Quantitative results for delay and crosstalk

Based on the simulations from the previous sections, some quantitative crosstalk results for the different situations are shown in Table 4.2 on the next page. The delay and swing of the aggressor wires themselves are also shown. The figures in the table reconfirm the benefits and drawbacks of the differential wires. On the one hand, differential signaling increases the delay and on the other hand it reduces the crosstalk. With conventional termination, the delay of a differential pair is a factor 1.30-1.34 larger than the delay of the single-ended interconnect (depending on the type of termination). With the addition of the twist, this delay difference can even become a factor 1.41 (for capacitive termination). With only miller-multiplication, the delay difference should be a factor 1.25. The remaining difference can be explained by the fact that differential signaling and the twisting cancels part of the charging of the neighboring wires and the capacitance to these neighbors thus behaves more like a grounded capacitance (as was also discussed in section 4.3.1).

Additional simulations were also carried out with symmetric termination impedances, meaning $R_S = R_L$ for the resistive receiver termination or $C_L = C_S$ for the capacitive transmitter and the results are also listed in Table 4.2. With symmetric termination, the crosstalk is even further reduced compared to the situation with the ‘idealized’ termination that was used in the previous sections ($R_S = 0$, $C_L = 0$). The resistive termination analysis in [82] even predicts total cancellation of the crosstalk with symmetric termination and a twist at 50%. In the lumped-element simulations that were used to create the data in Table 4.2, there is still a small crosstalk residue also with a twist at 50%. It is not entirely clear if this is due to simulation artifacts from the lumped model or if the analysis in [82] is too optimistic.

4. Termination, crosstalk and power consumption

Channel type	V_{swing}	t_d (50%)	$V_{\text{xtalk-peak}}/V_{\text{swing}}$	$t_{\text{xtalk-peak}}$
Single-ended				
Conventional, $R_S=0, R_L=\infty$	1	0.356	0.120	0.36
Conventional, $R_S=0.1R_{\text{wire}}, R_L=\infty$	1	0.425	0.119	0.44
Resistive, $R_S=0, R_L=0.1R_{\text{wire}}$	0.091	0.154	0.158	0.15
Resistive, $R_S=R_L=0.1R_{\text{wire}}$	0.083	0.179	0.155	0.17
Capacitive, $C_S=0.1C_{\text{wire}}, C_L=0$	0.102	0.139	0.288	0.23
Capacitive, $C_S=C_L=0.1C_{\text{wire}}$	0.092	0.169	0.254	0.29
Differential				
Conventional, $R_S=0, R_L=\infty$	2	0.467	0.047	0.50
Conventional, $R_S=0.1R_{\text{wire}}, R_L=\infty$	2	0.558	0.047	0.60
Resistive, $R_S=0, R_L=0.1R_{\text{wire}}$	0.182	0.201	0.063	0.19
Resistive, $R_S=R_L=0.1R_{\text{wire}}$	0.167	0.234	0.062	0.23
Capacitive, $C_S=0.1C_{\text{SE-wire}}, C_L=0$	0.155	0.187	0.113	0.30
Capacitive, $C_S=C_L=0.1C_{\text{SE-wire}}$	0.143	0.215	0.103	0.36
Differential, twist at 50%				
Conventional, $R_S=0, R_L=\infty$	2	0.473	0.023	0.65
Conventional, $R_S=0.1R_{\text{wire}}, R_L=\infty$	2	0.567	0.019	0.79
Resistive, $R_S=0, R_L=0.1R_{\text{wire}}$	0.182	0.204	0.0075	0.24
Resistive, $R_S=R_L=0.1R_{\text{wire}}$	0.167	0.237	0.0011	0.22
Capacitive, $C_S=0.1C_{\text{SE-wire}}, C_L=0$	0.148	0.196	0.0076	0.24
Capacitive, $C_S=0.1C_{\text{Diff-wire}}, C_L=0$	0.182	0.202	0.0086	0.25
Capacitive, $C_S=C_L=0.1C_{\text{SE-wire}}$	0.138	0.223	0.0018	0.44
Differential, twist at 70%				
Conventional, $R_S=0, R_L=\infty$	2	0.477	-0.0041	0.19
Conventional, $R_S=0.1R_{\text{wire}}, R_L=\infty$	2	0.570	-0.0052	0.26
Resistive, $R_S=0, R_L=0.1R_{\text{wire}}$	0.182	0.203	-0.025	0.21
Resistive, $R_S=R_L=0.1R_{\text{wire}}$	0.167	0.236	-0.028	0.25
Capacitive, $C_S=0.1C_{\text{SE-wire}}, C_L=0$	0.149	0.196	-0.035	∞
Capacitive, $C_S=C_L=0.1C_{\text{SE-wire}}$	0.139	0.223	-0.034	0.41

Table 4.2: Simulated swing, delay (t_d) and crosstalk data for wires in a bus with different configurations and different termination impedances. The RC product is normalized (and $C_{\text{side}}=C_{\text{top/bottom}}=0.25 \cdot C_{\text{wire}}$).

Interestingly, additional lumped-element simulations with slight tweaks of the twist position showed that the low-frequency part of the crosstalk for the capacitive transmitter reduces to zero when the twist is moved to 51% of the wire, but also for the lumped element models, the peak crosstalk is lowest with a twist at 50%.

In any case, the crosstalk is very small with a twist at 50%, both for the completely symmetric termination ($R_S = R_L$ or $C_L = C_S$), as well as for the case where $R_S = 0$ or $C_L = 0$. Non-zero R_S or C_L do increase the delay of the wire (as can be seen in Table 4.2). For the achievable data rate analysis in Chapter 6 and Chapter 7, the bandwidth-optimal situation is therefore used, and it is assumed that the transmitter has infinite drive strength ($R_S = 0$) and the receiver has no (parasitic) capacitance ($C_L = 0$).

4.4.6 Twists to reduce common-mode crosstalk

The two twists in the odd channels that were also shown in Figure 4.13 are used to also mitigate crosstalk from a differential aggressor to the common-mode signal of the neighboring victims. This common-mode crosstalk is minimized if the twists are placed at about 30% and 70% of the length [81, 82]. These twists would not be strictly necessary when the differential receiver can cope with common-mode variations. But practical receivers can benefit from inputs with a stable common-mode voltage and the double twists have little disadvantages (when enough vias are used to get a low-Ohmic twist) The following sub-section will discuss when it is desirable to further increase the number of twists.

4.4.7 Twisting patterns to reduce crosstalk in Multi-layer buses

In this project, we mainly focused on the mitigation of crosstalk between neighbors in a bus that is located in a single metal layer (as described in section 3.2). But in a number of cases, it can be desirable to use multiple metal layers for the bus (and deviate from perpendicular routing schemes). A NoC for example can benefit from a multi-layer bus, when dedicated portions of the chip area reserved for the network.

When one would simply stack multiple copies of the configuration from Figure 4.13 on top of each other, then one would create crosstalk between the bus channels in the different metal layers. Staggering the copies, such that a channel with a single twist would be above or below a channel with a double twist, as shown in Figure 4.16 does help, but still gives crosstalk between the diagonal channels as indicated in the figure.

Of course, crosstalk due to diagonal capacitive coupling is already much lower than when crosstalk would come from all sides (as in a single-ended multi-layer bus), but it will still have some degrading impact on the reliability of the communication. Applications that require really high data rates or very low interference levels would need more crosstalk reduction. These applications could be for example busses that use low-swing or multi-level signals or, in another application area, perhaps even busses that transport analog signals




Figure 4.16: Multi-layer twisted bus with single (1) and double (2) twists. The arrows indicate the remaining, non cancelled crosstalk. (a) Perspective view. (b) Frontal view

Fortunately, it is not difficult to extend the twisting approach and add more twists to cancel crosstalk between more channels than just the two horizontal neighbors. To this end, we should realize that a twist in a differential interconnect represents an inversion of the polarity of the crosstalk. Ideally, these inversions are spaced such that the crosstalk contribution of each part cancel when summed and observed at the receiver.

For an initial first-order approach, assume that a differential pair can be divided into N segments and that each segment has an equal contribution to the crosstalk at the receiver (a more detailed discussion of the crosstalk contribution is given in [82]). Furthermore assume that we can set the polarity of the crosstalk from each segment independently (with proper twisting). We can then search for polarity schemes that give a net zero crosstalk between different pairs.

For example consider again the twisted bus from Figure 4.13. The differential pairs in that bus can be assumed to consist of 4 segments. The first segment is the part up to the first twist position (the first 30% of the length, as was discussed in the previous section). The second segment is the part between the first twist position and the middle twist position (where the latter should be at 50% of the wire length for capacitive or resistive termination). The third segment is the part between the middle and third twist position (with the latter at 70% of the wire length) and the last segment is the remaining part up to the receiver.

By convention, assume that the polarity of a differential pair is +1 when the positive single-ended half is on the left side of the pair, and -1 when it is on the right side. For pair 0 in Figure 4.13, we can thus say that it has a polarity distribution over its segments equal to [1 1 -1 -1]. Pair 1 has a polarity distribution of [1 -1 -1 1]. The differential pairs next to each other will exhibit positive crosstalk for those segments where the single-ended halves with equal polarity are next to each other and negative crosstalk otherwise. The crosstalk between pair 0 and pair 1 cancels out as $[1 1 -1 -1] * [-1 -1 -1 1] = (-1 + 1 - 1 + 1) = 0$, as was also shown graphically with the capacitors between the pairs. However, the crosstalk between the non-direct neighbors pair 0 and pair -2 does not cancel out as $[1 1 -1 -1] * [1 1 -1 -1] = 4 \neq 0$.

In general, to quantify how much crosstalk a twisted differential pair with a certain polarity distribution v_i – from here on called polarity vector - receives from another pair (j), we can take the inner product of their polarity vectors:

$$X_{talk_{diff}} = -v_i * v_j' \tag{4.5}$$

The minus sign is not present for crosstalk between vertical neighbors as in Figure 4.16, because vertical neighbors have their positive terminal on the same side.

When v_i and v_j are mutually orthogonal, then their inner product is zero and the net crosstalk from a differential aggressor onto the differential signal of a victim will be zero. So, to create good twisting patterns (with a twist representing a change in polarity), we want to find polarity vectors that are orthogonal to the polarity vectors of the neighboring wire pairs, at least for those neighbors for which we desire to cancel the crosstalk.

To aid in the search for orthogonal polarity vectors, we can make use of the so-called Hadamard matrices, also known as Walsh matrices. These matrices consist of rows that contain only plus and minus ones (same as polarity vectors) and are orthogonal to each other. Hadamard matrices have applications in several different areas, including combinatorics, signal processing (CDMA communication), and numerical analysis [88]. A Hadamard matrix satisfies $H_N H_N^T = I$, or in other words, it is an orthonormal matrix (with all columns or rows being mutually orthogonal). Hadamard matrices with a size of $N \times N$, with N a power of two, can be constructed iteratively, starting from a simple unity scalar:

$$H^1 = 1$$

$$H^{2N} = \begin{bmatrix} H^N & H^N \\ H^N & -H^N \end{bmatrix} \quad (4.6)$$

As an example, Hadamard matrices of order two and four are shown below:

$$H^2 = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad H^4 = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix} \quad (4.7)$$

The third and fourth row of H^4 are equal to the polarity vectors of the wires with the single and double twists from Figure 4.13, as was discussed above. We could also have used e.g. the first row (no twists) and third row (one twist), which would also have given differential crosstalk cancellation, but then we would not have cancelled the common-mode crosstalk, as was also discussed in 4.4.6.

Cancellation of common-mode crosstalk

Although differential-mode crosstalk is more important, for many receivers it is still also of interest to keep the common-mode crosstalk low. The amount of common-mode crosstalk that is produced by a differential aggressor (a.k.a. the amount of differential to common-mode conversion) only depends on the polarity vector v_j of the aggressor and not of the victim, as both single-ended halves of the victim have equal sign for common-mode crosstalk:

$$Xtalk_{com} = [1 \quad 1 \quad \dots \quad 1]^* v_j \quad (4.8)$$

This means that common-mode crosstalk is only caused by aggressors that have polarity vectors with a non-zero sum. In a Hadamard matrix, that is only the case for the topmost

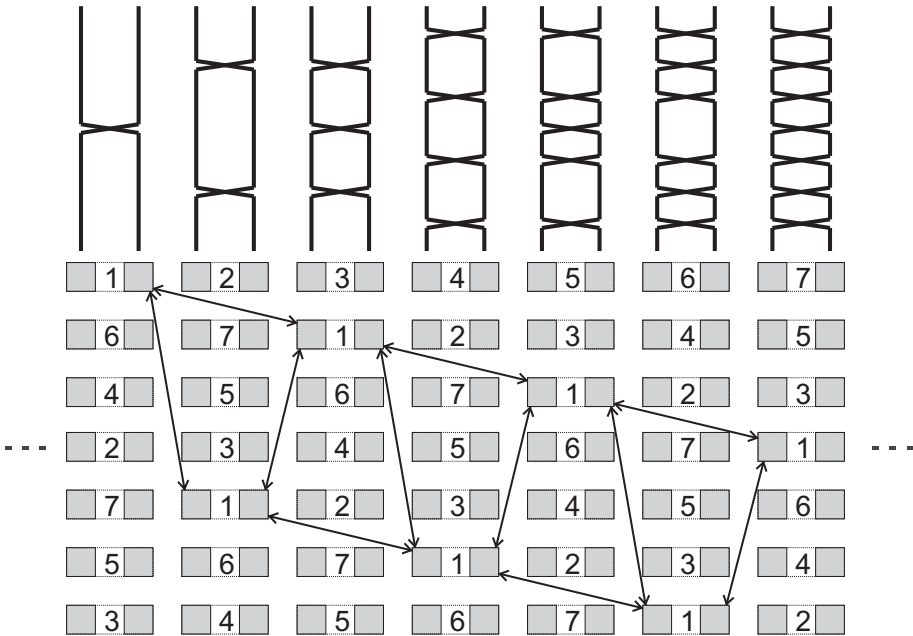


Figure 4.17: Frontal view of a multi-layer twisted bus with seven different twist types. The twist pattern for the different type-numbers is shown at the top of the figure. The remaining crosstalk is also indicated for the channels with one twist (type 1).

row-vector (v_1). If we do not use this first row, then we can use the remaining rows as twist specification, such that pairs with different polarity vectors will neither cause differential nor common-mode crosstalk onto each other. What remains is the decision how to distribute the available vectors over the channels, as differential pairs with the same vector will generate crosstalk onto each other ($v_i * v_i \neq 0$).

When a H^4 is used, then there are only three different twisting types available, which is not enough to arrange channels in a multi-layer bus such that all crosstalk between neighbors is canceled. An H^8 matrix is the next available Hadamard matrix which leaves seven possible twisting types. An example of a possible twisting arrangement with this H^8 matrix is shown in Figure 4.17. The type numbers in the figure have been chosen such that they represent the number of twists needed. The ordering of the rows in the H^8 matrix is different (type 1 to 7 correspond to row 5, 7, 2, 4, 8, 6, 2 respectively).

For the top layer, the different twisting patterns from type 1 to 7 are placed next to each other. So within this layer, there is no net crosstalk from any pair onto its 6 most adjacent neighbors. To also minimize crosstalk between the layers, the position of each type of twisting pattern is shifted two positions the right. This shift of two positions leaves only crosstalk between diagonal, non-direct neighbors (as indicated with the arrows in the figure for crosstalk between pairs of type 1).

This twisting arrangement does require seven ($N-1$) positions along the length of the bus to place the twists. So there are some layout costs, but the reward is a mitigation of almost all the crosstalk, except for a small remaining part between distant channels.

Twist positioning

As with the original single and double twists from section 4.4, the positions of the twists will not simply be equidistant and the type of termination will also have its influence on the ideal positions. But with the theory that is presented in [82], it should not be a problem to find N segments (of unequal length) that all have equal crosstalk contribution at the receiver. In [82], in the 'simplified low-frequency model' section, simple linear or quadratic equations are presented that describe the crosstalk contribution (for low frequencies) at every point along the interconnect. The segment lengths should be chosen such that the integral of the crosstalk contribution is equal for each segment.

So in conclusion, twisted differential interconnects provide a very simple way to mitigate crosstalk between neighbors in a bus and this method can be extended with the addition of more twists, to be able to also mitigate crosstalk between non-direct neighbors and between channels in a multi-layer bus.

4.5 Interconnect power

Power consumption due to (and in) interconnects is becoming an increasingly large part of the total power consumption of large scale digital IC's. Usually, most power in a chip is consumed in the active devices, for example because the capacitance density is much higher for active devices – with their very thin dielectrics – than the capacitance density of interconnects. But some of these active devices are actually only there to drive interconnects, such as repeater circuits (see section 8.5). Both the number of repeaters and their size is expected to rise significantly for future large scale designs [89, 90], at least for conventional circuit architectures.

The power in interconnects themselves is also projected to increase. Earlier issues of the ITRS [37] predicted that the power per layer of interconnect, per unit of bandwidth and area would plateau to about $1.5\text{W}/\text{GHz}/\text{cm}^2$ in future technologies, due to advancements in for example low- k dielectrics. However, the latest ITRS [3] predicts that this figure will increase to higher values (to about $2\text{W}/\text{GHz}/\text{cm}^2$ in 2024). The total power consumed in the interconnects will increase much more rapidly due to the increase in number of metal layers (section 2.6) and increases in frequency. In this last respect, actual scaling as predicted by the ITRS deviates from classical Dennard scaling that predicts a constant power per area over different technologies. That prediction mismatch can most likely be attributed to the fact that actual power supply scaling differs from the (constant field) Dennard model.

So in all, the power consumption due to communication over interconnects is expected to become significant in the future and it therefore makes sense to analyze this power - as will be done in this section - and investigate possible low-power circuits and signaling schemes (as discussed elsewhere in this thesis).

4.5.1 Classical interconnect power consumption

For interconnects with classical termination and with data communication at not too high speeds, power consumption is easy to estimate and well known. We repeat the equations here for completeness.

Each time a capacitor is charged with a (constant) voltage source V_{dd} , it costs a certain amount of energy E , proportional only to the new voltage and the initial voltage V_{init} as that voltage difference defines the amount of charge Q that is needed:

$$E_{rise} = \int_{t=0}^{t_{end}} V_{dd} I(t) dt = V_{dd} \int_{t=0}^{t_{end}} I(t) dt = V_{dd} Q = V_{dd} (V_{dd} - V_{init}) C \quad (4.9)$$

This same equation holds for an interconnect when the voltage on the interconnect capacitance reaches its steady state. Normally, the initial voltage on the capacitance is zero, simplifying (4.9) to CV^2 . As the energy stored on a capacitor is only $\frac{1}{2}CV^2$, the other half is dissipated in the resistance of the driver and of the interconnect (disregarding special cases where part of the power is radiated or where the driver has a source impedance that is also capable of storing energy, such as an inductor). The energy that is stored on the interconnect capacitance is dissipated when the interconnect is discharged.

So, assuming simple binary transmission, an amount of CV^2 is drawn from the supply each upward data transition (rising edge), of which half is dissipated during that transition and the other half during a subsequent falling edge. A power metric that is often used for communication circuits is the costs in energy per bit. The energy per transition can easily be converted to energy per bit when we assume binary data with stationary statistical characteristics [91], with a certain transition probability p_{trans} . As only half of the transitions (the rising edges) cost energy, the average energy cost will be:

$$E/bit = p_{trans} \frac{1}{2} E = p_{trans} \frac{1}{2} CV^2 \quad (4.10)$$

For fully random data, which has a transition probability density of 0.5, the energy per bit is $\frac{1}{4}CV^2$. The energy per bit is easily converted to an actual power-consumption by incorporating the data rate f_{clk} in bits/s into the equation:

$$P = E/b \cdot f_{clk} = f_{clk} p_{trans} \frac{1}{2} CV^2 \quad (4.11)$$

4.5.2 General model for interconnect power consumption

The simple CV^2 model from the previous section is only adequate in the most simple cases and is not sufficient to predict the power for many of the transceivers discussed in this thesis. When the line is for example not terminated with a capacitance, then additional current will flow into this termination and additional terms are needed. Or when the line does not reach a steady-state situation after each bit, then the power consumption also changes. To still be able to estimate power consumption in these more general cases, we developed a more complex model in this project.

To capture the frequency dependent behavior of the line (and also of the source-data), we start with general frequency-domain power formula's. For statistical data-sources, the power-spectral density (PSD) P_{xx} is usually used to characterize the data in the frequency-domain [92]. When we have a voltage source as transmitter, then the PSD is easiest specified as a 'voltage'-related power, in terms of V^2/Hz . It can be converted to actual power by the incorporation of the impedance:

$$P = \int_{f=-\infty}^{\infty} V(f)I(f)df = \int_{f=-\infty}^{\infty} \frac{P_{xx}(f)}{Z_{in}(f)} df = 2 \int_0^{\infty} \frac{P_{xx}(f)}{Z_{in}(f)} df \quad (4.12)$$

As we are only interested in the actual power consumption and not in the reactive power, we have to take the real part:

$$\begin{aligned} P_{in} &= 2 \int_{f=0}^{\infty} \text{real} \left(\frac{P_{xx}(f)}{Z_{in}(f)} \right) df = 2 \int_{f=0}^{\infty} P_{xx}(f) \text{real}(Y_{in}(f)) df \\ &= 2 \int_{f=0}^{\infty} P_{xx}(f) \frac{\text{real}(Z_{in}(f))}{|Z_{in}(f)|^2} df \end{aligned} \quad (4.13)$$

So when the PSD of the data is known and the impedance Z_{in} or the admittance Y_{in} of the interconnect (including the source and load impedance) is also known then the power can be evaluated (analytically or numerically).

The PSD of the data will depend on the chosen signaling technique and source coding, modulation and equalization can all have an impact on the PSD. Many PSD examples for different baseband and bandpass digital codes are given in [92], but they usually involve purely random data ($p_{trans}=0.5$).

As the power consumption in on-chip wires is quite dependent on the data activity (p_{trans}) and because a wide range of data activities can be found in on-chip signals, we have to incorporate p_{trans} into the PSD for meaningful analysis.

For the cases of random binary data transmission, we can use a Markov (also named Markoff) source to acquire the PSD. A binary Markov source is a statistical source that is characterized by:

$$\left. \begin{aligned} \Pr(a_{n+1} \neq a_n) &= p_{trans} \\ \Pr(a_{n+1} = a_n) &= 1 - p_{trans} \end{aligned} \right\} a = \{1, -1\} \quad (4.14)$$

A visual representation of such a source can be given by a two-state Markov-chain, with the corresponding probabilities associated to the state-transitions [91]. When this binary Markov source has zero mean, then it's auto-correlation function is given by:

$$R_{xx}(k) = (1 - 2p_{trans})^{|k|} \quad (4.15)$$

The PSD of the transmitted data with this auto-correlation can be computed directly when the transmitted pulse shape is also known [92]. However, it is also instructive to first look at some discrete-time properties of the signal.

Random data with an autocorrelation as in (4.15) can be represented as the output of a first-order all-pole filter that has white noise w with a power of σ_w^2 as input [58]:

$$H_m(z) = \frac{1}{z - (1 - 2p_{trans})}, \quad \sigma_w^2 = 1 - (1 - 2p_{trans})^2 \quad (4.16)$$

The discrete-time power spectral density of the signal is:




Figure 4.18: Power consumption versus transition probability at different data rates, for a 10mm single-ended interconnect in 0.13 μ m CMOS with $C=2.3$ pF ($V_{dd}=1.2$ V). Source: [2].

$$P_{data}(\Omega) = |H_m(e^{j\Omega})|^2 \sigma_w^2 = \frac{1 - (1 - 2p_{trans})^2}{1 - 2(1 - 2p_{trans})\cos(\Omega) + (1 - 2p_{trans})^2} \quad (4.17)$$

Not surprisingly, when equation (4.17) is evaluated, it shows that all power is concentrated at DC when the transition probability is zero and all power is found at half the sampling frequency ($\Omega=\pi$) when the transition probability is one. In between these extremes, the power is distributed over all frequencies.

With this discrete-time PSD, it is only a small step to the actual power-spectral density. The discrete-time PSD just replicates around every multiple of the sampling frequency f_{clk} and the spectrum of the pulse shape $F(f)$ acts as a filter that passes one of the copies and attenuates the others (albeit not perfectly):

$$P_{xx}(f) = |F(f)|^2 \cdot f_s \cdot P_{data}\left(2\pi \frac{f}{f_s}\right) \quad (4.18)$$

So, with equations (4.13), (4.17) and (4.18), one can numerically evaluate the power consumption, given a certain interconnect impedance (including Z_{source} and Z_{load}) and a certain pulse shape. The impedances can be determined with the s-domain equations from [2]. The desired pulse shape depends on the chosen signaling and equalization method. For plain binary signaling, the pulse shape is just the familiar square pulse with a sync function as frequency-domain description [92].

In [2], this power estimation model was applied to calculate the energy cost per bit for simple binary transmission with conventional termination at different data rates, as shown in Figure 4.18. At low data rates or low transition probabilities, the simple CV^2 model from equation (4.10) is fairly accurate, but at higher switching-activities the actual power consumption is significantly lower because the interconnect response is too slow to reach a steady state after each transmitted bit, meaning that not every part of the interconnect is

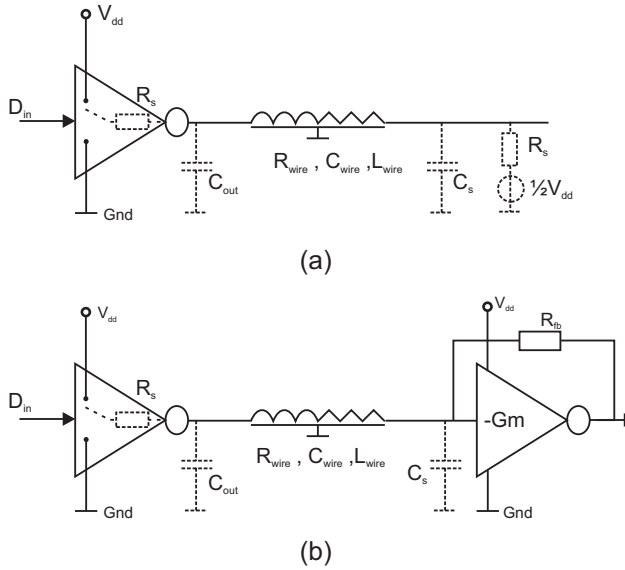


Figure 4.19: Transceiver with resistive termination with passive receiver impedance (a) and actively created impedance (b).

always fully charged or discharged, which saves some energy but complicates reliable signaling.

In section 10.2, the model is applied to calculate the power with a variety of termination impedances.

Additional power consumption in transceiver circuits

The model that is presented here calculates the power in the interconnect and in the termination impedances. Part of the power that is consumed in the transceivers is also implicitly included in the model. A source impedance R_s is for example often part of the transmitter, as e.g. the output resistor R_{DS} of a MOST in an inverter. In the transceiver circuit in Figure 4.19a for example, all the power consumption can be modeled by calculating the effective impedance Z_{in} as seen by the source.

Note that the voltage-source at the receiver does not consume or deliver power when the data has zero mean, as in the Markov source from equation (4.14). However, when the passive impedance at the receiver is replaced by an actively created one, as in Figure 4.19b, then the predictions from the model break down and one has to find other models to incorporate such powers.

The input load of a transmitter, its crowbar currents and other power-consuming transceiver properties differ between circuit implementations and can also not be captured in a general model. These factors are however not as dominant as the primary power path in the transceiver, the path through the interconnect itself, for which the power model is valuable.

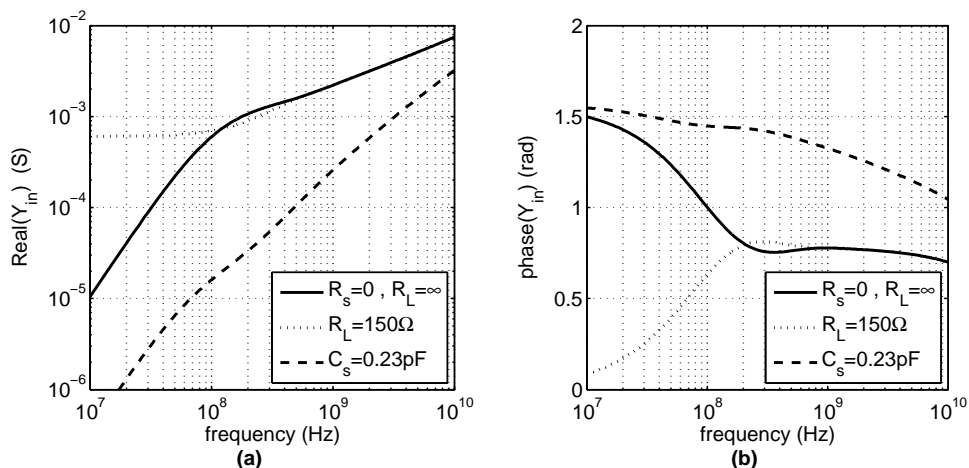


Figure 4.20: Real part (a) and phase (b) of interconnect admittance levels as seen from the source, for a 10mm interconnect in 0.13µm CMOS with different terminations.

4.5.3 Power efficiency versus signaling bandwidth

We will conclude the section on power consumption with a discussion on the challenge to both achieve very low power consumption and high data rates over on-chip interconnects. That it is difficult to achieve really high power efficiency can be seen when looking at the admittance levels of an interconnect; as shown in Figure 4.20, with different curves for different termination impedances. At low frequencies, the interconnect with conventional termination as well as the one with capacitive transmitter termination show a capacitive impedance/admittance behavior, so there is no power dissipation in the interconnect itself at these frequencies – according to equation (4.13). For the conventionally terminated interconnect however, a much more lossy admittance becomes visible above about 100-200MHz, which is essentially the frequency region where far-end reflections no longer play a role in the transfer and impedance [2]. The admittance phase decreases to about $\frac{1}{4}\pi$ (45°), meaning that in theory half of the total power that is injected in this frequency region is reactive power that can be recovered; the rest is consumed inside the interconnect. As the attenuation also rapidly increases for higher frequencies, one will have to inject considerable power to transmit information in the high-frequency region.

Interestingly, in the frequency range between about 100MHz and 1GHz, the resistively terminated interconnect actually has a lower loss component in the admittance than the conventionally terminated one. The interesting effect of this dip in the loss component is that for high data rates (with random data), both the conventional interconnect and the resistively terminated interconnect consume roughly the same amount of total power. This is because the static (DC) power of the resistive termination is compensated by the lower high-frequency loss (this is worked out in more detail in [2]).

Static power is however quite undesirable in large-scale digital IC's, as many communication channels are only occasionally active. So, this is a serious drawback of resistive termination. It can be alleviated by only using the resistance when it is actually

necessary [70, 93], but at the cost of increased complexity. A more desirable approach is to choose a scheme that intrinsically has a limited DC power consumption.

One could in theory make a transmitter that would not only be free of DC power consumption, but would also consume almost no power in the rest of the low frequency region. Such a transmitter would store and re-use the energy to and from the interconnect capacitance, with for example an inductive source impedance (not very dissimilar to a class-D amplifier). An inductor (especially a low-loss one) is however not very practical to implement for an on-chip transceiver. A simpler but slightly less efficient approach could be to re-use the charge that becomes available during discharging of one interconnect, for example to charge a capacitor that serves as the supply (battery) for another interconnect. With a bus of N interconnects, one could create $N-1$ virtual supplies, each with a lower voltage than the previous. Such a scheme would be very similar to conventional low-swing transceivers that use dedicated supply voltages (discussed in more detail in section 11.3), but without the need for a dedicated supply.

However, such a method still has the drawback of finite achievable data rates. At high frequencies, the wire impedance simply becomes too lossy and large voltages are needed to charge and discharge it sufficiently fast. In this respect, the capacitive transmitter termination (section 4.2.2) is quite ideal: it uses the energy that is build-up in the transmitter capacitance to create a large voltage pulse after each data transition (visible in Figure 4.6b) and thereby manages to charge and discharge the wire sufficiently fast while still enabling adequate power efficiency. This is also reflected in Figure 4.20, which shows that the reactive admittance part of the capacitively terminated wire extends to much higher frequencies than for other two wire examples.

In theory, capacitive transmitter termination would still be less efficient than low-swing signaling with a dedicated supply, as the former has to draw all its charging current from the normal voltage instead of a dedicated low supply. The practical overhead in power consumption can however be significantly reduced with capacitive transmitter termination, which can make capacitive termination even more power efficient than conventional low-swing signaling, as will be discussed in more detail in section 11.3.

4.6 Summary and conclusions

The list below shortly summarizes the results and conclusions from this chapter:

- For on-chip wires it is not efficient to use characteristic termination. Leaving the receiving end of the wire open – which is the classical method of on-chip termination – actually yields a 10 times higher bandwidth (or 2.8 times shorter delay).
- Resistive receiver termination or capacitive transmitter termination reduce the voltage swing and increase the bandwidth, with more than a factor three decrease of the dominant time constant compared to conventional termination. The second-order and third-order time constants decrease by only about 30% and 15% respectively.
- Compared to a resistive receiver, a resistive-inductive receiver can boost the bandwidth by another factor of three with no changes in the low-frequency behavior. However, such a receiver is less straightforward to design and not as power efficient as a capacitive transmitter.

- Resistive receiver (R_L) or capacitive transmitter (C_S) termination in combination with differential signaling over twisted wires is a powerful technique to both suppress crosstalk and improve the speed of the transmission. With single-ended transmission, the peak crosstalk from a single neighbor is high: 15% ($R_L=0.1R_{\text{wire}}$) or 25% ($C_S=0.1C_{\text{wire}}$) of the swing. With a twist at the optimal position of 50% of the length, this crosstalk is reduced down to 0.75% (when $C_L=0$, $R_S=0$) or even lower for symmetric termination (0.11% when $R_S=R_L$ or 0.18% when $C_L=C_S$).
- The number of twists to be used depend on which type of crosstalk has to be suppressed. Only one twist in the even pairs mitigates differential neighbor-to-neighbor crosstalk. Adding two twists in the uneven pairs also reduce the common-mode crosstalk. More twists can be added to reduce non-neighbor crosstalk or crosstalk in a multi-layer bus.
- For conventional termination and low data rates, the classical $E/\text{bit} = \frac{1}{2}p_{\text{trans}}CV^2$ formula can be used, but it overestimates the power at high data rates and high switching activities. For high data rates, and for other types of termination, the power model from equations (4.13), (4.17) and (4.18) is suitable.

Chapter 5

Data communication analysis

5.1 Introduction

With the communication medium analyzed in the previous chapters, this chapter will discuss the communication process itself and how the bandwidth of a channel affects the achievable data rate. The discussion in this chapter is kept general, to broaden the application area and simplify scaling of the results, independent of the specific communication medium that is used. Most parts of this chapter are therefore also applicable to other communication channels that suffer from bandwidth limitations. Compared to existing (general) communication literature, a different angle is taken, in the sense that bandwidth limitations is considered as the prime data rate bottleneck. Within this framework, an analysis method is presented and used to quantitatively predict the achievable data rate for various signaling techniques.

This chapter starts with a discussion of the general concept of data communication and how on-chip communication fits in this framework in section 5.2. Next, in section 5.3 it will be qualitatively explained how finite bandwidth and crosstalk impact communication. In section 5.4, a numerical method is presented to quantitatively analyze the finite bandwidth and crosstalk effects for a variety of transmission schemes. Synchronization is another aspect of data communication that is shortly discussed in section 5.5. The chapter is summarized and concluded in section 5.6.

5.2 General versus on-chip data communication

Digital communication in a general sense means the transmission of pieces of information from one location to the other, where the pieces of information are encoded in quantized quantities, often plain binary bits. Although manual transmission is possible (with e.g. a telegraph machine), nowadays the communication is usually handled by automated digital communication systems. Digital communication systems come in great many varieties and a lot of data has been published on this topic over the course of time.




Figure 5.1: General elements of a digital communication system.

Figure 5.1 shows a more or less standard example of elements present in a general digital communication system, as can be found in many communication textbooks [92, 94, 95]. Note that in this thesis, we will call the collection of elements that are dedicated for the communication a ‘transceiver’.

The blocks in Figure 5.1 all handle different aspects of the communication, starting with a source-encoder to compress the original information, followed by a channel encoder to add redundancy to the signal to enable error detection and correction. The next phase is the modulator that converts the data into a form that is suitable for transmission over the channel. The channel itself is the actual medium over which the communication takes place, which can be anything from air (RF transmission) to on-chip wires. The physical properties of the medium often also include some imperfections that cause signal attenuation and the addition of noise. The receiving end of the communication system tries to recover the original data as good as possible given these imperfections. To achieve this, the receiving blocks operate on the signal with inverse functions of their transmitter counterparts.

Although the system breakdown in Figure 5.1 offers a standard and well accepted framework to break down and analyze different aspects of communication, it does not really capture the elements that are usually found in an on-chip communication system, at least not with enough detail. It does for example not show explicit transmitter and detector elements, which are the primary building blocks for on-chip transceivers. Parameters of the transmitter and detector such as bandwidth and offset have a high impact on the practical achievable data rate and data integrity.

Also, a number of blocks from the general communication system from Figure 5.1 are not (yet) found in on-chip communication systems because they are too complex or simply not useful. Channel encoding to detect or correct bit-errors is for example not (yet) very useful as the dominant error sources are not random noise, but deterministic interference which can be removed by other, more effective measures. Data modulation is an aspect that also has not received much attention in on-chip communication literature. Actually, (complex) contemporary modulation schemes, including multi-carrier schemes such as OFDM or CDMA, have little to offer for on-chip channels, at least not in terms of achievable data rate, as will be discussed in more detail later on.

A signal operation that is also not shown in Figure 5.1, but actually is quite beneficial for on-chip communication, is the use of equalization. Equalization can be performed either at the transmitter or receiver side or at both sides, and it can boost the achievable data rate by




Figure 5.2: Elements of an on-chip communication system (showing only a single channel for simplicity). Optional components are shown in dashed and dotted lines.

a significant amount. A schematic of the elements that can be found in a contemporary on-chip communication system is shown in Figure 5.2. These elements will be discussed in more detail in subsequent sections and chapters.

Next to a different emphasis on communication blocks, compared to traditional communication systems, a different approach to the analysis of data-detection is also used in this thesis. In traditional communication analysis, a statistical approach is usually used, under the assumption that the primary source for errors is additive white Gaussian noise (AWGN). It is also often assumed that the signal is optimally equalized, filtered with a matched filter and subsequently sampled at the proper sample-instant. This process yields a certain signal level (or signal-power E_b) in the presence of a certain noise power (N_0) and with some assumptions, the ratio E_b/N_0 , can be translated to a statistical bit-error-rate [92].

However, in a number of communication systems such as wireline, backplane and especially on-chip communication, the really stochastic error sources such as thermal noise are by far not the most dominant error-signals. Instead, error sources such as crosstalk, inter-symbol interference (ISI), receiver offset, etcetera, are much more important. It can not be assumed that the signal is perfectly bandlimited, optimally equalized, nor sampled at the perfect location, so we have to be able to analyze the effect of finite channel bandwidth on the received signal properties to determine the margins for the detector.

In most communication textbooks [92, 94, 95], some attention is paid to bandlimited channels. In [92], the discussion is limited to Nyquist's theories about perfect equalization filters, such as raised cosine-rolloff filters (section 3-6). In [94], the discussion is extended with 'partial response coding' methods (such as duo-binary signaling) to control ISI (section 4.6). The most comprehensive discussion is found in [95], where substantial attention is also given to linear equalization (chapter 10 and 11), including imperfect equalization (section 10-2) and adaptive equalization (chapter 11). But the material in [95] is still not readily applicable to on-chip wires.

Overall, these textbook treatments of finite bandwidth effects are too general to be sufficient for the analysis of on-chip communication (and probably the same holds for wireline and backplane). Apart from this generality, they usually also focus on only one channel, omitting crosstalk between channels, whereas crosstalk between (physical or virtual) channels can have quite a detrimental effect.

At the other end of the spectrum, there is the classical way of on-chip communication analysis, which focuses primarily on the delay of signals [7]. That is also a somewhat limited approach that does not really lead to cutting-edge communication speeds. In this thesis we will try to approach on-chip communication from an angle somewhere in between classical communication and traditional on-chip communication analysis and try to find the

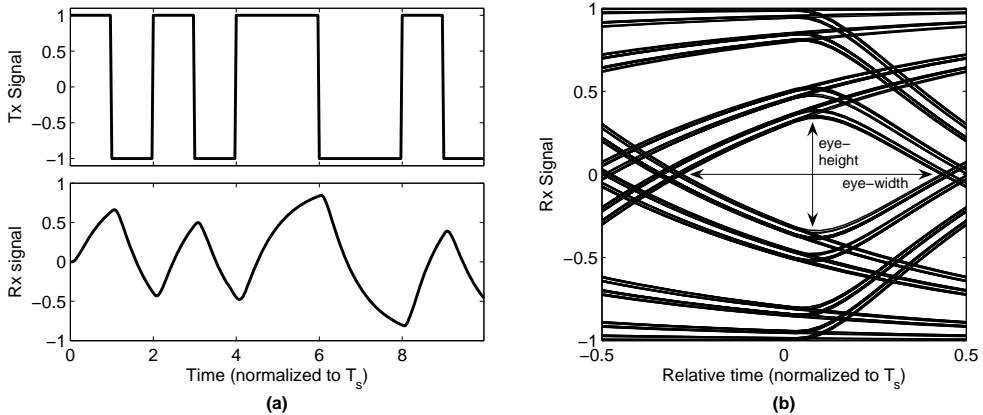


Figure 5.3: Transmitted and received symbol stream (a) and eye diagram (b) with a symbol-time of $T_s = \frac{1}{2} R_{\text{wire}} C_{\text{wire}}$. This would amount to $f_s = 1/T_s = 580 \text{ Mbit/s}$ for a single-ended or $f_s = 490 \text{ Mbit/s}$ for a differential interconnect, 10mm in length in $0.13 \mu\text{m}$ CMOS.

quantitative limits for communication speed and reliability, within the framework of practical on-chip transceivers.

5.3 Data transmission with finite bandwidths and crosstalk

5.3.1 Reliable data detection and eye diagrams

Traditionally, engineers who design on-chip data links are mainly interested in the (50%) delay of a signal edge when passed (transmitted) through a circuit (or interconnect). Delay of cascaded non-clocked (combinatorial) circuits and interconnects can simply be added and the total delay of the cascade is hence easily found. After the data has passed through the combinatorial circuits or interconnects, it is clocked or latched in a storage element (a latch or a flip-flop) where it is held and used in the next period. In the common synchronous design style, this total delay of a chain of non-clocked elements should in no circumstance exceed the clock-period, or else timing violations occur with bit-errors as a result [41]. In this design-style, interconnects, together with their driving and receiving circuits also have to abide to this constraint.

However, from a communication theory standpoint, it is not the delay of a channel that counts, it is whether data is detectable at the receiver side. For this purpose, a visual aid is often used in the form of a so-called eye diagram (or eye pattern) [41, 92] which shows the data signal, usually at the receiver side, with the consecutive symbol periods overlaid on top of each other. An example of such an eye diagram is shown in Figure 5.3b, using plain binary transmission (also known as ‘bipolar’ ‘NRZ’ transmission, ‘on-off keying’, etc). Note the different bands in the eye-diagram, which are caused by the dominance of the first-order component in the transfer, with an exponential decrease in influence of previous symbols. What is also shown, in Figure 5.3a, are the original signals with normal non-overlaid time-axis, both at the transmitter and at the receiver end, where the latter is used to create the eye-diagram.

5.3.2 Eye diagram properties

An eye-diagram can immediately show if it is possible, - without further signal conditioning - to detect the original bits from the received signal. As long as there is a clear difference between positive and negative received symbols – an open eye – then a detector should be able to decide if a one or zero was transmitted.

Such a detector should do a few things. First it should sample the received signal in the open part of the eye and next it should decide whether this sample is a one or zero (amplify, or regenerate the received sample). So the detector should have knowledge of the best sampling instant within the symbol, which requires some form of time-reference, linked to the time-reference of the transmitter (a common clock-signal for example). Ideally, the time-reference should generate equidistant timing events, centered precisely in the most open part of the eye. Unfortunately, practical timing references have some amount of timing-jitter, so it is important that the eye has a certain width in which data can be sampled.

For similar reasons, a vertical opening is desirable, as the detector will inevitably have only a finite precision for the regeneration (decision-making) of the sampled value, due to circuit error sources such as offset.

Note that the time-location of the largest eye opening, relative to the start of the transmission of the symbol, defines the delay or ‘latency’ of the communication. In off-chip communication, this latency can span many symbols despite the fact that the propagation-velocity is close to the speed of light; which means that several symbols can be ‘in-flight’ at any given time. In on-chip communication, the latency relative to the symbol-period T_s is much lower because the propagation delay is not an independent quantity, but is linked to the bandwidth limitation (as $t_d \sim 0.5R_{\text{wire}}C_{\text{wire}}$). So, the traditional focus on delay is not that strange for on-chip communication. However, when the data rates start to approach the limits, as in Figure 5.3, or with data rate enhancing techniques that will be discussed later, the latency can exceed one symbol period T_s . This means that the clock-timing for the receiver should ideally be skewed from the clock of the transmitter. This is good to keep in mind, as it requires a slightly different approach than the classical synchronous design styles that focus on skew-less clock distribution across chips.

So vertical and horizontal eye-openings, and latency are important properties for the analysis of communication over bandlimited channels. These properties can be extracted a posteriori from a simulated or measured eye-diagram, but that is a time-consuming operation that requires many overlaid symbols before it is possible to determine the eye-properties with a reasonable accuracy.

It will be shown in the next section that it is also possible to determine these eye-diagram properties a priori by analyzing the response of the channel to a single isolated symbol – the symbol response - with an approach that has some resemblance to the approach used in [96].

Apart from the adverse effect that a finite channel bandwidth has on the eye diagram, there is also the potential problem with crosstalk, which will be discussed next, before the actual quantitative symbol response analysis.




Figure 5.4: Symbol streams for three adjacent single-ended wires in a bus (a) and eye diagram (b) with the same symbol-time as in Figure 5.3 ($T_s = \frac{1}{2}R_{\text{wire}}C_{\text{wire}}$).

5.3.3 Eye diagrams and crosstalk

As was explained in Chapter 4, special layout measures such as twisted differential interconnects can be used to largely mitigate crosstalk between neighbors in a bus. However, these measures do come at a power and area cost and it is therefore also valuable to analyze the effect of crosstalk on reliable detection when these layout measures are not used (which will be done quantitatively in section 6.2). That crosstalk can have a significant impact is shown in Figure 5.4, where three adjacent single-ended wires in a bus all transmit data. Transitions in the data clearly affect the receiver response of neighboring wires, with the most influence if two neighbors switch in the same direction. The same RC-limited wires are used as in Figure 5.3, with the same data rate, but the eye is clearly more closed.

The eye diagram in Figure 5.4b also looks much more chaotic (or random) than the eye in Figure 5.3b, and one could argue that stochastic analysis techniques could be useful here. However, crosstalk between neighboring channels in a bus is in essence still an entirely deterministic process, and it can also be analyzed by examining symbol responses, as will be shown in the next section.

5.4 Symbol response analysis

5.4.1 Symbol response introduction

The effects of finite bandwidths and the effect of crosstalk create interference on the wanted signal. The effect of finite bandwidth creates so-called ‘inter-symbol interference’ (ISI), because previous (or sometimes following) symbols interfere with the current one [92, 94, 95]. The effect of crosstalk is called inter-channel interference (ICI), where channel could mean a physical channel, such as a neighboring wire, but it could also be a virtual channel, as in a multi-access or multi-carrier system. The same abbreviation is therefore also used to denote ‘inter-carrier interference’.

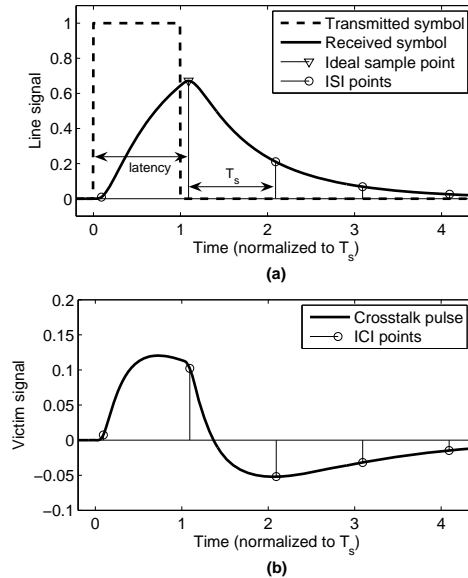


Figure 5.5: Symbol response (a), and crosstalk response (b) as received by the far-end of an adjacent wire, using $T_s = \frac{1}{2}R_{\text{wire}}C_{\text{wire}}$ (same as in Figure 5.4).

A way to analyze ISI and ICI is to look at the response of a channel to a single isolated symbol [97, 98]. An example of such a symbol response is shown in Figure 5.5a, with a symbol pulse shape that is representative for simple pulse-amplitude modulation (PAM) with normalized amplitude. The symbol duration chosen for the example is $\frac{1}{2}R_{\text{wire}}C_{\text{wire}}$, the same as used for Figure 5.3 and Figure 5.4 (this duration is chosen for no other reason than that it leads to visually pleasing results). Figure 5.5b shows the crosstalk symbol response from one channel to its neighbor. Note that, in case of binary bipolar signals, a zero level is actually never transmitted and is only present in the symbol response; a sequence of bipolar symbols would switch between one and minus one.

The value of the symbol at time t_d (the delay or latency) is the actually wanted signal. The ideal t_d is the instant of maximum vertical eye-opening ($t_d = t_{\text{ideal}}$). For simple channels and modulation schemes, this ideal sample instant is usually found at the time where the symbol response is the highest. Even with the ICI, which is quite high at this instant as shown in Figure 5.5, the ideal sample instant deviates only $0.015T_s$ from the time of the maximum. But for situations with more irregular ISI or ICI components, the ideal location can deviate somewhat more from this maximum.

The wanted value at the sample-instant is often called the ‘cursor’ sample when it is discussed in relation to the unwanted pre-cursor or post-cursor ISI samples [99, 100]. Post-cursor ISI is the most common (at least for on-chip communication) and is caused by non-zero levels of the symbol response at multiples of the symbol-time ($t = t_d + k \cdot T_s$, $k=1,2,3\dots$) after the cursor. But for many communication channels, it is also well possible that the symbol response is non-zero at samples before the cursor ($t = t_d - k \cdot T_s$), creating pre-cursor ISI.

With regard to ICI, all samples of the crosstalk-response samples contribute to the interference, including the sample at the cursor, as is visualized in Figure 5.5b.

Simple numerical programs can be developed that take these symbol and crosstalk responses and evaluate the amount ISI and ICI as a function of sampling instant (either worst-case or as e.g. a statistical variance). For simple binary signaling, such an analysis is also discussed in [96] (the topic is also briefly touched upon in chapter 10 of [95] in relation to non-perfect equalization).

In this project a somewhat similar numerical analysis method was developed that can be applied to a variety of signaling techniques, including multi-access signaling such as OFDM and CDMA. The numerical program is based on Matlab functions that take both symbol responses and transmitted pulses as matrices and are simple to use both for simple binary signaling, as well as for more complex situations that involve crosstalk and/or multi-access situations.

In these functions, the ideal sample instant is found by numerically evaluating the vertical eye opening for many different sample instants (sweep t_d). The eye-width can also be found by searching for the sample instances where the eye-height becomes zero.

When such an analysis is carried out for multiple data rates, then one can find the rate at which the eye-opening is just zero (a zero vertical eye-opening usually coincides with a zero horizontal opening), which defines the absolute maximum data rate that is achievable, given the chosen channel and communication approach.

This analysis was used to evaluate the merits or drawbacks of different signaling techniques for on-chip communication. Results of this evaluation are given in the next chapters, where the different signaling techniques are discussed. First, the analysis method itself is discussed in more detail below.

5.4.2 Linear models for communication systems

For the analysis, we assume that we have a number of communication channels that may have crosstalk to each other (for example a bus of wires as was shown in Figure 3.1). So, there is a transfer from transmitter TX_i to receiver RX_j , with $i=j$ is the wanted transfer and $i \neq j$ the crosstalk. Without loss of generality, we concentrate here on the signal received at channel zero: $j=0$. For further analysis, we define the following quantities: Let $s_{00}(t)$ represent the symbol response, $b_0(n)$ the symbol sequence, $b_i(t)$ the symbol sequence from an interfering channel and $s_{i0}(t)$ the crosstalk response of that interfering channel on the desired channel. With all the terms combined the input signal at the receiver $y_0(t)$ becomes:

$$y_0(t) = \sum_{i=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} b_i(k) s_{i0}(t - kT_s) \quad (5.1)$$

Note that it is assumed here that the interfering channels are clocked at the same instants as the channel under investigation, such as is the case between signals in a unidirectional bus. Other types of crosstalk - from e.g. unrelated perpendicular wires - are much more difficult to analyze deterministically (an impression of random crosstalk from perpendicular wires can be found in [2]). But this last type of crosstalk should be of much less concern in a well-designed on-chip communication system that uses e.g. a bus with (twisted) differential wires, as was discussed in section 4.4.

Equation (5.1) also implicitly assumes that we are dealing with linear modulation, which means that the shape of the transmitted pulse is independent of the source symbol, except for its phase and amplitude.

When we assume that the receiver samples once every symbol interval to detect the data, then the received sample sequence $r(n)$ becomes:

$$r_0(n) = y_0(t_d + nT_s) = \sum_{i=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} b_i(k) s_{i0}(t_d + (n-k)T_s) \quad (5.2)$$

Both equation (5.1) and (5.2) are simple convolution equations, representing a linear time-invariant system (LTI), with (5.2) being the discrete time equivalent of (5.1). The equations describe a MISO system, with Multiple Inputs and a Single Output, but it becomes a MIMO system if we are also interested in simultaneous analysis of the other channels (\dots, y_{-1}, y_1, y_2 , etc.). All signal analysis tools that are available for normal LTI systems (as described for example in [59]) can thus also be applied here. The impulse responses of this LTI system are the symbol response and crosstalk responses, which are themselves equal to transmitted pulse shape $h_{TXi}(t)$ convolved with the channel-responses $h_{CHij}(t)$ and possibly also some receiver filter $h_{RXj}(t)$:

$$s_{ij} = h_{TXi} * h_{CHij} * h_{RXj} \quad (5.3)$$

For some communication systems, the receiver filter h_{RX} can just model the transfer of the linear parts of the receiver circuits, or it might be used to model a simple matched filter such as an integrated-and-dump filter [92, 94, 95]. This filter-response can however also be used to model more complex situations, such as a channel-filter in an OFDM receiver (with filter-bank models for the FFT's and IFFT's). This will be explained in more detail later on, when discussing complex modulation schemes in section 5.4.4 (An example of such a system and its corresponding linear model is also shown there in Figure 5.7).

The transmitted symbols are chosen from a source alphabet of scalar (possibly complex) values:

$$b_i(n) = \{a_1, a_2, \dots, a_M\} \quad (5.4)$$

All the alphabet elements together form the constellation for M-ary signaling [92, 94, 95], for example -1 and 1 for bipolar binary signaling. To keep further analysis simple, we assume that we are dealing with a discrete memoryless source (DMS) and memoryless modulation, without correlation between subsequent symbols (so each letter from the alphabet can occur with a certain fixed probability $P(b=a_k)$). We hereby also effectively assume that no error-correction coding (such as a convolutional code) is used and that a simple hard-decision detector will be optimal at the receiver [95]. It might well be that error-coding and soft-decision can be included in the analysis, but that is beyond the scope of this thesis, as it is not very practical for on-chip communication due to the complexity of the operations. An exception to this assumption is any form of LTI memory ($b'(n) = \alpha_1 b'(n-1) + \alpha_2 b'(n-2) + \dots + \beta_0 b(n) + \beta_1 b(n-1) + \dots$). LTI memory is equivalent to linear filtering, which can simply be included in the symbol response (as will be done in Chapter 7, when discussing equalization). We also assume that the constellation is symmetric around zero (e.g. bipolar signaling), but this is just for simplicity and without loss of generality.

The detector has to recover the original symbol sequence from the received symbol sequence. When there would be no ISI and ICI then the received symbol-sequence would simply be:

$$y_{\text{wanted}}(t_d + nT_s) = b_0(n)s_{00}(t_d) \quad (5.5)$$

All the other components of the received signal are part of the ISI and ICI, which can be separated into the following equations:

$$e(t_d + nT_s) = ISI(t_d + nT_s) + ICI(t_d + nT_s) \quad (5.6)$$

$$ISI(t_d + nT_s) = \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} b_0(k)s_{00}(t_d + (n-k)T_s) \quad (5.7)$$

$$ICI(t_d + nT_s) = \sum_{\substack{i=-\infty \\ i \neq 0}}^{\infty} \sum_{k=-\infty}^{\infty} b_i(k)s_{i0}(t_d + (n-k)T_s) \quad (5.8)$$

5.4.3 Maximum interference and eye openings

The ISI and ICI will start to create bit-errors when their magnitude becomes larger than the detector threshold allows. For normal signaling without memory, optimum detector thresholds (V_T) are spaced midway between signal levels [94]. The distance of the thresholds from the signal levels (y_{wanted}) can be found by computing the Euclidian distance $d(a_i, a_j)$ between a symbol from the alphabet and its direct neighbors [95]. When we are dealing with a transmission scheme with equidistant constellation points, or if we are only interested in the possibilities of bit-errors and not for which symbol they can occur, then we can just use the minimum distance between any of two points in the constellation as a measure for the error-threshold:

$$d_{\min} = \min \left(\|a_i - a_j\| \right)_{\substack{j=1..M \\ i \neq j}} \quad (5.9)$$

So, error-free detection in the presence of interference is possible when the maximum of the interference ($\hat{e} = \max(e)$) is smaller than half the minimum distance, which at the receiver side becomes:

$$\hat{e}(t_d) < s_{00}(t_d) \frac{1}{2} d_{\min} \quad (5.10)$$

From equations (5.6) to (5.8), we can compute the maximum interference by taking the worst-case situation for every variable in the equations. The result will no longer be dependent on the current sample n , but does depend on the sampling delay (t_d):

$$\hat{e}(t_d) = ISI_{\max}(t_d) + ICI_{\max}(t_d) \quad (5.11)$$

$$ISI_{\max}(t_d) = \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} \|b_0\|_{\infty} |s_{00}(t_d + kT_s)| \quad (5.12)$$

$$ICI_{\max}(t_d) = \sum_{\substack{i=-\infty \\ i \neq 0}}^{\infty} \sum_{k=-\infty}^{\infty} \|b_i\|_{\infty} |s_{i0}(t_d + kT_s)| \quad (5.13)$$

In the equations, the $\|b_i\|_{\infty}$ is the maximum of the absolute value of b_i (the so-called ‘infinity norm’), as the largest symbol from the alphabet will give the largest error. For constellations that are symmetric around zero, this value will be equal to half the Euclidian distance between the two most far apart symbols in the constellation. When we assume that every channel uses the same constellation, then we can use the following substitution for $\|b_i\|_{\infty}$:

$$\|b_i\|_{\infty} = \frac{1}{2} d_{\max} = \frac{1}{2} \max_{\substack{i=1..M \\ j=1..M \\ i \neq j}} \|a_i - a_j\| \quad (5.14)$$

With which equation (5.11) simplifies to:

$$\hat{e}(t_d) = \hat{s}_e(t_d) \frac{1}{2} d_{\max} \quad (5.15)$$

$$\text{with } \hat{s}_e(t_d) = \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} |s_{00}(t_d + kT_s)| + \sum_{\substack{i=-\infty \\ i \neq 0}}^{\infty} \sum_{k=-\infty}^{\infty} |s_{i0}(t_d + kT_s)| \quad (5.16)$$

These equations, together with equation (5.10) can be used to evaluate the suitability of communication schemes for band limited channels. Theoretically, k should run from minus infinity to infinity, but for numerical evaluation, only the few most dominant ISI contributors have to be evaluated for an acceptable result (e.g. $k=-1..6$ for the example in Figure 5.5).

Results of the evaluation predict if a communication scheme is error-free or not and how much margin (m) there is at a given data rate:

$$m(t_d) = |s_{00}(t_d)| \frac{1}{2} d_{\min} - \hat{s}_e(t_d) \frac{1}{2} d_{\max} \quad (5.17)$$

$$\text{eye-height} = \max(m(t_d))_{t_d=0..\infty} \quad (5.18)$$

The eye-height is the maximum of m , found at the ideal sample-instant t_{ideal} . Of course, when the eye-height is smaller than zero, then the eye is completely closed.

The eye-width (T_w) is the distance between the zero-crossings of $m(t_d)$:

if $m(t_{edge_1}) = 0$, $m(t_{edge_2}) = 0$ and $m(t_d) > 0$ between $t_{edge_1} < t_{ideal} < t_{edge_2}$ then

$$\text{eye-width} = t_{edge_1} - t_{edge_2} \quad (5.19)$$

Similar quantities that were used extensively during this project are relative eye-openings, which allow a more general (and normalized) analysis of communication systems. The relative eye-width is defined as the eye-width divided by the symbol period. The relative eye-height is defined as the eye-height divided by the maximum received value at the ideal sampling instant ($t_d=t_{ideal}$) and its value is an indication for the required dynamic range at the receiver (where the absolute eye-height is an indication for the amount of noise and offset allowed):

$$rel. \ eye - height = \frac{m(t_d)}{|s_{00}(t_d)| \frac{1}{2} d_{max} + \hat{s}_e(t_d) \frac{1}{2} d_{max}} = \frac{\frac{s_{00}(t_d)}{\hat{s}_e(t_d)} \cdot \frac{d_{min} - 1}{d_{max}}}{\left| \frac{s_{00}(t_d)}{\hat{s}_e(t_d)} \right| + 1} \quad (5.20)$$

The time-instant where the relative eye-height is maximized is usually equal to the time-instant t_{ideal} where $m(t_d)$ is the highest, but not always. For some type of channels – or with equalization included – one can optimize for either highest absolute or highest relative eye-height.

From here on, the analysis has to differentiate between different signaling methods, as the type of constellation defines for example the relation between d_{min} and d_{max} (see also [95], section 4-3-1 for a discussion of different constellations and their Euclidian distances).

Real constellations (PAM)

Pulse amplitude modulation (PAM) is the simplest form of modulation, using M-ary signaling with real, equidistant constellation points (spread out symmetrically around zero). In this case the relation between d_{max} and d_{min} is straightforward:

$$d_{min} = \frac{d_{max}}{(M-1)} \quad (5.21)$$

For PAM, the error-condition from (5.10) can be rewritten, using (5.15) and (5.16):

$$\sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} |s_{00}(t_d + kT_s)| + \sum_{i=-\infty}^{\infty} \sum_{\substack{k=-\infty \\ i \neq 0}}^{\infty} |s_{i0}(t_d + kT_s)| < \frac{s_{00}(t_d)}{(M-1)} \quad (5.22)$$

This result basically states that to avoid detection-errors, all the ISI and ICI contributions in the symbol responses s_{i0} should together be smaller than the desired value $s_{00}(t_d)$ divided by the number of detector thresholds for the constellation ($M-1$).

Complex constellations

Another popular modulation category for many applications is band-pass modulation with complex alphabets. In that case, the symbols from the alphabet get mapped to cosine (real or in-phase component I) and sine (imaginary or quadrature component Q) components [92, 94, 95].




Figure 5.6: Received samples with ISI/ICI for a PSK modulated signal (a) and a QAM signal (b). For visual simplicity, the constellations are not rotated ($\text{angle}(s_{00}(t_d))=0$).

With complex modulation, the analysis also becomes a bit more complicated as the interfering signal (e) will become an error vector instead of a scalar, as illustrated in Figure 5.6. The received signals can no longer be captured easily in eye diagrams, but a scatter diagram that shows y at a certain sampling instant t_d , as is for example done in Figure 5.8 below, is now an often used substitute [95].

The error vector can have many angles, depending on the angle of previous symbols and the angles in the symbol response (as s_{i0} most likely becomes complex as well). It will still have a maximum value (\hat{e}), but that value will not necessarily be reached in every direction. For complete information one would need to compute the maximum as a function of the angle $\hat{e}(\theta)$, but that would be quite cumbersome. For an approximation of the constellation properties, the interference can be approximated as an error vector that is the same in every direction (as shown schematically in Figure 5.6b).

Note that the concepts of error-vectors and minimum or maximum constellation distances are also often used to quantify the performance of measured or simulated constellation data. The ‘error-vector magnitude’ ($\text{EVM} = \text{rms}(e)/\frac{1}{2}d_{\max} \cdot 100\%$) is for example an often-used performance figure [101].

PSK

For constellations with a circular shape, such as PSK [92, 94, 95], we can reasonably assume that a single value for \hat{e} is adequate to predict the error threshold from (5.10) with high precision, because e will have the same magnitude in nearly every direction (at least in as many directions as there are angles in the constellation points), as is shown in Figure 5.6a.

For normalized PSK with the absolute value of the symbols equal to 1, the distance between two points is [95]:

$$d(a_1 - a_n) = \left| 1 - e^{j\frac{n}{M}2\pi} \right| = \sqrt{2 - 2\cos\left(2\pi\frac{n}{M}\right)} = 2\sin\left(\pi\frac{n}{M}\right) \quad (5.23)$$

When M is even, then d_{\max} is 2, so:

$$d_{\min} = d_{\max} \sin\left(\frac{\pi}{M}\right) \quad \text{for even } M \quad (5.24)$$

For $M=3$, $d_{\min}=d_{\max}$, but for larger M , the formula above is also a good approximation for odd M .

For $M=2$, the results are equal to PAM, but if M increases, the ratio of d_{\min} over d_{\max} is higher for PSK than for PAM (which is desirable for a good eye-opening), because PSK exploits the additional orthogonal quadrature component.

QAM

For QAM modulated signals with a rectangular constellation (see Figure 5.6b), it can be that \hat{e} has a rectangular part in its contour, especially if the ISI of the previous symbol $s_{00}(t_d-T_s)$ heavily dominates over other ISI/ICI components (the rectangular contour will be rotated by the angle of $s_{00}(t_d-T_s)$). A similar argument holds for complex constellations with other non-circular outer boundaries. But for most of these constellations, a single value for \hat{e} should still be adequate, with only minor deviations from the real error threshold.

For M -ary rectangular QAM, with $K=\sqrt{M}$ points in each direction, the relation between d_{\min} and d_{\max} is:

$$d_{\min} = \frac{d_{\max}}{\sqrt{2}(K-1)} \quad (5.25)$$

Between diagonal neighbors, the distance is $\sqrt{2}$ times larger. If the worst-case interference happens to be directed along this diagonal – because the ISI is dominated by one component that has a rotation of $k-45^\circ$ compared to $s_{00}(t_d)$ – then the error-threshold will also become $\sqrt{2}$ times larger. That is however not a likely situation. In practice, the received constellation will at least have a somewhat round scatter, as in Figure 5.6b, due to other error signals such as noise or the non-dominant ISI components.

When compared to PSK, QAM becomes better with respect to d_{\min}/d_{\max} for K is three or higher, meaning $M = 9$ or higher. At $M=16$ ($K=4$), QAM has a d_{\min}/d_{\max} of about 0.24, PSK has a d_{\min}/d_{\max} of 0.195 and PAM has a ratio of only 0.067. The better distance ratios of the more complex modulation types are good reasons to use them in e.g. radio applications..

An important drawback to use complex modulation in on-chip communication (next to the complexity to implement) is the prerequisite that the modulated signals have to be strictly band limited. When the signals are not strictly band limited, then mirror-frequencies can enter the band of interest and interfere with the wanted signal (thereby increasing the error signal e). To analyze the effect of such interference, it is not possible to model the modulated signal as one complex signal, but the I and Q components have to be analyzed separately, as will be discussed below.

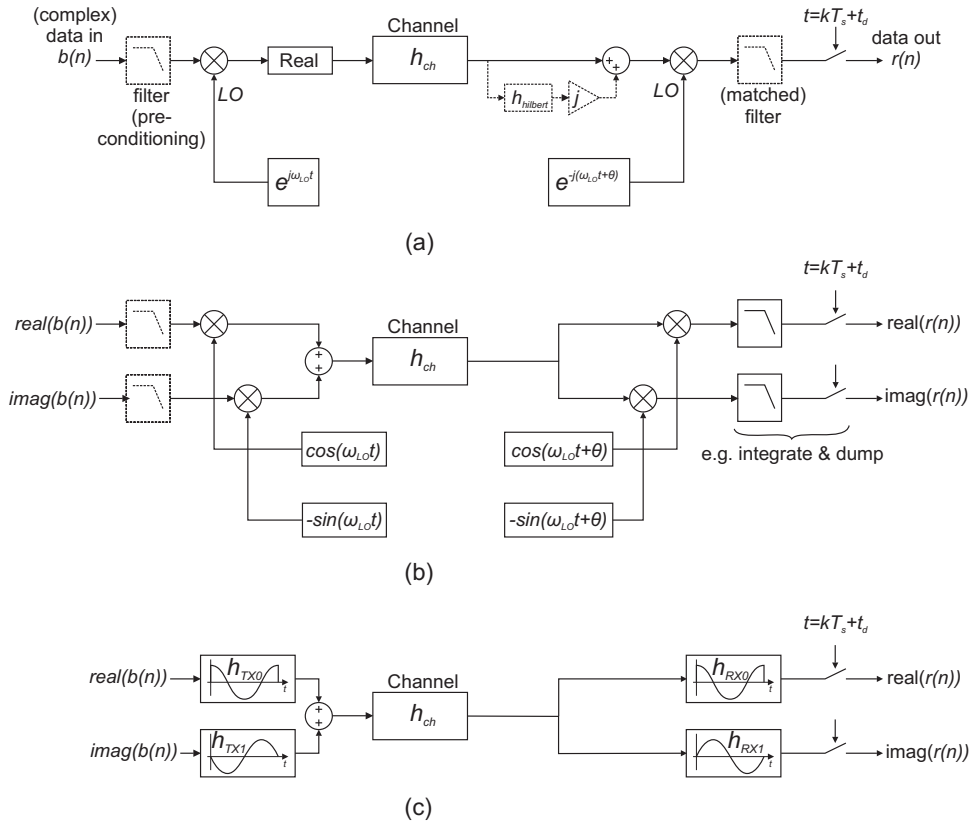


Figure 5.7: Band-pass transceiver models. (a) General complex signal model, (b) practical (direct downconversion) receiver. (c) linear model for analysis.

5.4.4 Complex signal analysis versus separation of I and Q

To evaluate whether complex constellations such as PSK and QAM have merits for on-chip communication, we have to analyze band-pass transceiver models (as a baseband transceiver can only transmit real signals)

A general band-pass transceiver model with complex signals is shown in Figure 5.7a. It is shown in for example [95] that any bandlimited band-pass signal can be described by a complex envelope that modulates a carrier signal, which forms the basis for the complex signal analysis (and complex constellations) as used in textbooks. Taking the real part of the signal just before the channel and using the hilbert transform after it, as shown in the figure, are the means to move from a complex single-sided spectrum to a real spectrum (that is symmetric around zero) and back again [95]. These two operations are completely transparent (and can even be omitted for analysis) as long as the modulated input signal has a strictly single-sided spectrum, which means that the data signal before the multiplication with the carrier has to be strictly band limited.

However, band-limiting the input signal with sufficient attenuation of the aliases requires quite complicated filters (such as a raised cosine filter), which is not desirable for on-chip

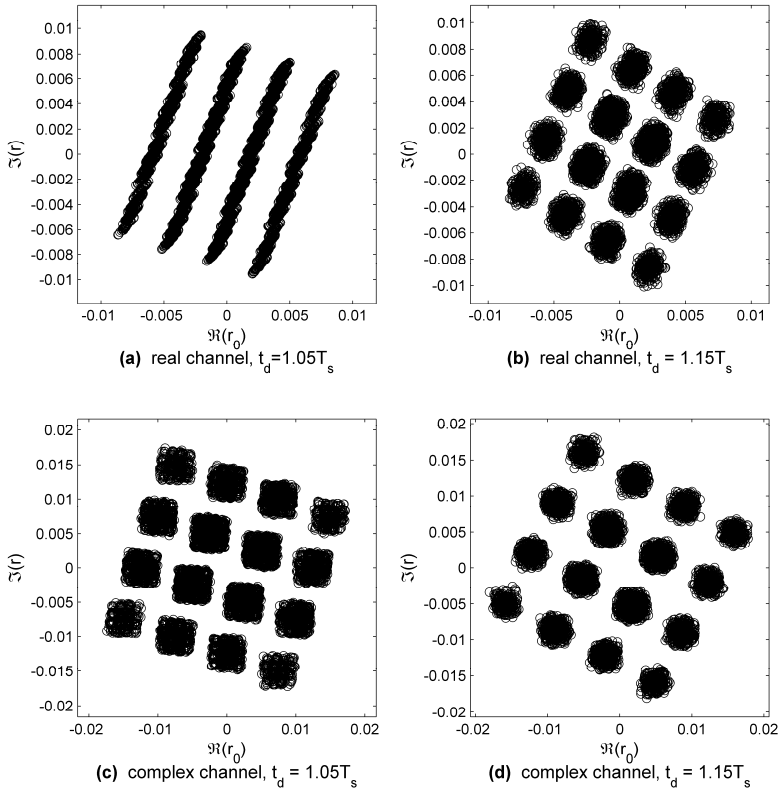


Figure 5.8: Constellation examples for 16-QAM signaling over an on-chip wire (a,b) and its conceptual complex counterpart (c,d), with $f_s=3$ symbols/ $R_{wire}C_{wire}$ and $f_{LO}=3f_s$.

transceivers. But with filters that are simple to implement, such as zero-order hold (which will generate rectangular pulses), the input signal will still contain much sidelobe energy. When such signals are modulated with the carrier, then the sidelobes from the mirror image can mix with the wanted signal itself and the system can no longer be modelled with linear operations on a complex signal.

In those cases, the IQ model is needed that is shown in Figure 5.7b. This is also a model that better resembles actual implementations of (direct conversion) band-pass transceivers. This MIMO system can linearly analyze input signals that are not strictly bandlimited.

Mirror images of the signal will manifest themselves as (not necessarily symmetric) crosstalk between the I and Q channel. The major difference with the complex model is that a complex s_{00} coefficient can only rotate or scale the constellation, whereas an IQ MIMO model can also model separate scaling of the I or Q axis and skewing of the constellation (e.g. a square becomes a trapezoid).

These effects can be quite significant for on-chip band-pass communication, as is for example visible in the constellation examples in Figure 5.8. In this figure, QAM

constellations (at the receiver-side of the model from Figure 5.7b) are shown for two different receiver sample times in (a) and (b). Clearly, (b) has the more ideal sample time, as (a) suffers from ISI that is quite severe in the I-direction. Subfigure (c) and (d) show how the constellations would look like when the channel would accept complex signals (the model from Figure 5.7a without the ‘real’ and ‘hilbert transform’ blocks). As discussed earlier, that would be a functionally correct model when the band-pass signals would be strictly bandlimited. The constellations in (c) and (d) not only remain rectangular, as opposed to the skewed constellation in (b), but the shape of the ISI/ICI is also less influenced by the timing of the receiver. The result is that the data is still (just) detectable in (c), while it is clearly not in (a).

Proper detection of the data in such IQ systems thus becomes more complicated: for strictly bandlimited signals, a single complex coefficient can be used before the detector to scale the constellation back to its proper angle and size. For the IQ MIMO system, a matrix of 2×2 coefficients is needed to reshape the constellation before the data is detected/discriminated at the receiver.

For systems with uncorrelated I and Q data signals (such as a rectangular QAM constellation), a separate treatment of the I and Q components as different (virtual) channels does have an advantage for the interference analysis: It eliminates the whole angular independence approximation for the error vector \hat{e} and just yields two scalar values, one for the I and one for the Q channel. So it can predict with good accuracy, what the eye opening will be at the input of the I and Q detectors. Note that the error components for the I and Q detector are not necessarily independent (due to the crosstalk between the channels), but that does not matter as long as the detectors are separate.

A functionally equivalent model of Figure 5.7b which contains only filtering and sample actions and is suitable for the analysis in this chapter, is shown in Figure 5.7c. At the transmitter, filtering an impulse train of data symbols with a sine wave is equivalent to multiplication of a PAM signal with the same sine wave. At the receiver, demodulation and filtering with an integrate and dump filter is again equivalent to the convolution (filtering) of the data with the sine wave (but now with the time reversed). The same holds for other periodic modulator signals as convolution is equivalent to multiplication and integration with the second variable time reversed (see e.g. [95]). So, for such modulated systems, it is still possible to construct symbol responses as in (5.3). One thing that has to be taken into account is that the detection instant t_d in the linear model represents both the phase of the receiver demodulation signal ($\theta = 2\pi t_d/T_s$) as well as the phase of the sampler itself. To analyse these two parameters independently, different receiver responses h_{RX} have to be used for different $\theta T_s/2\pi - t_d$ values.

5.4.5 Statistical analysis

In the analysis above, the worst-case interference \hat{e} , also called the ‘peak distortion’ ([95] section 10-2-1), was analyzed. However, in some transmission schemes, it can be that the probability for a bit-error is quite small, despite a worst-case interference that is higher than the error-threshold.




Figure 5.9: Histogram of 10000 received samples with binary signaling for either a single active wire (a) as in Figure 5.3 or with a wire in a bus (b) as in Figure 5.4.

In Figure 5.9 for example, a histogram is shown of received samples with simple binary signaling. In Figure 5.9a, only a single wire is active and the ISI creates a very distinct histogram pattern that is also very clear in the eye diagram of Figure 5.3, with small ISI bands around the dominant ISI components. In this case, when the data rate is increased, bit-errors will suddenly become abundant when the worst-case interference becomes higher than the wanted signal, or in other words when the condition in equation (5.10) is violated.

However, in cases where there are more than only a few dominant components in the interference, the edges become less well defined. In Figure 5.9b for example, ICI between wires in a bus is included, which changes the picture and creates a much smoother histogram and consequently a more gradual dependency between the occurrence of bit-errors and data rate. One would expect for example from the eye diagram in Figure 5.4 that the eye is still marginally open (with the eye-diagram based on 400 received symbols). The histogram in Figure 5.9b however shows that there is definitely a small probability of a bit-error. To analyze such situations, statistical methods that predict the probabilities of bit-errors – the bit error ratio (BER) – can be quite useful and can provide more information than just the worst-case interference. This is especially true for situations where bit-errors are allowed, for example when error-correction codes are used.

When the probability density function (PDF) of the error samples $e(t_d+n*T_s)$ is known, then the BER can be computed by integrating those parts of the PDF that exceed the error-threshold of $\frac{1}{2}d_{\min}$. The PDF of a random variable that is composed of a superposition of multiple independent components can be found by convolving the PDFs of the individual components with each other ([91], pp. 195).

So one can find the PDF of the received samples (or of the error part of the samples) by convolving scaled copies of the PDF of the original constellation; with the scaling factors being the samples of the symbol response $s_{i0}(t_d + kT_s)$. In [96], this method was proposed for the analysis of binary signaling in the presence of inter-symbol interference and crosstalk, using recursive repeated convolution to find the final PDF and numerically integrate this PDF to find the probability of a bit-error. Although this method is able to produce very accurate BER predictions, it is also quite involved and computationally intensive.

In this project, a simpler, more approximate statistical method was used. It makes use of the central limit theorem ([91], pp. 214), which states that the PDF of the sum of many independent random variables will approach a Gaussian shape. This effect is clearly visible in the shape of the interference Figure 5.9b.

The variance of the total error signal (σ_e^2), which is also the variance of the Gaussian approximation, equals the sum of the variances of the components because these components are independent [91]. When we assume that all the data sources b_i in the system have the same statistical characteristics ($\text{var}(b_i) = \text{var}(b)$, $\text{mean}(b_i) = \text{mean}(b)$), then based on (5.6)-(5.8), the error variance becomes:

$$\sigma_e^2(t_d) = \|s_e(t_d)\|^2 \cdot \text{var}(b) \quad (5.26)$$

$$\|s_e(t_d)\|^2 = \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} s_{00}(t_d + kT_s)^2 + \sum_{\substack{i=-\infty \\ i \neq 0}}^{\infty} \sum_{k=-\infty}^{\infty} s_{i0}(t_d + kT_s)^2 \quad (5.27)$$

$$\text{var}(b) = \sum_{i=0}^M |a_i - \text{mean}(b)|^2 p_i \quad (5.28)$$

These equations are the stochastic equivalent of the maximum error signal from (5.15).

One can make some simplifications for the computation of the variance of the constellation $\text{var}(b)$, starting for example with the assumption that the mean of the transmitted symbols $\text{mean}(b)$ is zero. A similar assumption is that all symbols in the alphabet have equal probability, meaning that $p_i = 1/M$.

Under these assumptions, the computation becomes simple for those constellations where all symbols have the same energy, such as binary signaling or PSK, where the variance of the constellation will simply be: $\text{var}(b) = (1/2 d_{\max})^2$. For multi-level constellations, the variance will be smaller than the maximum level. For 4-PAM for example $\text{var}(b) = 1/4(1+1/9+1/9+1) \cdot (1/2 d_{\max})^2 = 5/9 \cdot (1/2 d_{\max})^2$. For M-PAM signaling with higher M , the variance will approach the variance of a uniform distribution: $\text{var}(b) = 1/12 \cdot d_{\max}^2$. Variances for other constellations can be found in [95] (section 5-2), but then specified in terms of energy (which is equivalent to variance in the absence of a mean) and as a function of d_{\min} .

When the variance σ_e^2 is computed, then it can be used to approximate the probability of a symbol detection error (and subsequently the BER) with the aid of the Gaussian cumulative distribution F or the related Q-function ($F(x) = Q((\text{mean}-x)/\sigma)$). For many types of constellations, such procedures can be found in textbooks ([95] section 5-2). The equations are not always straightforward, for example because constellations such as PAM and QAM can have inner points that are completely surrounded, with a different error probability than

the points at the perimeter of the constellation. For constellations with only two points (binary signaling), the probability of an incorrectly detected symbol is straightforward:

$$P_e = Q\left(\frac{\frac{1}{2}d_{\min}}{\sigma_e}\right) \quad (5.29)$$

And in this case, the BER equals P_e as only one bit is encoded in each symbol.

When applying the BER equations, it should always be kept in mind that the error-signal due to interference is just an approximation of a Gaussian distribution. In reality, the tail of the distribution will be limited to a value given by \hat{e} and the results will not be accurate if \hat{e} approaches d_{\min} . For situations where \hat{e} is substantially higher than d_{\min} good predictions have been obtained with this analysis. For example for the CDMA analysis that will be discussed in section 6.5.3, the predicted BER deviated from the BER that was obtained with time-domain simulations by less than 1% for $\text{BER} > 10^{-1}$, less than 20% for $\text{BER} > 10^{-3}$ and still in the same order of magnitude for $\text{BER} > 1e-6$.

For the analysis of most practical on-chip communication systems, the worst-case analysis based on \hat{e} is probably preferable. It is not only simpler to use, but it is also more accurate at the boundary of zero or negligible BER. A zero BER (or at least a bit error rate lower than the maximum on-time of a device) is important for contemporary on-chip communication for which error-correction is not (yet) practical.

5.4.6 Remarks on symbol-response analysis

This section concludes the description of the symbol-response analysis with some observations that were made during its usage and some remarks on its applications and limitations.

Constellation attenuation

For baseband transmission, one might intuitively be inclined to use the transmitted levels of the constellation points times the DC gain of the channel as the wanted values at the receiver. One might assume for example that y_{wanted} be -1 or 1 for normalized bipolar binary transmission or $-1, -1/3, 1/3, 1$ for quaternary transmission. But actually, the wanted levels are attenuated by $s_{00}(t_d)$, as discussed and as is visible in Figure 5.5. In the binary example from Figure 5.3, a received level of 1 is actually an extreme case, being y_{wanted} plus maximum interference from previous positive symbols. So the high-frequency transfer of a channel has an influence on the attenuation of the constellation.

The attenuation of the constellation also has an influence on the absolute eye-height. Ideally, the eye-height should be as large as possible (to have a suitable margin at the receiver for e.g. offset and noise). However, attenuation of the y_{wanted} is not always easy to compensate. For example most forms of transmitter equalization, such as the pre-emphasis methods that will be discussed in Chapter 7, can only attenuate the low-frequency part of a non-flat channel transfer, instead of emphasizing the high-frequency part.

ICI between virtual or physical channels

For most types of multiplexed systems (OFDM, CDMA, etc), the exact same analysis can be used as the analysis of ICI from neighboring crosstalk channels, as there is no

conceptual difference between ICI from virtual or from actual channels. The aspect that is different is the shape of the different channel responses and cross-channel responses. In a wide bus with many neighboring channels, the symbol-responses are very similar ($s_{00} \approx s_{11} \approx s_{22} \dots$), but this will usually not be the case with multiplexed systems that use multiple virtual channels on a shared physical medium, such as the CDMA example that will be discussed in section 6.5.3.

Interference analysis with nonlinear modulation

The analysis method described here can be used to analyze linear types of digital communication schemes. It is less suitable for nonlinear modulation schemes such as frequency modulation (FM). Still, in some cases, one can ‘linearize’ the modulation and still apply the analysis. For binary FM, one can for example treat the two different frequencies as two separate channels, both using on-off keying and use ICI analysis to estimate the probability that a bit is detected erroneously. The results can not be used directly however, one should first account for the fact that both channels are not independent but are each others inverse.

Interference analysis and transmitter errors

An aspect that can not be directly analyzed with the method above is ISI due to timing-jitter at the transmitter. This is because such timing jitter alters the position and width of a transmitted symbol. An indication of the influence of transmitter timing jitter can be obtained by examining the sensitivity of the ISI to changes in sample instant ($dISI/dt$), but for a more accurate analysis, a more complex model is needed, as is discussed in [102]. We do not use those more complex models here because the effects of random timing jitter at the transmitter are assumed to be much lower than the effect of a finite channel bandwidth (given that a proper clock distribution network is present).

Another aspect that is not fully captured by the linear analysis are non-linear effects such as slewing in the transmitter. However, such non-linear effects can sometimes be incorporated via linearization with an additional modeling step. Slewing effects that lead to finite rise- and fall-times for example can, given plain binary signaling, be modeled with an additional filtering step, as can of course be done with the linear part of the low-pass characteristics of a transmitter. To model slewing, one can use a moving average filter to create a more realistic trapezium shape from the original idealized square wave. However, in contrary to the linear transmitter effects, slewing is not proportional to the step-size, so one can not use the filtering method to analyze slewing correctly for multi-level signaling, but it can still be used for approximations.

To fully capture non-linear behavior from a transmitter when it is used to output more than only binary levels, either due to a multi-level symbol alphabet or due to equalization, requires more elaborate extensions to the model.

Linearization of communication systems

As was also discussed in the previous two topics, the key enabling factor for the analysis as presented here is the linearization of the communication system. Many aspects of communication can be modeled with linear systems. It was for example shown that a modulated (possibly multi-carrier) signal can also be modeled with only linear elements (as

in Figure 5.7b). It is not always immediately clear how to linearize a system and not every method exploits linearization to the fullest. In [96] for example, the peak distortion analysis is differentiated with respect to the two values in the binary constellation and only the ISI that works destructively for the constellation point under analysis is counted. But, as was presented in the previous sections, ISI can also be viewed as an additive signal independent of the current symbol and the numerical analysis becomes faster and just as accurate if it is treated as such. So, for successful application of the analysis presented here, it is important to spend some effort in finding linear analogies for different parts of the communication.

5.5 Synchronization

Synchronization is the process of finding and using the correct timing information for detection of the bits at the receiver, which usually means finding the frequency of the symbol stream and finding the optimal detection instant (t_d) within a symbol. Many different synchronization methods exist and have been published. For a theoretical treatment, see for example [95].

For on-chip communication, synchronization is usually easier than with other types of data communication (such as RF, optical or wireline), as derivatives of the same global clock are often available to both the transmitter and the receiver. Perhaps large-scale digital IC's will migrate away from global skew-free clocks, but a clock that has at least the same frequency (so-called mesochronous) as the transmitted symbol stream is likely to be present at the receiving end (or otherwise, it can always be transmitted, as in a source-synchronous scheme). This reduces the synchronization process to finding the correct phase of the receiver-end clock. Again, many different possibilities exist.

The first and most simple version is to fix the receiver phase during the design, with delay circuits (buffers/inverters); this option fits best to classical digital design-styles, but it requires a priori knowledge of the channel delay.

A second, quite simple option is to use source-synchronous schemes that rely on a match between the time delay for the clock-channel and the data-channels; such a system (for a NoC application) will be discussed in section 11.5. But simple source-synchronous schemes only work when the channel has enough bandwidth to enable the transmission of a clock with a frequency of least half the data rate (such that there is one clock edge for every data symbol). When this is not the case and a clock with lower frequency is transmitted, then a clock-multiplier (e.g. PLL) is needed at the receiver.

The third, most complex option is to use a clock-data recovery (CDR) circuit at the receiver. A CDR circuit typically contains a phase-detector that controls some form of clock generator, such as a PLL or a DLL [95] to center the clock in the middle of the received eye-diagram. The phase-detector can be used to either lock on to a (possibly low-frequency) reference clock (similar to source-synchronous), or it might be used to directly derive a clock from the incoming data. For this last case a few very simple phase detectors have become popular for plain binary (NRZ) signaling such as a Hogge [103] or Alexander [104] phase detector.

Alternatively to a CDR with a PLL or DLL, oversampling can also be used to discriminate between the edges and center part of the eye diagram. Such oversampled receivers also often use variants of the Alexander phase detector.

A short overview of how CDRs are applied in the related field of wireline and backplane transmission will be given in section 7.2.5. Synchronization with CDRs might become interesting for on-chip communication someday, but is probably much too complex to be viable in the near future.

5.6 Summary and conclusions

In this chapter, high-speed data communication over bandlimited channels was discussed. It was shown how eye diagrams can be used to evaluate the quality of communication with threshold detectors, and it was briefly discussed that some form of synchronization is also necessary to set the timing of the detector.

The main part of the chapter focused on the presentation of an analysis method that can a priori predict communication properties such as the eye-opening (noise margin) and achievable data rate, without requiring time-domain system simulations. The analysis can be applied to all data communication systems that can be linearized, ranging from simple binary transmission to multi-channel band-pass communication. Crosstalk between physical (wires) or virtual channels can also be incorporated in the analysis. Either deterministic (eye-height) or stochastic (BER) metrics can be extracted with the analysis. Which of the two is more appropriate depends on the application and on the number of sources that can interfere with each other.

In the next two chapters, the analysis method will be applied to on-chip communication for a variety of signaling, modulation and equalization techniques, to investigate what types of communication schemes are most appropriate.

Chapter 6

Signaling and Modulation techniques

6.1 Introduction

This chapter uses the analysis method that was developed in the previous chapter to analyze eye properties and achievable data rates for various modulation techniques when applied to on-chip wires. The quantitative results are also summarized in Appendix B.

Normalized wire parameters will be used to express the results. The achievable data rate is for example scaled by the RC product of the wire ($R_{wire}C_{wire}$). This enables more generalized statements, independent of the line length and specific parameters of the technology. Findings from section 3.3 can be used to translate the normalized data to absolute numbers for a specific wire.

For the numerical analysis, the third-order interconnect models from section 3.8.5 (and section 4.2.2 for resistive receiver or capacitive transmitter termination) are used where possible, to keep the computation time to a minimum. Only for the analysis of wires with crosstalk, the more complex lumped RC models (as described in section 3.8.6) are used with 100 lumps per wire.

The chapter starts with an analysis of plain binary signaling over on-chip interconnects, including an analysis how crosstalk affects the achievable data rate (for both conventional and special termination) and what the improvements are with twisted differential wires. Next, in section 6.3 it is discussed how analysis of baseband communication over on-chip wires can be simplified which serves as a prelude to the discussion of multi-level signaling in section 6.4. Before the chapter is concluded, more complex band-pass signaling schemes are also shortly examined in section 6.5.

6.2 Plain binary signaling

6.2.1 Achievable data rate with and without crosstalk

This section presents some quantitative results for plain binary signaling with and without crosstalk. For this signaling type, the symbol response is simply the response of the channel to a rectangular pulse, as was shown in Figure 5.5. The analysis is carried out for different

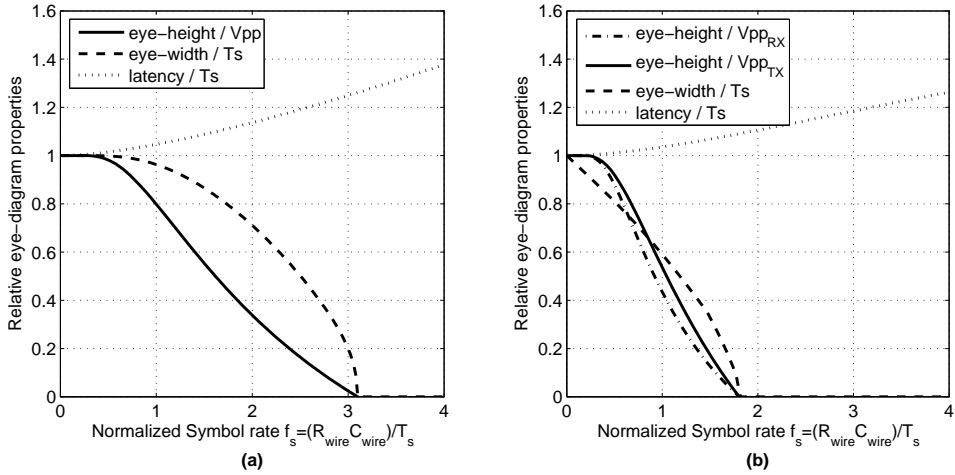


Figure 6.1: Eye properties as a function of data rate with plain binary signaling, with (a) a single, shielded wire as in Figure 5.3 or (b) an unshielded wire in a bus as in Figure 5.4. The wires have conventional termination ($R_s=0, R_l=\infty$).

types of wire termination, as the wire termination has an influence on both the speed of the response as well as on the crosstalk, as was discussed in Chapter 4.

Eye properties with conventional termination

In Figure 6.1, results of the analysis are shown for a conventionally-terminated ($R_s=0, R_l=\infty$) single-ended interconnects, both for a shielded wire (no crosstalk) and for an unshielded wire inside a bus (with crosstalk from left and right neighbor). As was explained in the introduction, the symbol rate is normalized to $R_{wire}C_{wire}$ to make it independent of the actual line-length and cross-section.

Without crosstalk, virtually no ISI effects are present below $f_s=0.4 \text{ bit}/R_{wire}C_{wire}$. At higher rates, the eye-height starts to drop, accompanied by an increase in delay and reduced eye-width. At a data rate of $f_s=3.1 \text{ bit}/R_{wire}C_{wire}$ the eye is completely closed. Note that the dominant time constant of an on-chip interconnect is $0.41 \cdot R_{wire}C_{wire}$ (as was discussed in section 3.8.5), so the achievable data rate relative to the time constant is about $f_s=1.27/\tau_{ch}$.

With crosstalk, the eye is already closed at $f_s=1.8 \text{ bit}/R_{wire}C_{wire}$ ($0.74 \cdot \tau_{ch}$), which is 42% lower than without crosstalk. An interesting visible effect of crosstalk is the rapid degradation of the eye-width compared to the situation without crosstalk. This is because the crosstalk is even more severe well before the ideal detection instant (as visible in Figure 5.5b). Note that two versions of the eye-height are shown in Figure 6.1b, one relative to the transmitter swing and one relative to the receiver swing. This is done because crosstalk not only degrades the opening of the eye, but also increases the peak-to-peak swing (when there is no crosstalk as in Figure 6.1a, then the two curves are the same). In case a linear receiver would be used to receive this signal (for example for further signal processing), then this increased swing requires additional dynamic range.

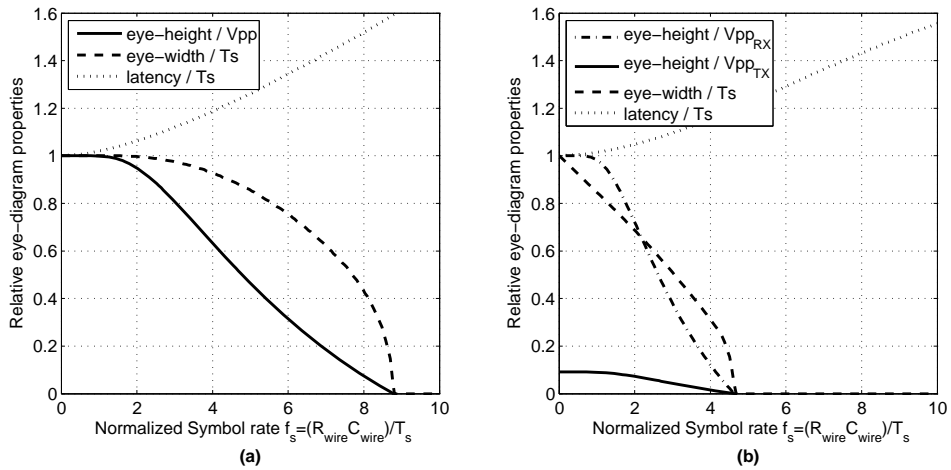


Figure 6.2: Eye properties as a function of data rate, as in Figure 6.1, but now with resistive termination ($R_s=0$, $R=R_{\text{wire}}/10$), with (a) a shielded wire or (b) a wire in a bus.

Note that Figure 6.1b also reconfirms that the eye is indeed closed at a normalized data rate of $f_s=2 \text{ bit}/R_{\text{wire}}C_{\text{wire}}$ in the presence of crosstalk, as was discussed earlier in section 5.4.5, despite the fact that it appears still marginally open in Figure 5.4.

Eye properties with resistive receiver or capacitive transmitter termination

Figure 6.2 shows results of the same analysis, but now applied to wires with resistive receiver termination (as discussed in section 4.2.2) with $R_s=0$, $R_1=R_{\text{wire}}/10$. For the single shielded wire in Figure 6.2a, the shape of the eye-properties is very much the same as in Figure 6.1a, except for a scaling of the x-axis. The achievable data rate with resistive receiver termination increase to $8.8 \text{ bit}/R_{\text{wire}}C_{\text{wire}}$. This is 2.8 times higher than with conventional termination, even slightly better than the factor 2.5 speed-up that the first-order model from equation (3.38) predicts.

Figure 6.2b shows the achievable data rate with the wires in a bus, with the reduction in the swing due to the termination clearly visible in the trace that shows the eye-height as a ratio to the transmitter swing. The figure also shows that the achievable data rate with resistive terminated wires with crosstalk is only $4.6 \text{ bit}/R_{\text{wire}}C_{\text{wire}}$. This is a reduction of 47% compared to the shielded wire, more than the reduction with conventional termination. This reconfirms the results from section 4.4.4 (in Table 4.2), that crosstalk has more impact on the performance of a resistively terminated wire than for a conventional terminated wire.

The analysis was also repeated for wires with capacitive transmitter termination ($C_s=1/10 \cdot C_{\text{wire}}$, $R_1=\infty$) and the results are shown in Figure 6.3 on the next page. The results for the shielded wire in Figure 6.3a are indeed exactly the same as for the resistive terminated version from Figure 6.2a, as was established in section 3.2.2. However, for the wires in a bus, the result is different as can be seen in Figure 6.3b. Even at DC, the eye-height is only 35% of the signal swing at the receiver, instead of the 100% for the other types of termination. This is because the capacitive transmitter termination also creates a low-frequency component for the crosstalk, as was discussed in section 4.3.1. This also

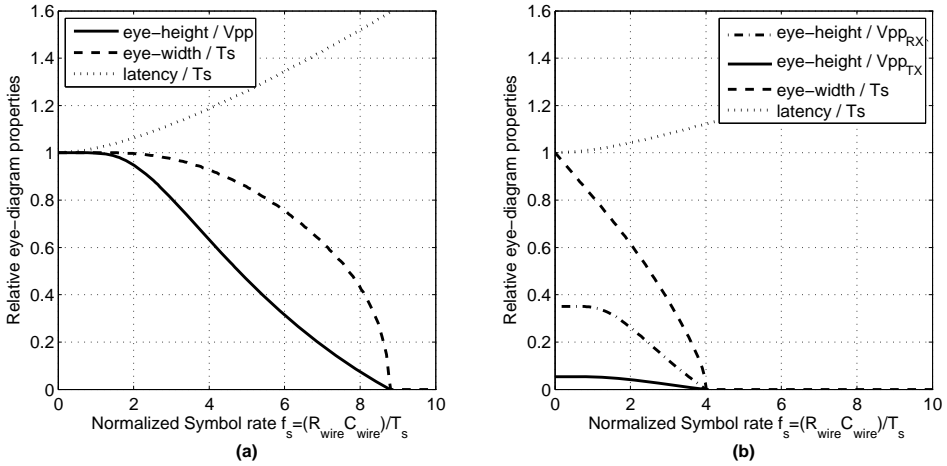


Figure 6.3: Eye properties as a function of data rate, as in Figure 6.1, but now with cap. Tx termination ($C_s = C_{wire}/10$, $R_T = \infty$), with (a) a shielded wire or (b) a wire in a bus.

results in an eye height of only 5.3% of the transmitted swing at low frequencies, instead of the 9% one would expect from $C_s/(C_s+C_{wire})$.

At high frequencies, the system with the capacitive transmitter is also slightly more susceptible to crosstalk than a system with a resistive receiver, as the eye in Figure 6.3b is already closed at a data rate of 4 bit/ $R_{wire}C_{wire}$ (versus the 4.6 bit/ $R_{wire}C_{wire}$ for the resistive receiver). For both techniques, it is clear that they only reach their full potential when the crosstalk is mitigated.

Crosstalk from non-direct neighbors

Note that the results obtained above only include crosstalk from the direct neighbours on the left and right side, which simplifies the analysis and enables a direct comparison with the results in section 5.3.3.

For conventional termination and resistive termination, the direct neighbors are also by far the most dominant source of crosstalk: when crosstalk from the next two wires is also taken into account, then the achievable rates drop to 1.7 bit/ $R_{wire}C_{wire}$ for the conventional terminated wire and to 4.4 bit/ $R_{wire}C_{wire}$ for the resistive receiver terminated wire. This is a reduction of only 5% compared to the situation with only direct neighbor crosstalk.

However, for the wire with the capacitive transmitter, the achievable rate drops by another 18%, to 3.3 bit/ $R_{wire}C_{wire}$ when the crosstalk from the second neighbors are taken into account. So with capacitive transmitters, the crosstalk is spread out over more wires than with the other types of termination, making it even more important to properly mitigate its effects by using e.g. twisted differential wires.

6.2.2 Achievable data rate with differential twisted wires

The plots without crosstalk in Figure 6.1a - Figure 6.3a were made by simulating a single wire and not a bus, but the results are also valid for twisted differential wires, as they have

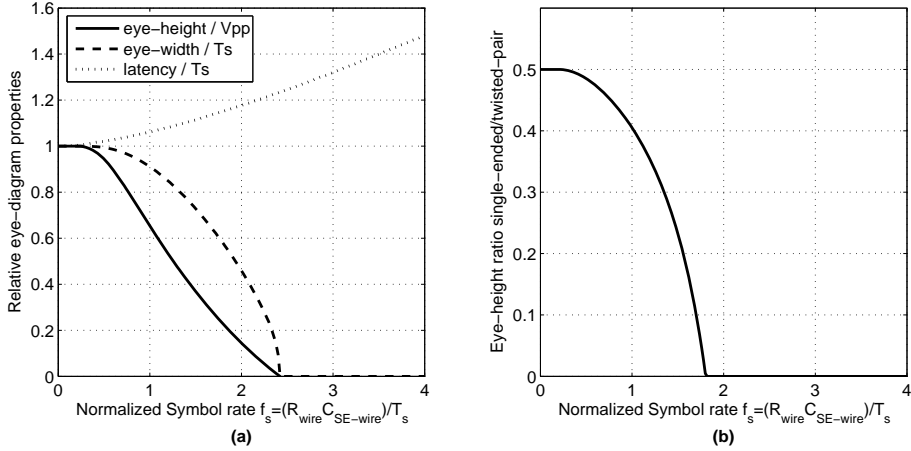


Figure 6.4: (a) Eye properties for a differential pair in a bus with $R_S=0$, $R_L=\infty$ and a twist at 70%. (b) Ratio of the single-ended eye-height to the differential eye-height.

the same response-shape and (virtually) no crosstalk. The main difference is that the response of the twisted wire is slower because its capacitance is 1.25 times higher due to the Miller multiplication of the mutual capacitance (see section 4.4.4). This assumed similarity was verified by analyzing the eye diagram properties of a twisted differential pair in a bus with the same physical properties as the wires in the analysis above.

For the case with conventional termination, the results of this analysis are shown in Figure 6.4, using a differential pair with a single twist at 70% of the length (the optimal position for conventional termination [81]). The eye is closed at $f_s = 2.4 \text{ bit}/R_{\text{wire}} C_{\text{wire-SE}}$ ($C_{\text{wire-SE}}$ denotes the capacitance of one single-ended half). This is a factor 1.29 lower than the achievable data rate for an isolated wire, only slightly more than the factor 1.25 increase from the mutual capacitance. This remaining difference is not caused by crosstalk, which is insignificant with the twist, but due to small differences between the simulated response of a single wire (with its one-dimensional lumped model) and of a bus (with its two dimensional model).

The analysis was also repeated for resistive receiver termination (with $R_S=0$, $R_I=R_{\text{wire}}/10$, as was also used in the previous section), now with a twist at 50%. For this case, the achievable data rate is $6.8 \text{ bit}/R_{\text{wire}} C_{\text{wire-SE}}$, again a factor 1.29 lower than for an isolated wire.

For the twisted wire with a capacitive transmitter, the twist was also placed at 50% of the wire length. With $C_S=C_{\text{wire-SE}}/10$, the achievable data rate becomes $7.24 \text{ bit}/R_{\text{wire}} C_{\text{wire-SE}}$. However, in this case the swing is lower than for the resistive receiver, as the effective wire capacitance of the differential wire is increased (also see section 4.4.4). The transmitter capacitance was updated to match the increase: $C_S = 1.25 \cdot C_{\text{wire-SE}}/10$ ($R_I=\infty$). The achievable data rate for this case is $7.0 \text{ bit}/R_{\text{wire}} C_{\text{wire-SE}}$, a factor 1.26 lower than for an isolated wire. This is slightly better than with resistive termination because the step response with capacitive transmitter termination, when used in a twisted differential bus, has slightly faster settling in the tail (as was also discussed in section 4.4.4).

But, small differences aside, it can be said that both the capacitive transmitter and resistive receiver variants again have very similar achievable data rates with the twisted wires. So the twisting helps to significantly increase the achievable data rate for all types of termination. Whether, or more precisely when, this increased performance outweighs the increase in costs for the differential wires is discussed next.

Break-even data rate for differential twisted wires

The above analysis showed that the achievable data rate of a twisted differential wire is between 41% and 112% higher than an unshielded wire in a bus with crosstalk from its direct and indirect neighbors. It is 41% higher for conventional termination, 55% for resistive receiver termination and even 112% for capacitive transmitter termination (also see Appendix B for the results).

One could argue that from an aggregate data rate/area viewpoint (as described in section 2.7), twisted differential wires would only be viable if they would enable more than twice the data rate than a single-ended wire as the latter occupies only half the area. The same holds for shielding, but shielding is less robust than differential signaling, as was discussed in section 4.4.1. So for the capacitive transmitter termination, the benefit of 112% data rate increase outweighs the cost of the doubling of the area, but for conventional or resistive termination, this is not the case.

However, area is not the only cost factor, power is another one. The most straightforward implementation of a differential transmitter is to simply add a duplicate of the original single-ended transmitter but then with an inverted data input (pseudo differential). This automatically doubles the power of the circuit (even more because of the miller multiplication of the mutual wire capacitance), but it also doubles the swing at the receiver as $V_{RXdiff}=V_+-V_-$.

For a fair comparison from a receiver noise-margin (and reliability) perspective, one should scale the swings until the same eye-height at the receiver is obtained for both the single-ended and the twisted interconnect. Note that for the sake of simplicity, we assume that for the single-ended situation a good reference voltage V_{ref} is available at the receiver, which defines the decision threshold, such that $V_{RXse}=V_+-V_{ref}$ (such a receiver is also known as a pseudo-differential receiver [85]). Figure 6.4b shows the ratio of the eye height for a single-ended unshielded wire in a bus to the eye height of a twisted differential wire pair, with conventional termination and with both transmitters using the same swing. The figure effectively shows how much the swing of the differential transceiver can be reduced to get the same eye-height as the single-ended transceiver.

Scaling down the swing saves power, with the power ideally being proportional to the swing squared as it takes CV^2 to charge an interconnect. Under this ideal scaling assumption, the differential transceiver is always beneficial, as it has a twice higher swing at only 2.5 times the power consumption (2 times because the two interconnects and an additional 1.25 times due to the higher capacitance).

However, for simple driver-circuits, the power consumption is more likely to be linearly dependent on the swing when we assume that the supply is fixed (as the charge current is linearly proportional to the swing). In this case, the differential transceiver becomes favorable at data rates that exceed $f_s=1.02 \text{ bit}/R_{wire}C_{wire-SE}$. Above this rate, the eye-height/swing becomes more than 2.5 times higher for the differential interconnect, or

inversely said, the single-ended eye-height is less than 40% of the differential one, as visible in Figure 6.4b.

For the resistive receiver termination, this break-even point can be found at $2.7 \text{ bit}/R_{\text{wire}}C_{\text{wire-SE}}$, but the power consumption contains more components in such a receiver, making the analysis more difficult. For a capacitive transmitter, the eye opening of the single-ended wire is always less than 40% of the eye height of the twisted differential version because of the crosstalk at low frequencies, but here a parallel G_m/R_L path could reduce crosstalk at low frequencies.

In the end, the best conclusion that can be drawn from this simple power model is that twisted differential interconnects can be more power efficient than single-ended unshielded alternatives at data rates above about one $\text{bit}/R_{\text{wire}}C_{\text{wire}}$ (assuming conventional termination).

In section 11.3, the power and optimal swing are calculated for a practical (capacitive pre-emphasis) transceiver for a NoC system, using a more elaborate model that also includes receiver power. The results found there support the conclusion that differential interconnects are favorable at high data rates.

When we would include other benefits of differential signaling, such as the robustness to (common-mode) crosstalk from other metal layers, then the break-even point shifts to lower data rates, but that is more difficult to quantify. For this project in which achievable data rate is a central issue the benefit is already clear, so for the subsequent analysis of other signaling methods, crosstalk will no longer be regarded, but it will be assumed cancelled by the twisted interconnects.

6.3 Analysis simplifications for baseband signaling

The results from the previous section are accurate but also require quite some numerical analysis to obtain. The fact that we can disregard crosstalk with twisted interconnects as discussed above, enables additional simplifications for the analysis of baseband PAM signals in combination with on-chip interconnects. The symbol response of an on-chip wire is usually entirely positive (at least for most common termination impedances), which helps to quickly find the ideal sample point t_{ideal} and relate the desired sample $s_{00}(t_{\text{ideal}})$ to the amount of ISI $\hat{s}_e(t_{\text{ideal}})$, as will be discussed below.

All LTI systems that have only real poles and no zeros (as is the case for the parametric interconnect models that were discussed in section 3.8.5 and section 4.2.2) will have a purely positive response to a positive input. As each individual pole has a purely positive exponential decaying impulse response, the cascade of the poles will also have a positive impulse response. That the symbol response of an on-chip wire to a PAM symbol is indeed entirely positive is for example visible in Figure 5.5a.

Because the symbol response of the wires is entirely positive, all ISI contributions will be positive and the absolute value operation from equation (5.16) can be omitted. So the maximum ISI is simply the sum of all the samples in the symbol response, except for the wanted sample $s_{00}(t_d)$. When the wanted sample is also added, then we get:

$$s_{dc}(t_d) = s_{00}(t_d) + \hat{s}_e(t_d) = \sum_{k=-\infty}^{\infty} s_{00}(t_d + kT_s) \quad \text{if } \forall_k s_{00}(t_d + kT_s) > 0 \quad (6.1)$$

This sum over all samples s_{dc} is the DC-transfer of the sampled response [59]. For normal PAM, with a pulse that has a width equal to T_s the s_{dc} equals the DC-transfer of the underlying continuous-time channel (in Figure 5.3b for example $s_{dc}=H(dc)=1$). This is because such a PAM-pulse has no energy at multiples of the sample frequency f_s [92] so no possibilities for aliases to alter the DC-transfer of the sampled response. For pulses that have this property, the s_{dc} will be independent of t_d :

$$s_{dc}(t_d) = S_{dc} \quad \text{if} \quad S(j\omega) = 0 \quad \text{for} \quad \omega = n \frac{2\pi}{T_s}, n \neq 0 \quad (6.2)$$

Under these conditions, it is easy to determine the ideal sample instant, which is the instant where $s_{00}(t_d)$ is maximum, as that is automatically the instant where the ISI is smallest.

$$s_{00}(t_{ideal}) = s_{max} = \max(s_{00}) \quad (6.3)$$

The ISI also immediately follows from the value of s_{00} and does not have to be computed separately:

$$\hat{s}_e(t_d) = S_{dc} - s_{00}(t_d) \quad (6.4)$$

These simplifications are applied in the next sections to analyze eye diagram properties of PAM signaling. The relations above can be applied more widely, also for non-PAM signals, as long as the conditions of an entirely positive symbol response and no energy at multiples of the sample-rate hold.

6.3.1 Eye properties for PAM with first-order channel models

As on-chip interconnects have a single dominant pole, it is possible to analyze first-order systems and use the results as approximations for the actual behavior. For a first-order low-pass channel, it is not difficult to analytically compute the eye-height and eye-width. To start, we write the transmitted PAM pulse shape as a summation of two step responses:

$$g_0(t) = \text{step}(t) - \text{step}(t - T_s) \quad (6.5)$$

The step response terms of a first-order channel are simple exponential functions (valid from the start of the step), so the symbol response can be written as:

$$s_{00}(t) = g_0 * h_{00} = \left(1 - e^{-\frac{t}{\tau_{ch}}} \right) \Big|_{t \geq 0} - \left(1 - e^{-\frac{-t+T_s}{\tau_{ch}}} \right) \Big|_{t \geq T_s} \quad (6.6)$$

This symbol response is shown in Figure 6.5a. The ideal sample instant is found at T_s , as that is where the symbol response is maximum (s_{max}), as discussed in section 6.3 above. The value of s_{max} follows directly from the equation above:

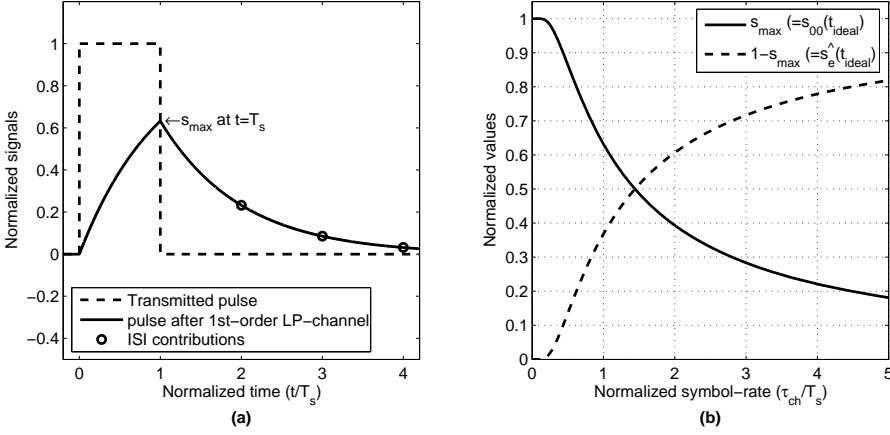


Figure 6.5: (a) Symbol response for PAM signals for a first-order channel model with $T_s = \tau_{ch}$ and (b) eye-properties as a function of the normalized symbol-rate.

$$s_{\max} = s_{00}(T_s) = \left(1 - e^{-\frac{T_s}{\tau_{ch}}}\right) \quad (6.7)$$

The DC transfer of the symbol response is always one ($S_{dc} = 1$, which can also be proven by evaluating the infinite sum $s_{00}(t_d + k \cdot T_s)$, $k = -\infty \dots \infty$). Using this and (6.4), it follows that:

$$\hat{s}_e = 1 - s_{\max} = e^{-\frac{T_s}{\tau_{ch}}} \quad (6.8)$$

The values of these variables are plotted in Figure 6.5b, as a function of the normalized sample-rate τ_{ch}/T_s .

The s_{\max} and \hat{s}_e define the eye-height and relative eye-height as defined in equations (5.17) to (5.20). For PAM, the absolute and relative eye-height differ by only a scaling factor as the maximum received value $s_{\max} + \hat{s}_e$ is always 1. With the substitution of (6.8) and some reworking, this yields:

$$rel. \text{ eye-height} = s_{\max} \frac{d_{\min}}{d_{\max}} - \hat{s}_e = \frac{d_{\min}}{d_{\max}} - e^{-\frac{T_s}{\tau_{ch}}} \left(1 + \frac{d_{\min}}{d_{\max}}\right) \quad (6.9)$$

For M-PAM, d_{\max} equals $M-1$ times d_{\min} , as stated in eq. (5.21):

$$rel. \text{ eye-height} = \frac{1}{M-1} - e^{-\frac{T_s}{\tau_{ch}}} \left(\frac{M}{M-1}\right) \quad (6.10)$$

The first term in the equation is the relative eye-height when there would be no ISI, which decreases proportionally to the number of levels (minus one) in the constellation. The second term is caused by the ISI. That this term is also slightly dependent on M (for small M) is because it accounts for the scaling of the constellation points due to ISI.

An equation can also be formulated for the eye-width from (5.19). The edges of the eye are found at the time instants where the eye-margin m is just zero. Substituting $m=0$ in (5.17), it follows that:

$$m(t_{edge}) = 0 \rightarrow s_{00}(t_{edge}) \frac{d_{\min}}{d_{\max}} = \hat{s}_e(t_{edge}) \quad (6.11)$$

Using again (5.21) and (6.4) with $S_{dc}=1$, we can rewrite this to:

$$s_{00}(t_{edge}) \frac{1}{M-1} = 1 - s_{00}(t_{edge}) \Leftrightarrow s_{00}(t_{edge}) \frac{M}{M-1} = 1 \Leftrightarrow s_{00}(t_{edge}) = 1 - \frac{1}{M} \quad (6.12)$$

In case the eye is open, there will be two solutions for this equation, at both sides of the ideal sample instant T_s with $t_{edge1} < T_s < t_{edge2}$. With the definition of $s_{00}(t)$ from (6.6) it follows that:

$$s_{00}(t_{edge1}) = 1 - e^{-\frac{t_{edge1}}{\tau_{ch}}}, \quad s_{00}(t_{edge2}) = e^{-\frac{-t_{edge2}+T_s}{\tau_{ch}}} - e^{-\frac{t_{edge2}}{\tau_{ch}}} \quad (6.13)$$

Substituting this in (6.12) and solving for t_{edge} yields for the first edge:

$$1 - e^{-\frac{t_{edge1}}{\tau_{ch}}} = 1 - \frac{1}{M} \Leftrightarrow e^{-\frac{t_{edge1}}{\tau_{ch}}} = \frac{1}{M} \rightarrow$$

$$\frac{t_{edge1}}{\tau_{ch}} = \ln(M) \quad (6.14)$$

And for the second edge:

$$e^{-\frac{t_{edge2}}{\tau_{ch}}} \left(e^{\frac{T_s}{\tau_{ch}}} - 1 \right) = \frac{M-1}{M} \rightarrow -\frac{t_{edge2}}{\tau_{ch}} + \ln \left(e^{\frac{T_s}{\tau_{ch}}} - 1 \right) = \ln(M-1) - \ln(M)$$

$$\frac{t_{edge2}}{\tau_{ch}} = \ln(M) - \ln(M-1) + \ln \left(e^{\frac{T_s}{\tau_{ch}}} - 1 \right) \quad (6.15)$$

The eye-width is $t_{edge2} - t_{edge1}$ which is now easy to compute (relative to the time constant):

$$rel. \text{ eye-width} = \frac{t_{edge2} - t_{edge1}}{T_s} = \frac{t_{edge2} - t_{edge1}}{\tau_{ch}} \frac{\tau_{ch}}{T_s} = \frac{\tau_{ch}}{T_s} \ln \left(e^{\frac{T_s}{\tau_{ch}}} - 1 \right) - \frac{\tau_{ch}}{T_s} \ln(M-1) \quad (6.16)$$

When the symbol time T_s is large compared to the channel time constant τ_{ch} , then the first term is roughly equal to one and the second term is zero (so fully open eye), but the terms rapidly decrease when T_s approaches or becomes smaller than τ_{ch} . The second term accounts for the fact that the sensitivity to ISI increases with a higher number of levels M and it is zero for binary signaling ($M=2$).

6.3.2 Eye properties for binary signaling with first-order channel models

In the next sub-section, the result of these equations will be discussed in more detail for different M . Here, we first focus on the result for plain binary signaling ($M=2$). In this case, the eye-height and eye-width become:

$$\text{binary rel. eye-height} = 1 - 2e^{-\frac{T_s}{\tau_{ch}}} \quad (6.17)$$

$$\text{binary rel. eye-width} = \frac{\tau_{ch}}{T_s} \ln \left(e^{\frac{T_s}{\tau_{ch}}} - 1 \right) \quad (6.18)$$

The smallest symbol time that still has an open-eye (both vertically and horizontally) is thus:

$$\frac{T_s}{\tau_{ch}} = \ln(2) \quad , \quad f_s = \frac{1}{T_s} = \frac{1}{\ln(2)\tau_{ch}} \approx \frac{1.44}{\tau_{ch}} \quad (6.19)$$

In section 6.2, the more accurate line-model predicted an achievable data rate of $f_s = 1.27/\tau_{ch}$ (with $\tau_{ch} = 0.41RC$ being the dominant time constant of the wire). So the higher-order terms in the transfer degrade the achievable data rate by only 12%. This shows that the simple first-order model is already quite suitable for a rough estimate of the achievable data rate.

6.4 Multi-level signaling

In a number of off-chip communication papers [97, 98, 105, 106] it is argued that, under certain conditions, multi-level signaling can improve the achievable data rate. In this section it will be examined if that is also true for on-chip communication.

Using multiple levels in the constellation enables to increase the symbol-time (T_s), while keeping the bit-rate (R) the same, because [92]:

$$R = \frac{\log_2(M)}{T_s} \quad (6.20)$$




Figure 6.6: Eye diagram for (a) 3-PAM and (b) 4-PAM with an on-chip interconnect. In both cases, $T_s = \log_2(M) \frac{1}{2} R_{wire} C_{wire}$ to get the same equivalent bit-rate as in Figure 5.3b.

Ternary signaling ($M=3$) or quaternary signaling ($M=4$) can for example reduce the symbol-time by a factor 1.58 and a factor 2 respectively, while keeping the same data rate. This does not automatically imply that the eye properties also improve, as is visible in the eye-diagrams in Figure 6.6. The magnitude of the ISI is lower at larger symbol times (Figure 6.6b has a larger symbol-time than Figure 6.6a), but the constellation simultaneously becomes more sensitive to ISI as the constellation points are more closely spaced. As a consequence, the eye-height in Figure 6.6b is substantially smaller than that of Figure 6.6a.

6.4.1 Eye properties for M-ary signaling with first-order channel models

The eye properties can also be analyzed more quantitatively, using the first-order model from section 6.3.1. When we use equation (6.10) for the relative eye-height and solve for zero eye-height we get:

$$\frac{1}{M-1} = e^{-\frac{T_s}{\tau_{ch}}} \left(\frac{M}{M-1} \right) \rightarrow e^{-\frac{T_s}{\tau_{ch}}} = \frac{1}{M} \rightarrow \frac{T_s}{\tau_{ch}} = \ln(M) \quad (6.21)$$

If we now substitute (6.20) to translate this to achievable data rate, the following relation is obtained:

$$\frac{T_s}{\tau_{ch}} = \ln(M) = R \ln(2) T_s \rightarrow R = \frac{\ln(2)}{\tau_{ch}} \quad (6.22)$$

So, the threshold of zero eye-height – which defines the achievable bit-rate - is independent of the number of levels in the constellation in the first order model.

The same boundaries hold for the eye-width from (6.16), but computing it is a bit more involved. Of course, the achievable data rate that is found by solving for zero eye-width should be the same as when using the eye-height, because for normal continuous signals, an

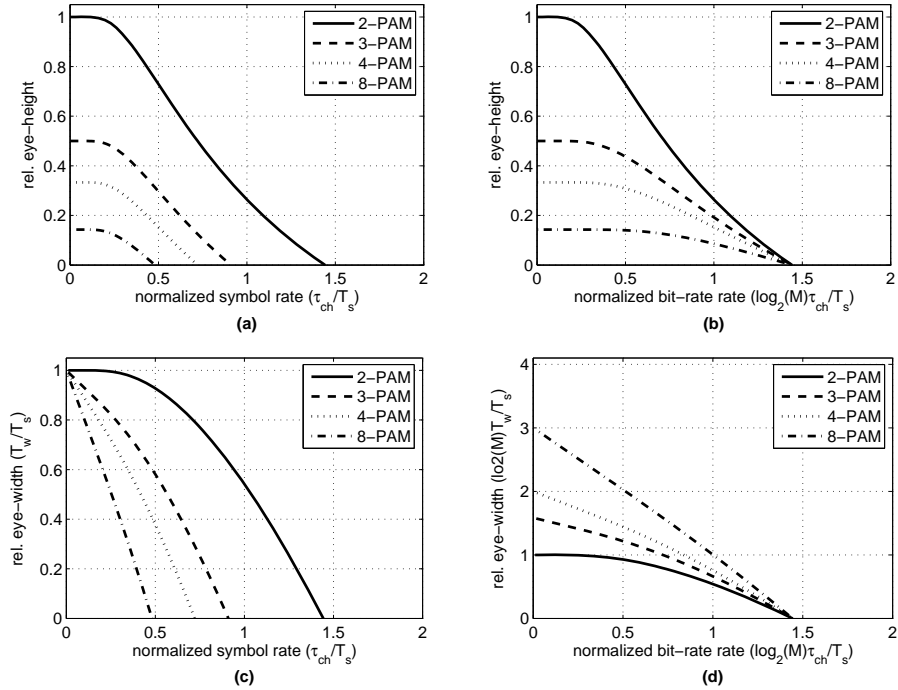


Figure 6.7: Eye-height (a,b) and eye-width (c,d) for M -ary PAM signaling over a first-order channel, as a function of the normalized symbol rate (a,c) or bit rate (b,d).

eye can not be closed in only one direction – at least not with the definitions from (5.18) and (5.19).

In Figure 6.7, these results are shown graphically, by plotting the eye-height from (6.10) and eye-width from (6.16), both as a function of the normalized symbol-rate τ_{ch}/T_s and as a function of the normalized bit-rate $R \tau_{ch}$ for a number of different M .

From these figures, it becomes clear that, at a given bit-rate R , multi-level signaling enables a trade-off in eye-height and eye-width. A higher M gives more eye-width, which is advantageous when time-dependent error-sources such as jitter are the dominant problem. A lower M gives more eye-height, which is better when additive disturbances are the dominant problem. For on-chip communication, the latter type of disturbance is usually more dominant so binary signaling is the best option.

6.4.2 M -ary eye properties with higher-order channel models

Note that the shape of the binary eye property curves in Figure 6.7 (a) and (c) quite accurately resemble the curves from Figure 6.1a and Figure 6.4a, which reconfirms that the first-order model is indeed suitable to estimate eye properties of actual on-chip channels.

Also note that the shape of the eye-height in Figure 6.7a is independent of M , except for a vertical scaling and a translation. This is not only true for the first-order model – with the

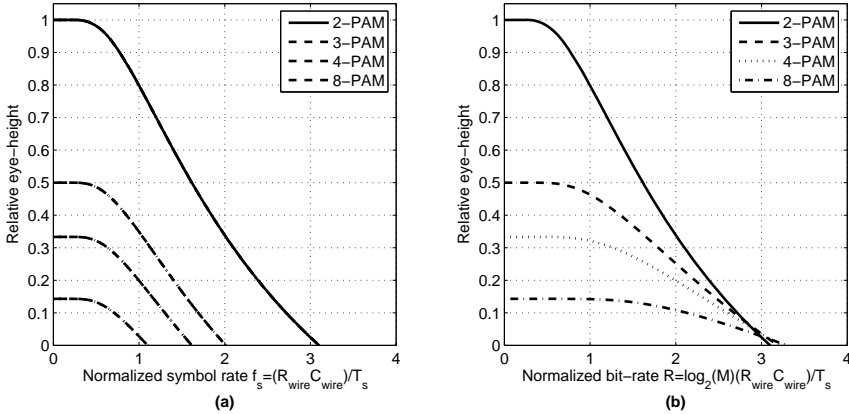


Figure 6.8: Eye-height for M -ary PAM signaling over an on-chip interconnect with conventional termination as a function of the normalized symbol rate (a) and bit-rate (b).

eye-height from equation (6.10) – but also for higher-order channels, provided that they have only real poles in their transfer. For these channels, the sum of the wanted symbol and the ISI is constant and equal to the DC-transfer of the channel, as was discussed in section 6.3. Assuming again that the DC-transfer is one, similar to the first-order model equations from section 6.3.1, the relative eye-height equation can be written as:

$$rel. \text{ eye-height} = s_{\max} \frac{d_{\min}}{d_{\max}} - \hat{s}_{e\min} = s_{\max} \left(\frac{M}{M-1} \right) - 1 \quad (6.23)$$

Thus the shape of the eye-height curve (as a function of e.g. T_s) tracks the shape of s_{\max} , apart from a scaling by $M/(M-1)$ and a translation over minus one. So, if one knows the shape of the eye-height for one M , one can deduce the eye-height for other values of M . This has been done to create the eye-height curves in Figure 6.8, which are based on the eye-height curve from Figure 6.1a.

The relation for the eye-width curve is not so simple between different M , as is visible in Figure 6.7c (and d). It is also expressed in the first-order eye-width equation (6.16), where the first term – which is non-linearly dependent on τ_{chl}/T_s – is dominant at low M , but the second term – which is linearly dependent on τ_{chl}/T_s – is more significant at high M .

The latency between start of transmission and the ideal sample instant is exactly T_s for the first-order model. For actual on-chip interconnects, the latency is a bit higher as the higher-order terms in the transfer do add some delay. The latency is however independent of M as the optimal sample instant is always found at the time where the symbol response is maximum (s_{\max}). But at a given bit-rate R , the absolute latency will be higher for higher M as the symbol time is higher.

6.4.3 Arguments for and against M -ary signaling ($M > 2$)

When higher-order interconnect models are used, then the achievable data rate is still quite independent of M , as is visible in Figure 6.8b. There is however a small advantage for higher M . But this benefit is only obtained at very small eye openings (3-PAM only

exceeds 2-PAM for rel. eye-height < 0.07) and only in a small symbol rate region, which does not make it a very practical advantage.

Only when other time constants – from e.g. a finite receiver bandwidth – start to approach the dominant interconnect time constant, such that the order of the dominant part of the transfer increases, then higher data rates can be obtained with multi-level signaling.

The benefit that multi-level signaling can have with higher-order channel transfers are exploited in some wireline communication systems. In [98, 107], it is explained that it is beneficial to go to multi-level when the slope of the channel transfer exceeds a certain threshold. For example, when the slope exceeds 10dB/octave then 4-PAM would be beneficial over 2-PAM according to [98]. This is because 10dB more SNR at Nyquist more than compensates the factor three reduction in initial eye-height (as the eye-height is proportional to $1/(M-1)$, and $20\log_{10}(1/3) = -9.5\text{dB}$), as was described in a general sense in [107].

Note that the eye openings and possible benefits of multi-level signaling can also be modeled analytically for second-order systems by solving t_{ideal} and $s_{max} = s\theta\theta(t_{ideal})$ from $ds_{\theta\theta}/dt = 0$, which implies solving $h_{\theta\theta}(t) = h_{\theta\theta}(t - T_s)$ because of equation (6.5). But these equations are rather lengthy and little additional information is gained for practical situations, so they are omitted here.

One other possible motivation to go to multi-level signaling is that, at a given maximum swing (d_{max}), the power is lower for signals with multiple levels. The power consumed in a wire is proportional to the variance of the transmitted code (assuming fully random codes) [92]. The variance of a normalized binary code is 1, while the variance of a ternary code is $1/3 \cdot ((-1)^2 + 0 + 1^2) = 2/3$. A quaternary code has a variance of $1/4 \cdot ((-1)^2 + (-1/3)^2 + (1/3)^2 + 1^2) = 10/18$. As M increases, the variance eventually becomes equal to $1/3$, which is the variance of a uniform random signal ($1/12 \cdot d_{max}^2$).

When an ideal transmitter would be available, then one could scale up the transmitted levels of M -PAM until the same power is obtained as for 2-PAM. For ternary signaling for example, one could scale d_{max} by $\sqrt{3/2}$ to get an equivalent power as 2-PAM. However, even with such a scaled d_{max} , 3-PAM only gets a higher eye-height than 2-PAM at data rates above $2.4 \text{ bit}/R_{wire}C_{wire}$, where the eye-height has dropped to 20% of its DC value according to the curves from Figure 6.8b. For higher M the break-even rate is even higher. Given that practical transmitters do have constraints on the transmitted levels (e.g. limited to V_{dd}) and that a multi-level transmitter tends to be less efficient than a binary transmitter, it is thus unlikely that these theoretical power saving possibilities can become practical. In the end, binary transmission is likely to be the most power efficient, especially when the data rate is lower than $2.4 \text{ bit}/R_{wire}C_{wire}$.

6.5 Achievable rates for band-pass signals

Another signaling technique that is widely used in off-chip data communication is band-pass signaling. Many different modulation schemes exist to create band-pass digital signals, including band-pass coding such as Manchester coding, or non-linear modulation methods such as FM [92, 94, 95]. Here we stick here to the linear methods, which are the most widely applied, and analyze if they can also have their merits for on-chip communication.

6.5.1 Single carrier PAM modulation

Band-pass signaling does not immediately come to mind as a candidate for on-chip communication as the circuitry is typically more involved than baseband communication circuits. However, in 2002 Chang et. al. applied such circuits to on-chip communication [21, 108]. The motivation that was given was that the group velocity is higher at higher frequencies, where the wires behave more like a transmission line. Quite a large cross-sectional area was used to get transmission-line behavior in a usable frequency range, so it is not a really practical approach from a BW/area perspective (as was discussed in section 2.7). The group-delay indeed decreases at higher frequencies, even when skin-effect is modeled and at really high frequencies, the delay is ultimately limited by the speed of light in the medium, but data transmission in this region does however come at a price of significant attenuation (as was discussed in section 3.7).

The system as proposed in [21, 108] is basically the same system as was shown in Figure 5.7b, but then with only one of the two ‘sub-channels’ active, as a simple binary alphabet is used instead of quadrature modulation. As was explained in section 5.4.4, this system can be analyzed with the model in Figure 5.7c, with linear (filter) models for the various wave shapes to model the modulation and demodulation.

Examples of the transmitter’s wave shape h_{tx} , the result when it is passed through the channel ($h_{tx} * h_{ch} = h_{tx}$ convolved with h_{ch}), and the result after demodulation and filtering at the receiver (the symbol-response: $s_{00} = h_{tx} * h_{ch} * h_{rx}$) are all shown in Figure 6.9 for two different frequencies and phases of the local oscillator (LO). What is also shown in Figure 6.9 is the frequency-domain transfer functions of the channel H_{ch} . The effect it has on the spectrum of the transmitted band-pass PAM (that should have a sinc-shaped spectrum when it would not be distorted [92]) is visible in the plot of $|H_{ch}H_{rx}|$.

What is notable in the figures is that the start and end of the pulse differ significantly between cosine or sine wave pulse shapes. The sine wave shape has more wanted pulse, but creates also much more ISI, as is best visible in sub-figure (a1) and (b1). In that sense, this band-pass signaling approach can be viewed as a pulse-shaping method, that reduces ISI, but also reduces the wanted signal level ($s_{00}(t_{ideal})$).

The achievable data rates for these band-pass systems are shown in Figure 6.10 on page 140. The figure shows that the cosine-wave symbols (used for (a) and (c)) have a much higher achievable data rate than the sine-wave shapes (used for (b) and (d)), due to a lower ISI. Note that the jumps in the latency in Figure 6.10 are caused by the fact that the optimum sample point can change from around one local maximum in s_{00} to the next, depending on how the ISI sums up. That there are multiple local maxima also causes the irregular shape in the eye-height versus data rate. What is also shown in Figure 6.10 is the maximum of $h_{tx} * h_{ch}$ which is a measure for the amplitude of the signal at the receiver. This amplitude clearly decays very fast for higher data rates. This amplitude can be boosted in the demodulation process (modeled by h_{rx}) which would give higher amplitudes at the detector (to get higher amplitudes for s_{00}), but that will also boost error signals such as crosstalk.

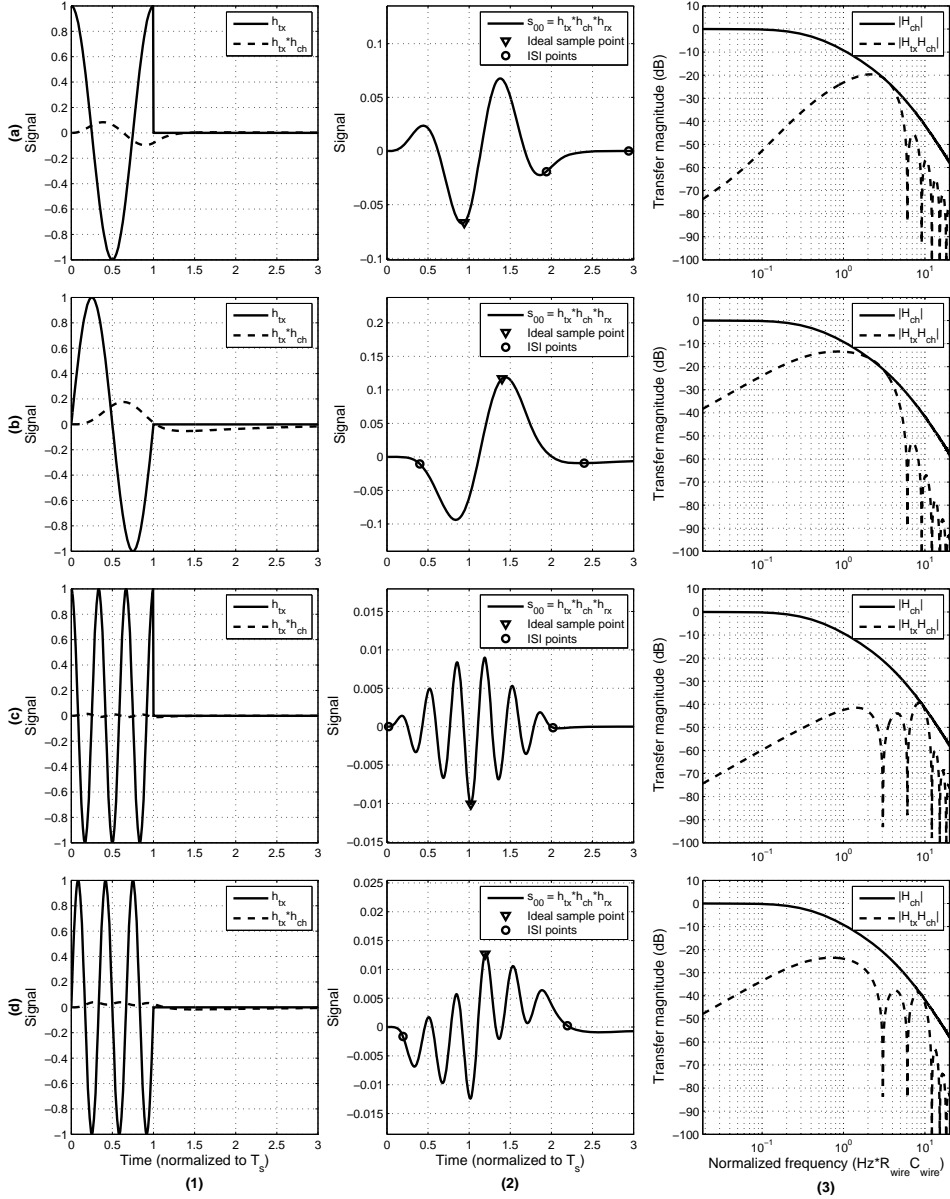


Figure 6.9: Symbol wave shapes (1,2) and the frequency domain transfers (3) for band-pass PAM signaling with $f_s=3$ symbols/ $R_{\text{wire}}C_{\text{wire}}$. In column (a) and (b) $f_{\text{LO}}=f_s$. In (c) and (d) $f_{\text{LO}}=3f_s$. The phase of the LO_{TX} is 0° in (a) and (c) and -90° in (b) and (d).

So it is in theory possible to achieve high data rates with band-pass signaling, but the system is complex, sensitive to interference, and requires careful tuning of transceiver parameters such as the phase of the local oscillators. Especially the phase of the transmitter LO has a large impact, as it changes the transmitted symbol wave shape. At low f_{LO}/f_s ratios,

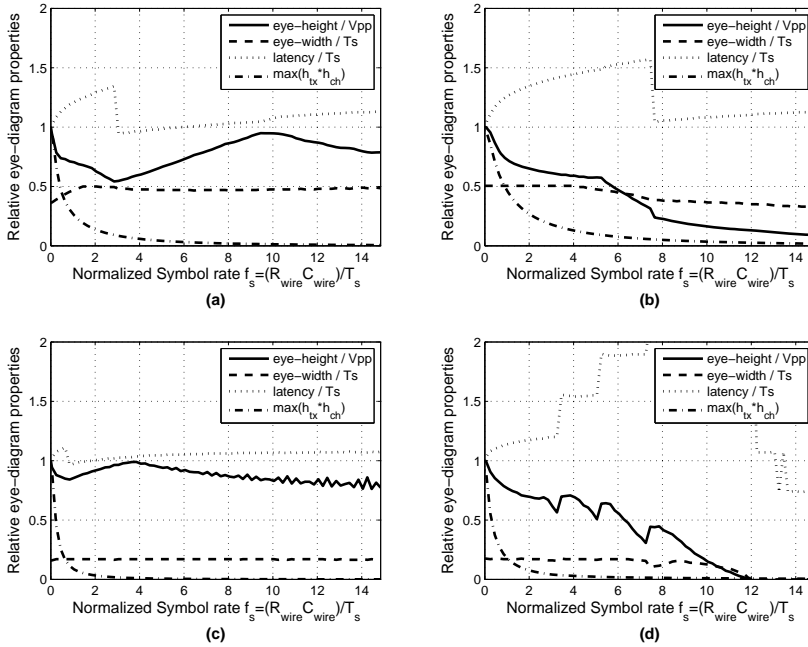


Figure 6.10: Eye properties for band-pass 2-PAM signaling. (a) and (b) use $f_{LO} = f_s$, while (c) and (d) use $f_{LO} = 3f_s$. The phase of the LO_{TX} is 0° in (a) and (c) and -90° in (b) and (d).

the phase of the receiver LO is of less importance and one could even omit the demodulation and directly detect the symbol from $h_{tx} * h_{ch}$. However, for higher f_{LO}/f_s ratios, it is important to set the receiver demodulation phase correctly, as is also visible in the small eye-width in Figure 6.10c and d.

As mentioned before in section 5.4.3, band-pass signaling can also be viewed as a pulse shaping method. However, equalization can be a much more effective method to shape the symbol response and reduce the ISI, as will be discussed in Chapter 7. Simple equalization methods such as first-order FIR or PW pre-emphasis are less complex to implement and more successful in achieving a good wanted signal level with low ISI. These pre-emphasis methods have a pulse shape that is not very dissimilar from the sine wave in Figure 6.9b1, but the pre-emphasis parameters can be tuned to remove the ISI from the tail of the response.

For a quantitative comparison between band-pass PAM and equalization, note that Figure 6.10a shows that band-pass PAM with sine wave pulse shape has 95% relative eye-opening at a data rate of 10 bit/ $R_{\text{wire}} C_{\text{wire}}$, but the received amplitude $\max(h_{tx} * h_{ch})$ is only 0.013V and after demodulation with the matched filter h_{rx} , the s_{00} is only $7.3e-3$ (for a h_{rx} with normalized amplitude). With FIR or PW pre-emphasis, the absolute eye opening at the same data rate is much larger: 0.041V or 0.047V respectively, as will be shown in Figure 7.6.

So, in conclusion, band-pass PAM can increase the achievable data rate, but equalization is a simpler and more effective solution.

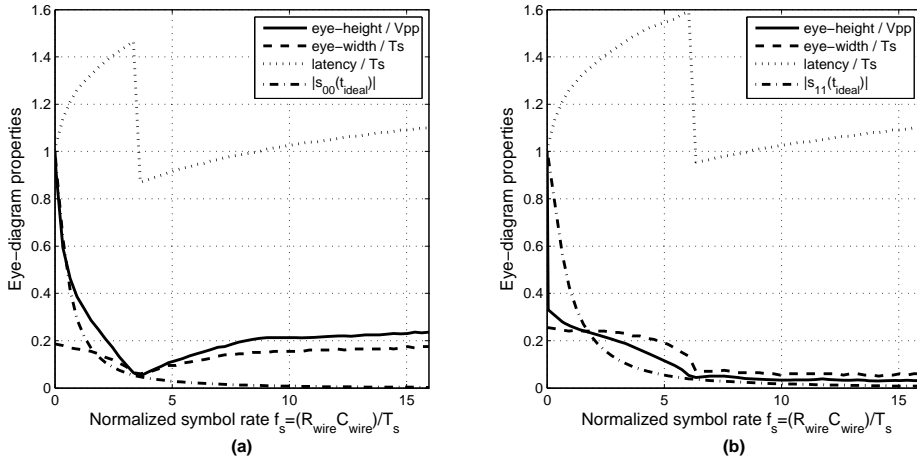


Figure 6.11: Eye properties for 4-PSK band-pass signaling with $f_{LO}=f_s$. The properties of the I channel is shown in (a) and the Q-channel in (b).

6.5.2 Single carrier quadrature modulation

Quadrature modulation can in theory double the achievable data rate in the same signal bandwidth, by using two quadrature phases of the carrier signal. However, as was explained in section 5.4.4, it does require that the two modulated components are orthogonal; or in other words: the signal should be a true band-pass signal, bandlimited between $0 < f_0 < f_1$.

This is clearly not the case for the simple modulation schemes for on-chip communication, as was also shown earlier with the constellation example from Figure 5.8. It is also visible in the symbol response examples from Figure 6.9, which show quite different behavior for the two 90degrees phase-shifted variants of the symbol response.

To see how much the crosstalk between the I and Q channel affect the eye properties, an analysis with 2-PAM for both the I and Q channel (equivalent to 4-PSK) was carried out and the results are shown in Figure 6.11. Compared to band-pass signaling with a single channel as was shown in Figure 6.10a and b, the eye is clearly much more closed. At a given symbol-rate, the data rate for the quadrature modulation is two times higher, but this advantage is more than negated by the smaller eye-opening. For example at a symbol-rate of 5 symbols/ $R_{\text{wire}}C_{\text{wire}}$, the relative eye-height is only 11% for both channels, and the s_{00} at this rate is only 0.029 and 0.055 for the I and Q channel respectively. So the eye opening is smaller than the eye-opening at 10 symbols/ $R_{\text{wire}}C_{\text{wire}}$ with non-quadrature band-pass PAM. So given that the PAM band-pass was already not a suitable option for on-chip communication, quadrature band-pass signaling is even less suitable.

As was mentioned in section 5.4.4, the crosstalk between the I and the Q channel could be reduced by transforming the constellation prior to detection, but that would require a 2x2 matrix of coefficients that have to be adapted for different symbol-rates, which is quite complex and was therefore not investigated further.

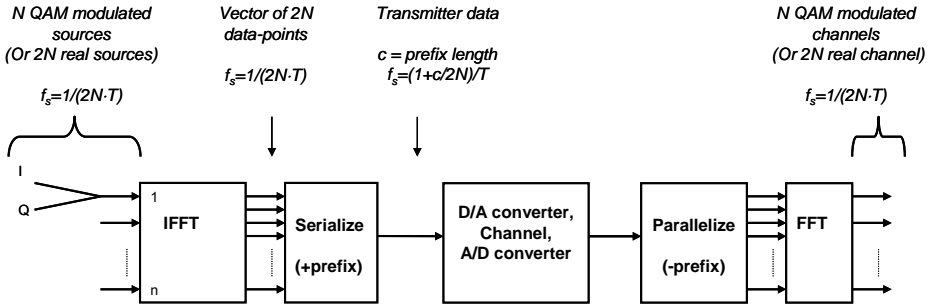


Figure 6.12: simplified OFDM system schematic.

6.5.3 Multi-Carrier and OFDM or CDMA

In the first phase of this project, it was also investigated whether simple variants of multi-carrier/multi-channel schemes could be used to improve the data capacity of the interconnects [109]. As was discussed in section 5.4, multi-channel transmission schemes such as OFDM or CDMA [95] can also be analyzed by transforming the signaling scheme to a functionally equal system that consists entirely of filters and samplers. Conceptually, such schemes use filter-banks for the transmission and modulation of different channels and matched filter-banks for the reception. An OFDM system as is shown in Figure 6.12 can be analyzed with the schematic from Figure 5.7c, by adding more input and output filters to the system. With this analogy, it is again possible to derive the eye-diagram properties at the receiving end.

As with IQ modulation, it is again vital that the inter-channel interference (ICI) is also taken into account, because the bandwidth limitations in the channel destroy the orthogonality between the channels.

For these multi-carrier systems, the worst-case eye-opening analysis that is useful in single-channel systems gives a pessimistic estimate of the achievable data rate. With multi-channel systems, the chance that this worst-case situation occurs can become arbitrarily small if the numbers of channels increase. Therefore, the statistical approach from section 5.4.5 was used to estimate the BER.

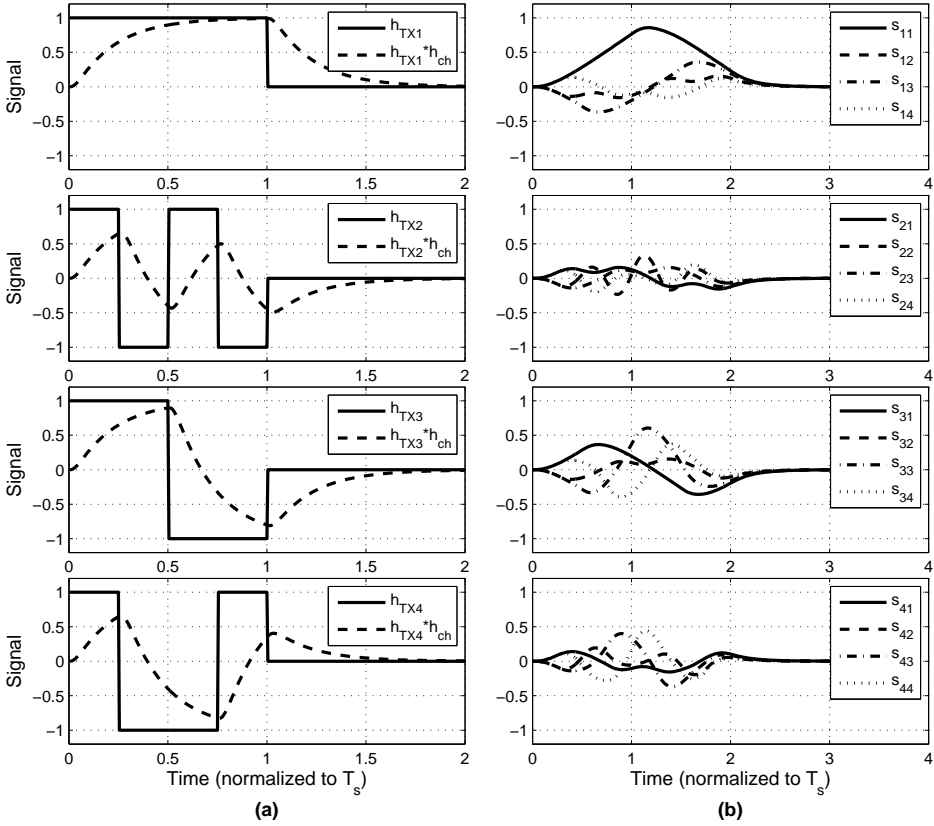


Figure 6.13: Symbol wave shapes for four channel CDMA, with $f_s=0.5$ symbols/ $R_{wire}C_{wire}$ (so chip-rate= $2/R_{wire}C_{wire}$). In column (a), the transmitted wave shapes are shown before and after the channel. Column (b) shows the symbol responses after the receiver filters.

An example of such analysis is 4-channel CDMA based on a 4×4 Hadamard matrix (also see section 4.4.7), of which the transmitted and received symbol responses are shown in Figure 6.13. Figure 6.14 on the next page shows the estimated BER for this system as a function of the detection instant (t_d) at a normalized data rate of 2 bit/ $R_{wire}C_{wire}$. The validity of the estimated BER was also verified with time-domain simulations, as mentioned earlier in section 5.4.5.

The BER numbers in the figure are far too large for reliable communication, while at the same data rate, plain binary signaling still has an open eye (Figure 5.3). Analysis with CDMA at other data rates or with another number of channels showed similar results: it is not a viable alternative for plain binary signaling.

The problem with multi-carrier schemes is mainly the ICI: the interconnect destroys the orthogonality between the different channels, as was also the case for the I/Q modulation. The achievable data rate for both CDMA and OFDM (even with long cyclic prefixes) is quite low due to this inter-channel interference and they are therefore not useful for reliable signaling across on-chip interconnects, at least not without additional channel equalization.




Figure 6.14: BER as a function of detection instant for four channel CDMA, with $f_s=0.5$ symbols/ $R_{\text{wire}}C_{\text{wire}}$ and binary signaling per channel.

A combination of multi-channel transmitter (similar to OFDM) and channel equalization could solve this problem. Such a combination is for example presented in [110] for off-chip communication. But this system is much too complex for implementation in on-chip buses. For the well-defined on-chip channel transfer, simple single-channel transmission with equalization, as discussed in the next section, is likely to perform at least as well.

6.6 Summary and conclusions

The list below shortly summarizes the results and conclusions from this chapter:

- Plain binary signaling over a single shielded wire is limited to a data rate of maximally $3.1 \text{ bit}/R_{\text{wire}}C_{\text{wire}}$ with conventional termination ($R_s=0, R_l=\infty$). This increases to $8.8 \text{ bit}/R_{\text{wire}}C_{\text{wire}}$ with resistive receiver termination ($R_s=0, R_l=R_{\text{wire}}/10$) or capacitive transmitter termination.
- For conventional or resistively terminated interconnects, twisted differential wires boost the data rate by 41% and 55% respectively, which does not necessarily justify the additional cost in area. But a transceiver that uses twisted differential wires and conventional termination can still be more power efficient above data rates of about one $\text{bit}/R_{\text{wire}}C_{\text{wire}}$.
- For capacitive transmitter termination (which is a very power efficient termination technique), the benefits of twisted differential wires clearly outweighs the area costs, as the achievable data rate is 112% higher than for an unshielded wires in a bus.
- An on-chip interconnect can also be approximated by a first-order channel model (as long as the crosstalk is mitigated). For first-order channels, simple closed-form analytical solutions were developed that give the eye height and eye-width as a

function of data rate (and have only 12% difference with the numerical results for the higher-order line model).

- For first-order channels and M-ary PAM signaling, the achievable data rate is independent of the number of levels. Multi-level signaling with $M > 2$ can be used to trade off eye-height for eye-width. For on-chip interconnects, eye height is more important, and in this respect plain binary signaling shows the best results. For the higher-order interconnect models, multi-level signaling can have a higher eye opening, but only at very high rates where the eye is already nearly closed, making it not a practical benefit.
- Band-pass signaling with real alphabets can increase the achievable data rates compared to plain binary signaling, but it is complex and less effective than other methods such as equalization. Band-pass signaling with complex alphabets suffers from crosstalk between the I and Q channel, resulting in lower eye-openings than with real alphabets.
- The investigated multi-channel signaling schemes show no benefit in achievable data rate. Instead, they achieve worse results than plain binary signaling while also being much more complex.

Chapter 7

Equalization techniques

7.1 Introduction

Equalization is a well-known topic to improve communication over a bandlimited channel [95]. It can be part of either the communication blocks at the transmitter side, or at the receiver side or both, as was shown earlier in Figure 5.2. It can also be implemented in many different ways, with either time-discrete or time-continuous filters.

This chapter will discuss equalization in more detail and how it can be applied to on-chip transceivers. To present some background, the chapter starts in the next section with a literature overview of the types of equalization that have been applied in recent off-chip and on-chip transceivers.

After the general overview, the different forms of equalization that were applied in this project are described in more detail. This includes two types of transmitter-side equalization, FIR and PW pre-emphasis, which will be described in section 7.3 and 7.4 respectively and compared to each other in section 7.5.

As a preview, Figure 7.1 on the next page shows how much improvement FIR and PW pre-emphasis can give in eye opening and hence in achievable data rate. Without equalization, signaling at a symbol rate of $5 \text{ bit}/R_{\text{wire}}C_{\text{wire}}$ is clearly not possible as the eye is completely closed, but with both FIR and PW pre-emphasis, the eye is still perfectly open.

After the discussion of the transmitter-side equalization, it will be described in section 7.6 how receiver-side equalization was applied in this project, using a special form of decision feedback equalization (DFE).

The last parts of the chapter discuss how tolerances in wire and equalization parameters can be dealt with in section 7.7 and whether multi-level signaling has benefits in combination with equalization in section 7.8. The chapter finishes with a short summary and conclusions in section 7.9. The quantitative results from this chapter are also summarized in Appendix B.

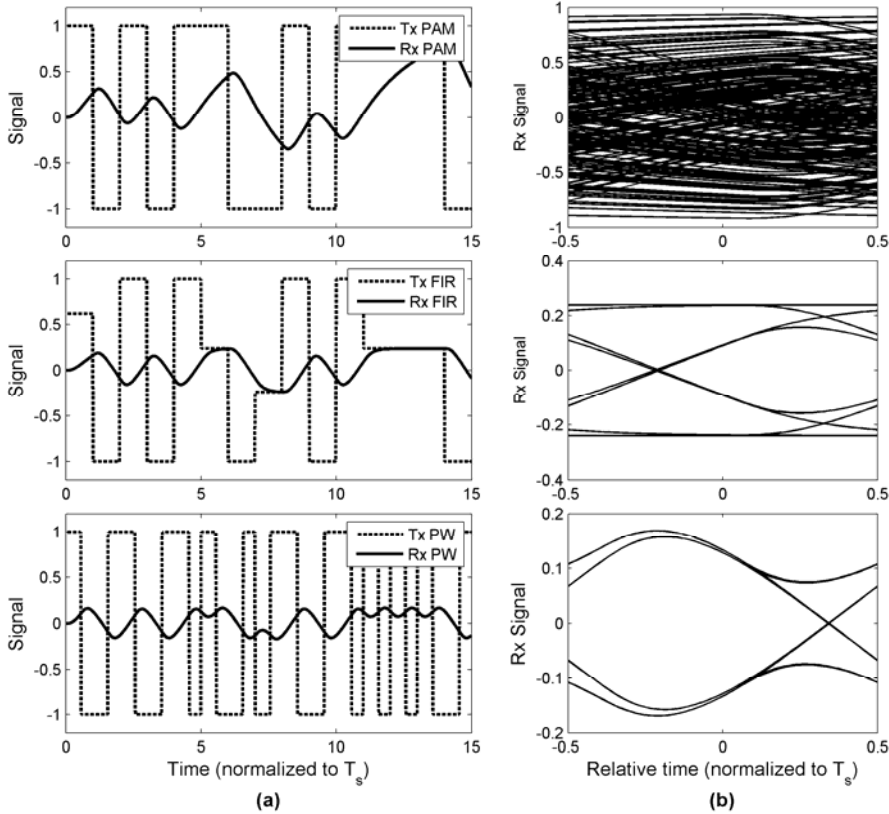


Figure 7.1: symbol streams (a) and eye diagrams (b) of plain binary signaling with a symbol-time of $T_s=1/5 \cdot R_{\text{wire}} C_{\text{wires}}$ without any equalization (top), with FIR pre-emphasis (middle) and with pulse-width pre-emphasis (bottom).

7.2 Equalization overview

This section gives a short overview of the various implementations found in recent literature, focusing on multi-gigabit per second implementations for wireline, backplane and on-chip communication. The former two fields have a long history in equalizing transceivers for bandlimited channels and the data rates are similar as for on-chip transceivers. They can therefore be a good source for inspiration for on-chip communication, with the difference that circuits for wireline and backplane communication are allowed to consume much more area and power and can consequently also be more complex. The types of equalization that were applied in this project are included in the overview, with later sections discussing on the technical details.

To complete the overview, the second part of this section also briefly discuss more advanced topics, such as popular adaptation algorithms and the combination of equalization with clock-recovery. Again, the goal of the overview is to draw some inspiration for

possible adaptation algorithms for on-chip communication, as discussed in more detail in section 7.7.2

7.2.1 Transmitter-side equalization

On the transmitter side, equalization is sometimes known as forward- or pre-equalization, but is more often called pre-emphasis (or de-emphasis), to indicate that the transmitter compensates the attenuation of the channel by emphasizing (or boosting) these frequencies. In practice, it is usually much easier to let the transmitter attenuate the other frequencies, which are not attenuated by the channel, for example because the swing at the transmitter is limited by the supply and emphasis is not an option. With respect to the flattening of the overall frequency transfer and the corresponding reduction in ISI, the two methods are the same, with the second method having the side-effect of attenuation. De-emphasis would be the most appropriate name for this second method, but in practice pre-emphasis is more commonly used, even when there is an overall attenuation [97].

In wireline and backplane communication, discrete-time pre-emphasis is very popular. In such pre-emphasis schemes, delayed copies of the symbols are weighted and summed to create equalizing finite impulse response (FIR) filters [97, 98, 111-117].

The tap-delays for the pre-emphasis do not have to be equal to the symbol time, but can also be a fraction of this time (a form of fractionally spaced equalization [95]). Half-symbol spaced delays can for example be used, which gives an additional degree of freedom to also cancel ISI at the zero-crossings [118, 119], thereby minimizing the so-called data-dependent jitter [118] and opening the eye in the horizontal direction.

Half-symbol delays are easy to create with digital logic, but more complicated techniques have also been reported to create other fractional delays, for example the use of LC delay lines [120]. When only two taps are used for the pre-emphasis, then clock skew can also be used to create a fractional delay, as was proposed for the on-chip communication scheme in [74, 121].

Transmitters with continuous-time equalizing filters are hard to find in backplane and wireline communication literature, with the possible exception for transmitters where the equalization is directly coupled to the line termination. These include for example AC-coupled pulsed chip-to-chip transceivers [122] and inductive termination to boost transmitted HF components [123]. These also include the capacitive on-chip transmitters that were used in this project [62, 124] and in [74, 78, 121], and with some modifications also in [77].

FIR pre-emphasis is also used in some on-chip transmitters [64, 71, 72, 78, 121]. In this project, we initially examined FIR pre-emphasis, as discussed in section 7.3, but moved towards another form of pre-emphasis: pulse-width pre-emphasis [33, 73] (discussed in section 7.4), because of its implementation advantages. Pulse-width pre-emphasis also turned out to be very suitable for wireline communication [39, 44].

7.2.2 Receiver-side equalization

On the receiver side, one can use a very similar equalization process: those components that are attenuated by the channel should be boosted or alternatively, the other frequency components should be attenuated by some type of equalizing filter. Equalization at the

receiver side is the most traditional one, or is at least treated as such by the textbooks. [92, 94, 95].

Two flavors of equalization are popular at the receiver side in wireline and backplane transceivers. The first flavor is continuous-time equalization, with (HF) boosting circuits that compensate for the attenuation of the channel, also called ‘linear equalization’ [117, 123, 125-133]. This flavor has the longest history and is still being used in many contemporary transceivers, but its popularity seems to diminish because of another equalization form, the so called ‘Decision-feedback equalization’ (DFE) [98, 111-116, 119, 134-138]. In DFE, a weighted sum of previously detected symbols is fed back to the detector input, which creates a discrete-time filter, but with the inclusion of the non-linear detector in the loop. It is therefore sometimes also called ‘nonlinear equalization’.

A motivation to use DFE instead of linear equalization is that DFE does not amplify crosstalk [112]. But DFE with more than a few taps tends to become complex and power hungry, so it is sometimes combined with a linear equalizer [111, 112, 116, 137, 138]. An often stated motivation is that linear equalizers can cancel long tails in the symbol response with lower complexity. But actually, DFE with a continuous-time feedback filter can also do this, as was shown in this project [62, 124]. Such a continuous-time feedback DFE was later also applied to a backplane communication receiver [139]. Its method of operation will be discussed in more detail in section 7.6.

Apart from complexity, DFE with discrete-time feedback filters also suffer from circuit delays within the loop, especially for the first filter tap. To circumvent this delay problem, loop-unrolling [134, 140] has become popular in multi-gigabit DFE receivers [111, 112, 114-116, 135, 141].

Apart from the special DFE that was used in this project, DFE has also been applied in some other on-chip communication receivers. This includes very simple receiver circuits where the threshold of an inverter is adjusted to create a hysteresis which acts as a single-tap DFE [90, 142] (in [142], the inverter threshold is also adapted to reduce crosstalk). More recently, DFE with loop-unrolling was also applied in on-chip transceiver experiments [71, 72], but was eventually not used as pre-emphasis alone turned to be more power efficient.

So far two popular categories of receiver equalization were discussed, but there exist a third category: discrete-time receiver equalizers without feedback [105, 143-145]. Such type of equalizers however have not become popular and do not seem to have clear advantages over the other options.

As with transmitter side discrete-time equalization, the tap-delays of discrete-time equalizers (whether it be DFE or FIR equalizers without feedback) do not have to be equal to the symbol time, but can also be a fraction of it, for example half-symbol spaced equalization [105, 119].

7.2.3 Transmitter and receiver equalization

The choice whether to equalize at the transmitter or at the receiver side depends on many factors, especially when circuit-level considerations are also taken into account. There does not seem to be a decisive preference for either of the two and it is in fact becoming increasingly popular to equalize at both ends of the channel [98, 111-116, 119, 126, 135].

When DFE is used at the receiver, an often stated motivation to also use transmitter pre-emphasis is the statement that transmitter pre-emphasis can correct pre-cursor ISI, while receiver DFE can not [111]. Also, pre-emphasis can cancel long tails in the symbol response with only a few taps (as discussed in section 7.3), while DFE can cancel only as many ISI points as it has taps [111]. DFE on the other hand is able to equalize the wanted signals without amplifying crosstalk [112].

A whole different reason to use both transmitter and receiver equalization is to enable an asymmetric system, with a controller with transmitter equalization for the outgoing data-paths and receiver equalization for the incoming data paths. This simplifies the transceivers at the other end of the channel, which can be of interest for for example memory interfaces [117].

7.2.4 Adaptive equalization

In many transceivers, the properties of the channel are not known a priori. This is especially true for wireline and backplane transceivers, as they have to be able to operate with various channel lengths. Also, the transfer functions of the channels can become quite complex and unpredictable due to e.g. variations in the dielectrics (especially true for backplanes with non-homogenous PCB materials such as FR4). Therefore, many backplane and wireline transceivers use some form of adaptation for the equalization. A short overview of popular adaptation methods is given in this sub-section.

Most transceivers that use discrete-time equalization at the receiver, combine equalization with variants of LMS algorithms to adapt the coefficients, most often so-called signed-signed LMS [106, 111, 113-115, 145]. Sometimes, the LMS algorithm is slightly modified to also minimize ISI at the left and right edge of the eye [119] and thereby reduce the data-dependent jitter, or it can contain additives such as an eye opening monitor [138]. Other adaptation algorithms have also been proposed, for example adaptations based on BER optimization [136, 146].

The adaptation algorithms are sometimes used primarily for the receiver (DFE) equalizer [113-115, 119], or in other implementations, there is return communication to also adapt the transmitter side equalizer [106, 111, 135, 136]. Often, a mix of startup calibration with a training sequence and online adaptation is used, especially when the transmitter is also adaptive.

For continuous-time receiver equalizers, other adaptation algorithms have been used, for example methods that compare the energy in different frequency bands before and after the detector [125, 128, 129, 131, 133]. A similar method is to compare the slope before and after the detector [130]. The adaptation can also be done digitally, by for example analyzing the histogram of the data-edges with an oversampled receiver [132].

The variants of adaptive equalization where the received power at different frequencies is the adaptation criterium, might also be of interest for (future) on-chip communication systems. Its implementation does not have to be complicated, especially not when a few interconnects in a bus can be dedicated to pilot tones, to simplify the energy measurement, as discussed in section 7.7.2.

7.2.5 Adaptive equalization and clock recovery

Note that, with adaptive equalization, it is still necessary to use a proper sample instant that minimizes ISI. A suboptimal sample instant will complicate the equalization and possibly even make it unstable and unusable. As was discussed in section 5.5, the clock-data recovery (CDR) circuit in the receiver has the task to find a suitable sampling instant (ensure phase-alignment).

Although full-fledged CDRs are not yet necessary for on-chip communication, it is still instructive to see how CDRs are applied in wireline and backplane transceivers to get an indication how to approach those situations where a global clock is no longer available or situations where the wires are so bandlimited (or the parameters so variable) that it is no longer possible to send a directly usable clock alongside the data. Thus, an overview of how CDRs are applied in wireline and backplane transceivers is given below.

In some wireline transceivers, the CDR has to recover both the frequency and the phase of the clock. But even when the clock is already known in the receiver, or in source-synchronous links where a (possibly low frequent) reference clock is sent along the data channels, CDR's can still be found to multiply the frequency of the clock and get phase alignment [127, 132, 145]. Once the frequency of the recovered clock is known, control of the clock generator phase is either done continuously or through the selection of one clock-phase from a finite set of possible phases (creating a sort of 'digital' CDR [132]).

In case of a continuous-time receiver equalizer, the CDR can simply be placed after the equalizer [131, 133]. In this case, CDR operates on data patterns with well-defined edges, simplifying the clock recovery and enabling the use of classical phase detectors for plain binary (NRZ) data such as the Hogge [103] or Alexander [104] phase detector. As the eye is already open at this point, the phase-detector can double as a data detector [131].

In case of receivers that use DFE, placement of the CDR after the equalizer is less straightforward. This is especially so when loop-unrolling (see section 7.6) is used, because then there is no single stable transition point for a phase-detector to use, as discussed in [116, 138]. So for many receivers with loop-unrolled DFE, the CDR is a standalone circuit that primarily uses the input signal to the DFE for phase detection [100, 111, 114], possibly also taking some information from the DFE output into account [135, 141].

However, it is still desirable to place the CDR after the DFE, as a properly equalized signal has much less data-dependent jitter on the edges. Some receivers with loop-unrolled DFE manage to do this by using more complex phase detectors than the classical Hogge or Alexander detectors, using for example multiple edge detectors with different thresholds (an 'unrolled' edge detector) [115] or an 'eye tracking' phase-detector that does not sample on the edge but just before and after it [116].

For the receivers that do not use loop-unrolled DFE, the CDR can be placed after the DFE with more ease [113, 119, 136, 138]. But still, care has to be taken, as the recovered clock itself is used in the DFE, which creates an additional feedback loop that can interfere with the control algorithms. This can for example be seen in [119] where, among other things, an eye-height comparison is given between CDR's with unequalized inputs and CDR's with DFE equalized inputs.

A method to avoid interference between the CDR and DFE control loops can be to either use very different convergence speeds for the two loops [112] or to combine the two into a single control algorithm [113, 115, 136]. A simple method to combine CDR with adaptive receiver equalization is for example presented in [113]. There, the eye is sampled both in the centre and at edges (as in a normal Alexander phase detector [104]) but the sample value at the edges is not only used for adaptation of the phase of the clock but also to update the DFE coefficients.

In general, combining adaptive equalization with adaptive clock-recovery is not always trivial. In [115, 136], it is for example shown that clock-recovery and transmitter FIR adaptation can interact, with suboptimal results. Note that interaction between adaptation loops is automatically prevented when the different adaptation steps are not done online, but are done sequentially in a calibration routine at startup [145].

So, the algorithms that combine adaptive clock recovery and adaptive equalization for off-chip communication are not really simple, especially not when all the adaptation loops have to use the same data stream. For on-chip applications, a big part of the complexity of clock and data-recovery can be avoided by using dedicated lines to transmit synchronization data, and using the assumption that the lines are matched to the other lines in the bus. In section 7.7.2, a method is discussed that should enable simple clock recovery and equalization with only a few wires overhead.

This concludes the literature overview of equalization and clock-data recovery. From here on, the chapter will focus on a more detailed explanation of the on-chip equalization techniques that were used in this project.

7.3 FIR-pre-emphasis

The dominance of the first-order roll-off in the transfer makes an on-chip interconnect a very suitable candidate for simple pre-emphasis transmission schemes, such as a two-taps FIR-filter. Such a two-taps filter can be expressed as a discrete time difference equation that acts on the transmitted symbol stream b_0 :

$$b'(n) = \beta \cdot b(n) - (1 - \beta) \cdot b(n-1) \quad (7.1)$$

The second filter coefficient is in this case taken to be $(\beta-1)$ to ensure that the maximum value of b' does not exceed the maximum value of b , so that the signal remains within the signal boundaries set by e.g. the supply.

In the text below, it is discussed how β can be optimized for maximum eye-width, starting with an analytical model for a first-order channel (as was also created for un-equalized signaling in section 6.3.1) and subsequently with a numerical analysis of more accurate on-chip wire models.

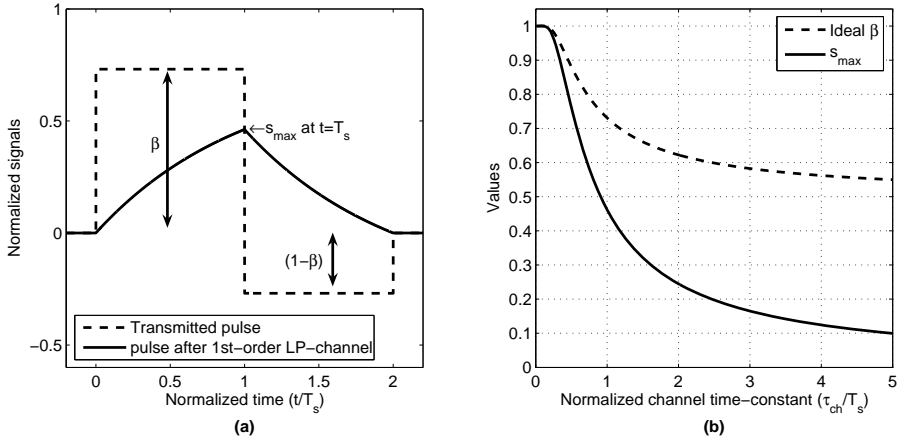


Figure 7.2: Symbol response (a) and ideal β and s_{max} (b) for FIR pre-emphasis with a first-order channel. The symbol response in (a) is an example with $T_s = \tau_{ch}$.

7.3.1 FIR pre-emphasis with first-order channel models

For first-order low-pass channels, FIR pre-emphasis can completely cancel the ISI as shown in Figure 7.2a. The value for β that gives zero ISI is a function of the symbol-time (T_s) and of the time constant of the channel (τ_{ch}) and can be found by writing the transmitted pulse shape as a summation of three step responses:

$$g_0(t) = \beta \text{step}(t) - \text{step}(t - T_s) + (1 - \beta) \text{step}(t - 2T_s) \quad (7.2)$$

The step response terms of a first-order channel are simple exponential functions (valid from the start of the step), so the symbol response can be written as:

$$s_{00}(t) = g_0 * h_{00} = \beta \left(1 - e^{-\frac{t}{\tau_{ch}}} \right) \Big|_{t \geq 0} - \left(1 - e^{-\frac{t - T_s}{\tau_{ch}}} \right) \Big|_{t \geq T_s} + (1 - \beta) \left(1 - e^{-\frac{t - 2T_s}{\tau_{ch}}} \right) \Big|_{t \geq 2T_s} \quad (7.3)$$

The response will be zero for $t \geq 2T_s$, meaning no ISI, on the following condition:

$$e^{-\frac{t}{\tau_{ch}}} \left(-\beta + e^{\frac{T_s}{\tau_{ch}}} - (1 - \beta) e^{\frac{2T_s}{\tau_{ch}}} \right) = 0 \rightarrow \beta \left(1 - e^{-\frac{T_s}{\tau_{ch}}} \right) = e^{\frac{T_s}{\tau_{ch}}} - e^{\frac{2T_s}{\tau_{ch}}} \rightarrow \quad (7.4)$$

$$\beta = \frac{e^{\frac{T_s}{\tau_{ch}}} - e^{\frac{2T_s}{\tau_{ch}}}}{1 - e^{-\frac{T_s}{\tau_{ch}}}} = \frac{e^{\frac{T_s}{\tau_{ch}}} \left(1 - e^{\frac{T_s}{\tau_{ch}}} \right)}{\left(1 - e^{-\frac{T_s}{\tau_{ch}}} \right) \left(1 + e^{\frac{T_s}{\tau_{ch}}} \right)} = \frac{e^{\frac{T_s}{\tau_{ch}}}}{\left(1 + e^{\frac{T_s}{\tau_{ch}}} \right)} = 1 - \left(1 + e^{\frac{T_s}{\tau_{ch}}} \right)^{-1} \quad (7.5)$$

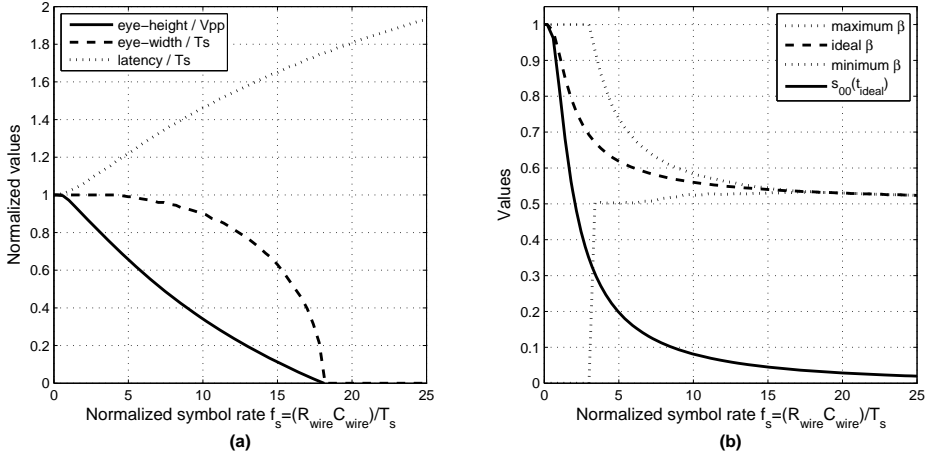


Figure 7.3: Eye properties as a function of data rate with FIR-pre-emphasis binary signaling (a) and s_{00} and the optimum β (b). The boundaries for β at which the eye just closes are also shown in (b).

The maximum value of the response is found at $t=T_s$, which is the ideal detection instant as there is no ISI. The maximum value s_{\max} ($=s_{00}(t_d)$ when $t_d=t_{\text{ideal}}=T_s$) can be found by substituting (7.5) into the first term of (7.4) and reworking the result:

$$s_{\max} = s_{00}(T_s) = \frac{e^{\frac{T_s}{\tau_{ch}}} \left(1 - e^{-\frac{T_s}{\tau_{ch}}} \right)}{\left(1 + e^{\frac{T_s}{\tau_{ch}}} \right)} = \frac{\left(e^{\frac{T_s}{\tau_{ch}}} - 1 \right)}{\left(1 + e^{\frac{T_s}{\tau_{ch}}} \right)} = 1 - 2 \left(1 + e^{\frac{T_s}{\tau_{ch}}} \right)^{-1} \quad (7.6)$$

In Figure 7.2b, the results of equation (7.5) and (7.6) are shown graphically. When the data rate relative to the time constant of the channel (τ_{ch}/T_s) rises above about 0.2, then equalization starts to become interesting and the ideal β starts to drop. As a consequence, the maximum sample value s_{\max} goes down, to only one tenth of the original value when the data rate is five times faster than the time constant of the channel. Without equalization, the eye would be closed at $\tau_{ch}/T_s=1/\ln(2)=1.44$, as was shown in equation (6.19), so FIR pre-emphasis can indeed significantly increase the achievable data rate. It does however come at the cost of reduced sample values at the receiver, as long as we assume that the transmitter swing is limited to a fixed value (e.g. the supply voltage). So with aggressive de-emphasis, other error-sources at the receiver, such as offset and ICI, will also have to be low to guarantee reliable signaling. Also, with an actual (on-chip) channel, higher-order time constants will start to have an effect at high τ_{ch}/T_s values, which will introduce ISI that can not be compensated with the simple two-taps FIR pre-emphasis.

7.3.2 Achievable data rate with FIR pre-emphasis for on-chip wires

Figure 7.3a shows the eye properties for FIR pre-emphasis using the third-order on-chip wire model from Table 3.3 and numerical analysis to find the optimum β and corresponding eye-properties. Now, the eye closes due to the higher-order time constants, but only at a rate of $18.1 \text{ bit}/R_{\text{wire}}C_{\text{wire}}$, which is nearly six times higher than without equalization (see section 6.2).

The optimum β , as is plotted in Figure 7.3b, equals the optimum predicted by the first-order model with less than 1% difference when $0.41 \cdot R_{\text{wire}}C_{\text{wire}}$ is substituted for τ_{ch} , so the first-order model is indeed accurate.

Boundaries for the β , at which the eye just closes, were also computed and plotted in Figure 7.3b. At low rates where un-equalized signaling is also still possible, the allowed values for β cover the complete range from 0 to 1 (at $\beta = 0$ or $\beta = 1$, the symbol will actually just be plain PAM). Not unexpectedly, the margin in the β reduces as the data rate approaches the maximally achievable data rate.

The same analysis was also carried out with a wire with either resistive receiver termination or capacitive transmitter termination, using the third-order model from Table 4.1. For this wire model, the eye properties highly resemble those from Figure 7.3 (and are therefore not repeated in a figure) but with the x-axis scaled, such that the achievable rate increases from $18.1 \text{ bit}/R_{\text{wire}}C_{\text{wire}}$ to $26.1 \text{ bit}/R_{\text{wire}}C_{\text{wire}}$. So there is an increase in data rate compared to conventional termination but only a factor 1.4. Without equalization, the special termination schemes give an achievable data rate of $8.8 \text{ bit}/R_{\text{wire}}C_{\text{wire}}$, 2.8 times higher than with conventional termination as was discussed in section 6.2. The higher order components in the wire transfer reduce the additional gain in achievable data rate that simple first-order equalization can bring on top of the gain from capacitive transmitter or resistive receiver termination.

To further increase the achievable data rate, higher-order FIR-filtering could be used, but at the cost of a continued reduction in absolute eye-height as the symbol period shortens (because $s_{00}(t_{\text{ideal}})$ reduces). In on-chip communication, a significant eye-height is needed at the receiver to guarantee a suitable noise-margin (at acceptable power-levels). This limits the benefits of higher-order FIR pre-emphasis and it was therefore not investigated further.

7.4 Pulse-width pre-emphasis

As an alternative to FIR pre-emphasis, pulse-width (PW) pre-emphasis was developed [33, 73]. As shown in Figure 4.2a, PW pre-emphasis can greatly reduce the amount of ISI by using the second part of the symbol-time to compensate for the remaining line charge. An advantage of a PW pre-emphasis circuit is that it only needs to switch between two voltage levels (for 2-PAM signals), which allows the use of simple transmitters (e.g. inverters) and reduces the influence of finite slew rates. The emphasis on timing accuracy instead of amplitude accuracy (for conventional pre-emphasis) also facilitates the scaling to future deep-submicron, high-speed, low-voltage CMOS technologies. A drawback of PW pre-emphasis is the fact that the power consumption does not scale with data activity, as the transition inside the symbol always consumes power. PW pre-emphasis was also

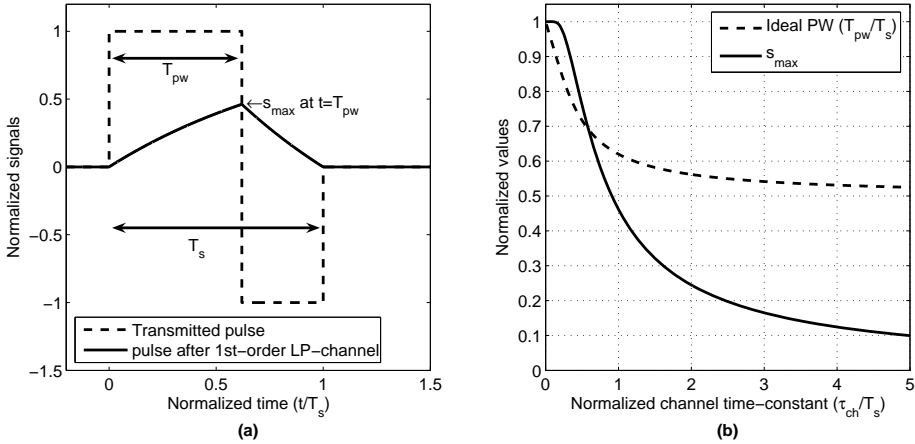


Figure 7.4: Symbol response example at $T_s = \tau_{ch}$ (a) and ideal pulse-width and s_{max} (b) for pulse-width pre-emphasis with a first-order channel.

successfully used for wireline communication [39, 44], where this last drawback is of less concern.

Before going into the details of PW pre-emphasis, it is interesting to note that another form of binary signaling with pulse-width control was proposed in [20]. There, current pulses with a duration shorter than the symbol time are used to transmit zeros (the line remains static for ones). The shorter pulse spread out the energy of the signal over a wider bandwidth, but the pulses are not shaped specifically to emphasize those parts of the wire transfer that are attenuated by the channel. The ISI is therefore not canceled, only reduced together with the signal amplitude itself which is a disadvantage compared to the PW pre-emphasis discussed here.

7.4.1 PW pre-emphasis with first-order channel models

For first-order low-pass channels, PW pre-emphasis can completely cancel the ISI, as shown in Figure 4.2a. The required pulse-width for zero ISI is a function of the symbol-time (T_s) and of the time constant of the channel (τ_{ch}) and can again be found by writing the transmitted pulse shape as a summation of three step responses:

$$g_0(t) = \text{step}(t) - 2\text{step}(t - T_{pw}) + \text{step}(t - T_s) \quad (7.7)$$

With the exponential functions for first-order step response terms, the symbol response can be written as:

$$s_{00}(t) = g_0 * h_{00} = \left(1 - e^{-\frac{t}{\tau_{ch}}} \right) \Big|_{t \geq 0} - 2 \left(1 - e^{-\frac{-t + T_{pw}}{\tau_{ch}}} \right) \Big|_{t \geq T_{pw}} + \left(1 - e^{-\frac{-t + T_s}{\tau_{ch}}} \right) \Big|_{t \geq T_s} \quad (7.8)$$

The response will be zero for $t \geq T_s$, meaning no ISI, on the following condition:




Figure 7.5: Eye properties as a function of data rate with PW-pre-emphasis and binary signaling (a) and s_{00} and the optimum PW (b). The boundaries for PW at which the eye just closes are also shown in (b).

$$e^{-\frac{t}{\tau_{ch}}} \left(-1 + 2e^{\frac{T_{pw}}{\tau_{ch}}} - e^{\frac{T_s}{\tau_{ch}}} \right) = 0 \rightarrow \frac{T_{pw}}{\tau_{ch}} = \ln \left(\frac{1}{2} + \frac{1}{2} e^{\frac{T_s}{\tau_{ch}}} \right) \rightarrow \quad (7.9)$$

$$PW = \frac{T_{pw}}{T_s} = \ln \left(\frac{1}{2} + \frac{1}{2} e^{\frac{T_s}{\tau_{ch}}} \right) \frac{\tau_{ch}}{T_s} \quad (7.10)$$

So, as with FIR pre-emphasis, the ideal pulse-width is a function of the ratio between T_s and τ_{ch} . The maximum value of the response is found at $t=T_{pw}$, which is the ideal detection instant as there is no ISI. The maximum value s_{max} ($=s_{00}(t_d)$ when $t_d=t_{ideal}=T_{pw}$) can be found by substituting (7.10) into the first term of (7.9):

$$s_{max} = s_{00}(T_{pw}) = 1 - \left(\frac{1}{2} + \frac{1}{2} e^{\frac{T_s}{\tau_{ch}}} \right)^{-1} \quad (7.11)$$

The ideal pulse-width and the resulting s_{max} are shown in Figure 4.2b. Again, for symbol times much smaller than the channel time constant, a high amount of de-emphasis is needed and the ideal pulse-width approaches 50% while s_{max} approaches zero.

Interestingly, the s_{max} is exactly the same as for (symbol-spaced) FIR. So, although the equalization methods are quite different, the results at the sample time are the same. Do note that the eye-diagram at times other than the sample time does differ between PW and FIR pre-emphasis, which can also be seen in Figure 7.1. For first-order channels, PW pre-emphasis will have a zero value at the receiver at $t=T_s$, while for FIR pre-emphasis, the value at this point depends on the current and next symbol.

7.4.2 Achievable data rate with PW pre-emphasis for on-chip wires

Figure 7.5a shows the eye properties for PW pre-emphasis using the third-order on-chip wire model from Table 3.3. Now, the eye closes at a rate of $22.8 \text{ bit}/R_{\text{wire}}C_{\text{wire}}$, which is more than seven times higher than without equalization (see section 6.2).

The optimum PW, as is plotted in Figure 7.5b, again equals the optimum predicted by the first-order model with less than 1% difference (when $0.41 \cdot R_{\text{wire}}C_{\text{wire}}$ is substituted for τ_{ch}), as was the case for FIR pre-emphasis. The PW boundaries also show the same familiar shape.

The analysis with the resistive RX or capacitive TX terminated wire model also gave the same conclusions as for FIR; again the eye properties highly resemble those from Figure 7.5, with the x-axis scaled by a factor of 1.4.

As for FIR, higher order PW pre-emphasis could in theory further increase the achievable data rate. Instead of only one pulse edge, multiple edges could be used to create more degrees of freedom. This ‘multitap PWM pre-emphasis’ is discussed in more detail in [44]. It was not used in this project as it is more complex to implement and, as with multi-tap FIR, a spectacular increase in data rate is not expected.

More results for PW pre-emphasis with actual interconnects will be given in Chapter 8, including an additional picture of the sensitivity to PW variations and a discussion of an actual implementation.

7.5 FIR versus PW pre-emphasis

7.5.1 Differences for on-chip and off-chip applications

The previous two sections showed that, for on-chip communication, FIR or PW pre-emphasis have very similar effects. For PW pre-emphasis, the eye remains open up to higher data rates, but that is in a region where the absolute eye opening is already very small, as is visible in Figure 7.6 on the next page. For the most part, the absolute eye-opening is nearly the same for both types of pre-emphasis, as was also predicted by the first-order model.

This is quite a different situation than what is described in [44] for off-chip wires. There it was found that PW pre-emphasis fitted better to wire line channel transfers than FIR pre-emphasis. However, as off-chip wires have no dominant pole, as was discussed in section 4.2.1, pulse-width pre-emphasis can not entirely nullify the long tail of the characteristic ‘diffusion’ step-response, nor can any other simple first-order equalization scheme. The complex shape of that tail, which is not easily captured by simple step responses from low-order transfer functions, requires either high-order equalization schemes or coded data. 8B/10B codes [147] are for example often employed in off-chip communication, to minimize the low-frequency content in the data stream and avoid the so-called ‘baseline wander’ effects from the long tail in the step response.

What was also observed in [44] is that half-symbol spaced FIR pre-emphasis [105] has quite some spectral resemblance to PW pre-emphasis. This is because the ‘overdrive’ part




Figure 7.6: Comparison of eye-heights for FIR and PW pre-emphasis.

of the pulse is only half a symbol wide. However, because the pulse is more narrow, the absolute eye height will also be lower, which makes it less interesting than PW pre-emphasis (especially in on-chip communication as absolute eye-height is of concern for e.g. offset).

For on-chip communication, a favorable aspect of PW pre-emphasis is the lower latency compared to (symbol spaced) FIR. In many digital circuits, delay is even the most important parameter, so the advantage can be quite significant.

An already briefly mentioned drawback of PW pre-emphasis is the ‘static’ power that is consumed as the line keeps switching even when there is no change in data. In theory, FIR pre-emphasis could have a benefit here. But power in a FIR pre-emphasis scheme will also depend heavily on the implementation of the transmitter, as is discussed next.

7.5.2 Implementation differences

Current-summing transmitters that are usually used for off-chip FIR pre-emphasis schemes (overdrive signaling) [97] are less ideal for on-chip implementation because they have a high driver impedance which lowers the bandwidth. Nevertheless, such pre-emphasis transmitters were presented in recent literature [63, 64, 71, 72] and even showed good results in combination with resistive receiver termination. The power overhead in the transmitter can be kept to a minimum when the current is not actually summed in the transmitter, but instead the data and currents are re-coded into a set of sources of which only one is turned on at a time [71, 72].

But a low-ohmic driver impedance is desired to obtain a high interconnect bandwidth, which requires either voltage-mode transmitters or capacitive transmitters. The option with the capacitive transmitter is discussed in the next sub-section. Voltage-mode overdrive transmitters have the drawback that they require the availability of additional supply voltages or require a large static current to create low output impedance. A transmitter circuit of the latter type was investigated for the first chip in this project, but was not used, because at realistic bias currents, finite slew-rate limited the equalization performance of the low-ohmic driver circuit.

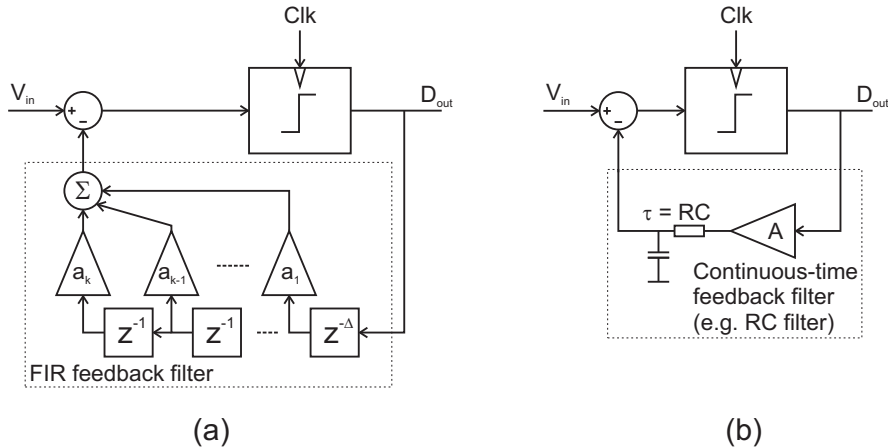


Figure 7.7: Schematics of a conventional discrete time decision feedback equalizer (a) and a continuous time alternative (b).

In this respect, PW transmitters are favorable, as they switch only between two voltage levels, which make them ideally suited for ‘digitally oriented’ deep sub-micron processes. The pulse-width scheme is also not difficult to implement, as the PW-pre-emphasis circuit in the next chapter shows.

7.5.3 FIR pre-emphasis and capacitive transmitters

One type of a pre-emphasis transmitter that can circumvent most drawback above, is the combination of FIR pre-emphasis with capacitive transmitter termination (see section 4.2.2). This combination was not tested in this project, but was introduced in [74] (and discussed in [121]). The use of different capacitors for the different taps of the filter enables a very simple transmitter that should scale well to future technologies. In [74, 121], an additional proposal was to vary the delay of the second tap to other values than simply one symbol (similar to fractionally spaced equalization, as mentioned earlier [95]). Not much data rate gain was observed from the FIR pre-emphasis in that particular experiment, but this seemed to be due to limitations in the test setup. In principle, simple 2-taps symbol-spaced FIR should give the same benefits as those discussed in section 7.3, assuming that the capacitors are given the correct values that match to the ideal beta.

7.6 Decision feedback equalization

As was discussed in section 7.2, a form of receiver-end equalization that has become very popular in wireline and backplane communication is Decision Feedback Equalization (DFE). The technique itself is already quite old [148] and it also has a rich history in other application such as hard-disk equalization [149], and many other applications.

DFE, as shown in Figure 7.7, uses previously detected bits to remove (post-cursor) ISI from the current sample. The fact that DFE uses previous decisions to filter the ISI has the advantage that it can boost the high-frequency signal component without amplifying noise or crosstalk components, this in contrast to linear equalization filters, as is explained in [112]. For noise-limited channels, this can give DFE an advantage in achievable data rates

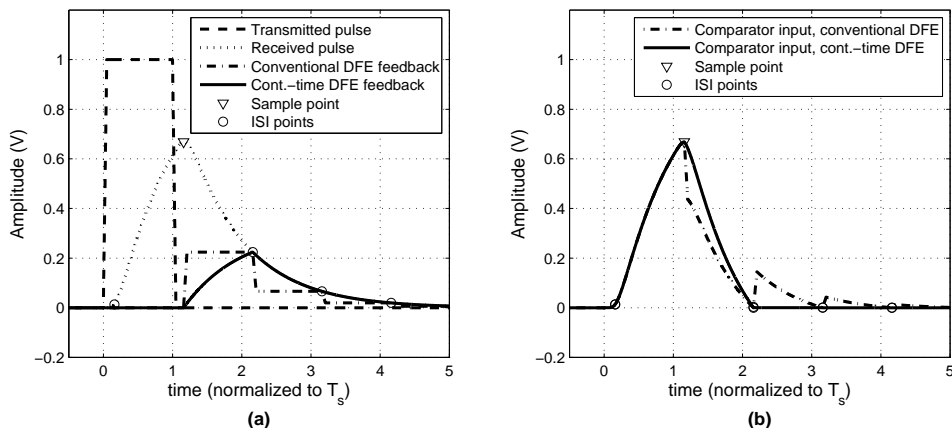


Figure 7.8: Symbol responses with different forms of DFE. Signals are shown at the summing node of the receiver(a) and at the input of the comparator (b).

[95]. Often, run-time adaptation is used to tune the coefficients of the DFE to the actual channel.

The feedback-filter that is most often used in DFE is a finite-impulse-response (FIR) filter, as shown in Figure 7.7a. However, the use of a FIR filter as feedback filter is not always the best choice, as many low-pass channels have a response with only one or two dominant poles. A simple analog or IIR filter can easily reproduce such a response. A FIR filter however needs many taps to cancel the ISI from the long (exponentially decaying) tail when the frequency of the dominant pole is low compared to the sample frequency.

7.6.1 DFE with continuous-time feedback filter

In [150] it was proposed to use IIR filters for the DFE in a fast Ethernet receiver, which needed only one to two taps instead of the eight to twelve that would be required if a FIR filter had been used. Another alternative that was already proposed quite some time ago in [151] (and recently applied in [139]) is to use a combination of a FIR and IIR filter. This enables both cancellation of high-frequency artifacts with the FIR and cancellation of dominant poles with the IIR.

However, in [150] and [151], discrete-time IIR filters are still used, while it is also possible to use an analog continuous-time filter. In this project (and later also in [139]), a continuous-time filter was used as feedback filters, as these types of filters can have simple implementations with good power-efficiency. A schematic of such a DFE is shown in Figure 7.7b. In the figure, a simple first-order filter is shown as an example, but more complex filters are also possible.

To illustrate its behavior, next to the behavior of the conventional FIR-feedback DFE, symbol responses are shown in Figure 7.8a. It is clear that the received pulse contains a lot of ISI and detection without equalization would not be possible. It is also visible that the response has a long tail, so the ISI is distributed over many samples. For the conventional FIR DFE, this would mean the requirements of a lot of taps, of which the first three are shown in the figure. In contrast, a simple first-order low-pass filter is sufficient to almost

completely cancel the ISI, because of the dominant first-order pole of an on-chip wire. Figure 7.8b shows that the ISI is indeed canceled at the input of the quantizer – provided of course that the correct filter parameters are used. It also shows that the analog feedback filter not only cancels ISI at the sample instants, but also in between these instants, which will ensure a wide eye-opening. Conventional discrete-time feedback would need additional complexity to also minimize ISI at the left and right edges of the eye [119].

Note that Figure 7.8a shows an idealized situation, with virtually no delay for the discrete-time feedback. In an actual case, there will be some delay (Δ , as schematically shown in Figure 7.7a), as the components in the loop are not infinitely fast. In high-speed DFE circuits, it is problematic to get this delay below the symbol time, which is why techniques such as loop-unrolling [134, 140] are used in many DFE circuits. Loop-unrolling could be combined with a continuous-time filter, if the filter would only be used to cancel the ISI of the 2nd and later post-cursor sample-points (as in [139]). However, in this project, timing closure was not a big problem, as a fast comparator was used and the number of delaying components in the loop was kept to a minimum (as discussed in section 10.3.2). Any delay that is there (such as the comparator delay) does not have to be problematic, as long as it is taken into account in the feedback factor A - by using $T_s' = T_s + \Delta$ in equation (7.12) below.

To get the perfect cancellation of ISI as shown in Figure 7.8, the continuous-time DFE circuit requires that two parameters are set correctly: 1) the feedback time constant τ , which should equal the dominant time constant of the channel and 2) the feedback gain-factor A , which value depends on the data rate. In contrast, the value of every tap of the conventional FIR DFE feedback has to be adapted as a function of the data rate, to get proper ISI cancellation.

7.6.2 Continuous-time DFE with first-order channel models

For first-order channel models, τ and A can be determined analytically, as was done earlier for pre-emphasis parameters.

The feedback time constant τ is easy to set: it should equal the (dominant) time constant of the channel. Only then will the response of the feedback perfectly follow the tail of the symbol response, as in Figure 7.8a.

For the gain factor, the analysis of the FIR pre-emphasis from section 7.3 can be re-used, because the requirement for the gain of the feedback pulse is exactly the same as for the gain of the second tap of the FIR filter: the magnitude of these compensating pulses should be such that they exactly compensate the tail of the response of the original symbol. The difference is that with pre-emphasis, the pulse is filtered by the channel and with DFE, the pulse is filtered by the feedback filter. For first-order channels, these filters are equal (assuming τ is set equal to τ_{ch}).

So the optimal β from section 7.3 can be used to derive the optimum A . For the FIR pre-emphasis, the second pulse with its magnitude $-(1 - \beta)$ is used to cancel the ISI from the first pulse that has magnitude β . But with DFE, the transmitted pulse has a magnitude of 1 (assuming a normalized transmitter), so we have to scale the magnitude of the compensation pulse by β . With equation (7.5) and some substitution, this produces:

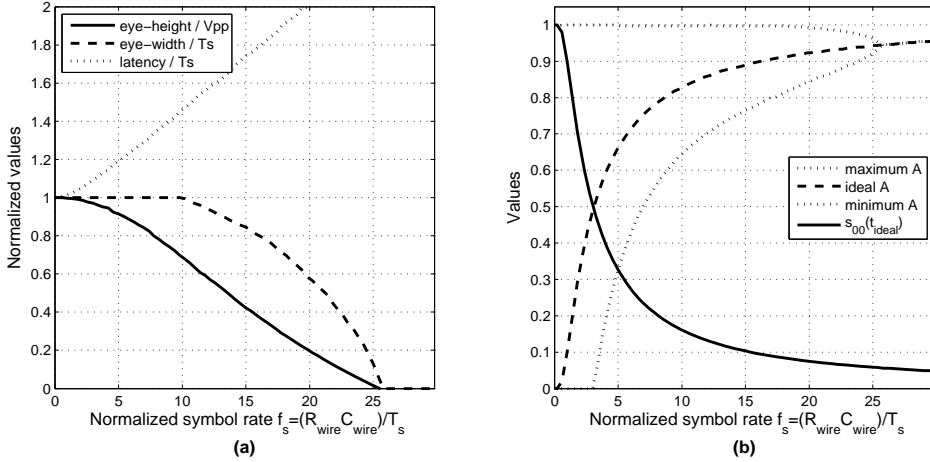


Figure 7.9: Eye properties as a function of data rate with DFE and binary signaling (a) and s_{00} and the optimum and boundaries for the feedback gain A (b).

$$A = \frac{1 - \beta}{\beta} = e^{-\frac{T_s}{\tau_{ch}}} \tag{7.12}$$

This is a simple and intuitive result. It could also have been obtained by realizing that the tail of the original pulse decays with $e^{-t/\tau}$. When we want to cancel this tail by a second pulse that has the same shape, but is delayed by one symbol time T_s then we should scale this second pulse by $e^{-T_s/\tau}$.

7.6.3 Achievable data rate with continuous-time DFE for on-chip wires

Figure 7.9a shows the eye properties with continuous-time DFE, using the third-order on-chip wire model from Table 3.3. Now, the eye closes at a rate of 25.5 bit/ $R_{\text{wire}}C_{\text{wire}}$, more than eight times higher than without equalization (see section 6.2). This result is also better than the 18.1 bit/ $R_{\text{wire}}C_{\text{wire}}$ obtained with FIR, or the 22.8 bit/ $R_{\text{wire}}C_{\text{wire}}$ from the PW pre-emphasis.

This benefit is obtained because the feedback pulse from the DFE is of true first-order (at least in the model), meaning that it can converge perfectly with the tail of the symbol response, already right at the first post-cursor ISI point, which can also be seen in Figure 7.8a. The FIR pre-emphasis filter can not perfectly cancel the first post-cursor ISI, as the peak of the compensating pulse is rounded off due to the higher-order effects of the channel. In practice, this benefit of the DFE over FIR will be less pronounced because the pulse for the DFE feedback filter will have some delay and the filter will contain some parasitic higher-order poles.

A benefit that will remain is that the received swing is higher for DFE ($s_{00}(t_{\text{ideal}})$ is larger), which creates more eye opening at lower data rates than with FIR (or PW) pre-emphasis. At a rate of 5 bit/ $R_{\text{wire}}C_{\text{wire}}$ for example, the $s_{00}(t_{\text{ideal}}) = 0.33V$ for DFE and the eye opening is




Figure 7.10: Tolerances in τ for the DFE feedback filter.

0.3V. For FIR pre-emphasis at this rate, the eye-opening is only 0.15V. This is because with DFE, the transmitter transmits the PAM symbols with full swing, while with FIR, the symbol magnitude (β) is reduced to create room for the overdrive.

A drawback for the DFE, compared to the FIR pre-emphasis, is that it requires setting of two parameters: the gain A and the time constant τ . With FIR, the time constant for the compensation is the channel itself.

The optimal value for A is plotted in Figure 7.9b. At high data rates, the optimal value is about 5% to 10% higher than the value that is predicted by the first-order model from equation (7.12). This is because the higher-order components in the wire transfer delay the first-order decay of the tail.

There is quite some tolerance in the optimal value for A , as can be seen in Figure 7.9b, but this figure does assume an optimal filter time constant τ . The tolerance for the time constant is plotted in Figure 7.10 (using the ‘ideal- A ’ curve from Figure 7.9b to set A).

The above analysis was also repeated with the resistive RX or capacitive TX terminated wire model from Table 4.1. When the time constant of the DFE was again set equal to the first-order time constant of the wire, then the data rate increased to 33.6 bit/ $R_{\text{wire}}C_{\text{wire}}$. With this termination, the performance of the DFE suffered from the higher-order behavior of the wire and the optimum gain-factor is a bit of a compromise, with a compensation pulse that has a bit too low magnitude to cancel the first post-cursor ISI and too high magnitude to cancel the rest of the tail. A small increase in the τ of the feedback filter improved the compromise: a τ for the feedback filter that is 8% larger than the dominant time constant was the optimum, with an achievable data rate of 37 bit/ $R_{\text{wire}}C_{\text{wire}}$. But, these small optimizations in τ set aside, on average, the increase in data rate is again a factor 1.4 over the case without special termination, just as with FIR and PW pre-emphasis.

In this project, in the actual transceiver circuit, the τ of the filter was fixed at design time, as the time constant of the on-chip wire is known quite well in advance, and quite some mismatch between the filter and the channel time constant is allowed, as can be seen from

Figure 7.10 and as also will be discussed in the next section. The A was set at runtime, for experimentation and to be able to turn the equalization off.

For applications with more variable channel time constants such as backplane and wireline transceivers, the continuous-time DFE can still be quite applicable, as it is not difficult to make feedback filters with adaptable gains and adaptable time constants [139]. For such channels, a combination of a discrete-time first DFE tap to cancel the first (higher-order) part of the ISI and a continuous-time filter for the long tail seems quite promising [139].

7.7 Equalization and process spread

This section discusses what the effects are of deviations between intended and actual parameters, such as the wire time constant. The RC time constant of a wire can for example easily deviate with +/-40% due to tolerances in the manufacturing steps.

7.7.1 Dealing with mismatch at design time

When equalization schemes are used without any tuning of the equalization parameters, except at design-time, then one needs to take the possible mismatch between the intended equalization and the actual wire time constant in to account.

To assess the allowed tolerances, the minima and maxima for the equalization parameters can be used, as plotted in Figure 7.3 for FIR, in Figure 7.5 for PWM and Figure 7.9 and Figure 7.10 for DFE. It is not surprising that the tolerance for the equalization parameters (the difference between the minimum and maximum value) decrease for higher data rates. What is however interesting and useful, is that the minimum value of the equalization parameters are monotonically rising functions of the data rate and the maximum values are monotonically falling, right up to the point where they converge to the ideal value (where the eye closes). This is true for all three tested equalization methods.

So, when a certain setting of the equalization parameter is adequate for a certain data rate, then that same setting is also adequate for all lower data rates. This implies that one can set the equalization parameter at design time, based on the worst-case (maximum) value of the $R_{\text{wire}}C_{\text{wire}}$ time constant, instead of being based on the expected (mean) value.

The implication is that in the nominal case, the eye will not be as open as it could have been, as the equalization is not optimal, it is however now robust against wire time constant deviations.

Mismatch in low-frequency time constants

There is also another type of mismatch that can manifest itself specifically in the capacitive transmitter, and that is mismatch between the low-frequency path, with the transfer defined by $G_m \cdot R_L$ and the high-frequency path, defined by $C_s / (C_s + C_{\text{wire}})$, as was discussed in section 2.6.2. Mismatch between these paths will result in a pole and a zero in the transfer that no longer cancel each other. This is very much the same as with mismatch in equalization, where the equalizer also tries to cancel the poles of the wire, which will not be perfect in the presence of mismatch. A difference is that mismatch in low-frequency time constants will result in responses with very long tails (so called 'slow settling component'). The worst-case ISI depends on the relative difference between the pole and the zero time constants. When the pole time constant is a factor two higher than the zero time constant,

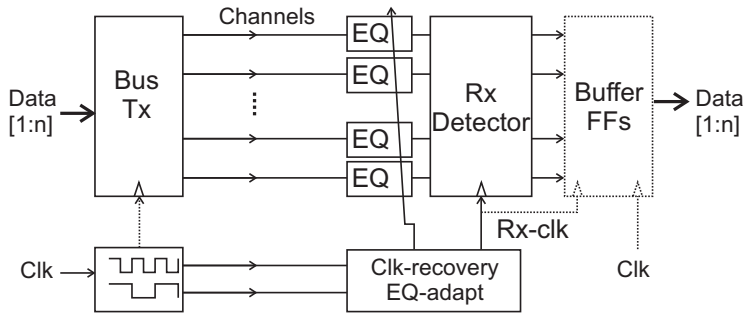


Figure 7.11: Block diagram of bus transceiver with adaptive equalization.

then the energy will be divided in two equal parts: half for the PAM pulse itself (s_{00}) and half for the tail (ISI_{max}), which will just close the eye. So, as long as the time constant of the compensating zero deviates less than a factor two from the pole, the eye will still be open.

7.7.2 Adaptive equalization for on-chip transceivers

Setting the equalization parameter at design time, based on the worst-case wire time constant is one approach, but does require that the equalization parameter itself is well defined and that the receiver can cope with smaller than optimal eye-openings due to imperfect equalization.

Another approach would be to use circuits with equalization parameters that somehow track the properties of the wire, either with a feedforward approach (copying of the wire time constant) or with a feedback approach (adaptive equalization). A feedforward approach could for example be PW equalization with a circuit that translates the phase-shift of a dummy interconnect into a pulse-width that is approximately equal to the optimum from equation (7.10).

Adaptive equalization that uses feedback from the receiver to change the equalization parameters is a well-known technique for backplane and wireline communication, as was discussed in section 7.2.4. But, as far as the author is aware, such adaptation algorithms have not yet been implemented for on-chip transceivers, as they tend to be quite complex and area consuming.

One advantage of an on-chip bus that can help the future introduction of adaptive equalization is the good match that can be obtained between delays of different wires in a bus (provided the bus is laid out properly). This means that a single adaptation algorithm can be used to set the parameters for all the receivers in the bus. Such an adaptation algorithm might also be used to control the sampling phase of the receiver, similar to the combination of CDR and adaptive EQ in off-chip transceivers, as was discussed in section 7.2.5.

A high-level block diagram how such an adaptive equalization can look like is shown in Figure 7.11. Assuming there is a good match between the delays of the various wires in the bus, a few of the wires can be used for the transmission of dedicated periodic patterns that are used by the adaptation algorithm and clock recovery.

Simple patterns are for example an alternating 10101010... sequence together with a 11001100... sequence. The 10101010... pattern is suitable for multiple purposes. First, it can serve as a half-rate clock as in other simple source-synchronous transmission schemes (see section 11.5). Second, it can be used for comparison with the other sequence: If the magnitude of the received symbols of the 10101010... sequence is different (smaller or larger) than the symbol magnitudes of the 11001100... sequence, then the equalization can be adapted (increased or decreased respectively) to reduce the ISI.

More complex patterns that include more wires can also be envisioned. Patterns with staggered switching edges could for example be combined in the clock-recovery circuit to re-create a clock without needing the transmission of a half-rate pattern. This can be of value for those cases where the attenuation or phase-shift of a half-rate clock becomes problematic.

Assuming that the bus is wide (large N), the wire overhead for the transmission of these additional signals can be kept low. The circuits for the clock-recovery and EQ adaptation can also be kept much simpler than normal adaptation circuits as they do not have to distill parameters from random data.

Such adaptive algorithms were not yet used in the actual transceivers build in this project. Instead, in our test-chips the equalization parameters were defined externally during the measurements, to be do more extensive experiments and get a better feeling for tolerances in actual circuits. The results obtained (as discussed in the next chapters) show that simple practical circuits can indeed achieve significant gains in achievable data rates, without explicit need for tuning.

7.8 Equalization combined with M-PAM

It was shown in the previous chapter that the more advanced signaling techniques such as multi-level signaling, pass-band signaling and multi-channel signaling have little benefits for on-chip communication when applied by themselves. The only technique that showed a very small increase in achievable data rate was multi-level signaling. In this section it is shortly investigated whether the benefit of multi-level signaling can improve when it is combined with equalization.

To this end, the symbol response parameters that were obtained from the quantitative equalization analysis (specifically $s_{00}(t_{ideal})$ and $s_{emax}(t_{ideal})$) in the previous sections have been re-used in combination with equations (5.17), (5.18) and (5.21) to compute the eye-height for M-PAM with equalization.

An example of this analysis is shown in Figure 7.12, for the case of M-PAM in combination with FIR pre-emphasis over a conventionally terminated wire. Compared to M-PAM without equalization, as was shown in Figure 6.8, the achievable data rate is much higher, but the curves also have a different shape: first there is a small flat part, where no equalization is necessary, then there is a steep decrease because the increase in equalization reduces the swing, and finally there is a more slowly decaying tail up to the point where the eye-height passes through zero (the eye closes).




Figure 7.12: Eye-height for M-ary PAM signaling in combination with FIR pre-emphasis over an on-chip interconnect with conventional termination as a function of the normalized symbol-rate (a) and bit-rate (b). Compare to Figure 6.8

Interestingly, Figure 7.12b shows that 2-PAM has the largest eye opening and highest achievable data rate (of 18.1 bit/ $R_{wire} C_{wire}$) when it is combined with FIR pre-emphasis, with the achievable data rate monotonically decreasing when the number of level increases (down to 11.3 bit/ $R_{wire} C_{wire}$ for 8-PAM).

However, the same is not true for PW pre-emphasis. Analysis showed that for PW pre-emphasis, 3-PAM outperforms 2-PAM for data rates above 6.8 bit/ $R_{wire} C_{wire}$ and even more levels are favorable at higher data rates. The differences are not large (which is why no additional figure is shown), with achievable data rates of 22.7 bit/ $R_{wire} C_{wire}$ for 2-PAM up to 25.7 bit/ $R_{wire} C_{wire}$ for 8-PAM, but the tendency is just the other way as for FIR.

Analysis with DFE equalization showed yet another slightly different behavior, with 3-PAM becoming favorable above 18.1 bit/ $R_{wire} C_{wire}$, but with 4-PAM just showing the highest achievable data rate of 27.2 bit/ $R_{wire} C_{wire}$, compared to 25.5 bit/ $R_{wire} C_{wire}$ for 2-PAM and 26.5 bit/ $R_{wire} C_{wire}$ for 8-PAM.

These different outcomes for different forms of equalization can be explained because the tail of the eye-opening is slightly different, as was shown in Figure 7.6 for FIR and PW pre-emphasis. The small differences in the slope of this tail can just tip the balance between 2-PAM or multi-level having the highest achievable data rate.

Very similar results (but at higher data rates) were found for the case of M-PAM with equalization in combination with resistive receiver or capacitive transmitter termination. Apparently, it does not make a big difference that higher-order terms have more influence in the transfer of the specially terminated wires.

In general, the original conclusion that 2-PAM is better than M-PAM – in terms of eye-opening and simplicity of application – still holds, also when equalization is taken into account. Only in very special circumstances, at particular data rates and with particular equalization schemes, can multi-level have any benefit.

In wireline communication, it has also been observed that plain binary signaling can give higher eye openings than M-ary signaling for about 90% of the channels surveyed [114], when they are both combined with rigorous equalization. And the popularity of multi-level signaling does indeed seem to diminish after the initial boom in [97, 98, 105, 106, 135].

Combinations of other techniques have not been investigated, so it might be that pass-band or multi-channel signaling techniques improve when the line is also equalized. A more important reason to discard such a combination, even in the unexpected case that it could give some data rate advantages, is the complexity of such a transceiver which often also translates to power consumption.

7.9 Summary and conclusions

The list below shortly summarizes the results and conclusions from this chapter:

- For channels that are dominated by a first-order time constant, simple analytical equations predict the required equalization parameters (for first-order FIR, PW or DFE) and eye opening.
- With conventional wire termination, simple first-order equalization can boost the data rate by a factor of 6, 7 or 8 depending on whether FIR pre-emphasis, PW pre-emphasis or continuous-time DFE equalization is used, also see Appendix B.
- The combination of first-order equalization with resistive receivers or capacitive transmitters yields only a factor 1.4-1.5 additional increase in achievable data rate, compared to equalization with conventional termination. This is because the higher-order terms in the wire transfer become more dominant.
- For on-chip interconnects, it is possible to set the equalization parameters at design time and still be robust against process variations. Adaptive equalization can be used to improve the eye-openings for speed critical applications. Whether (or when) the additional complexity of such adaptation schemes is justified is left as a topic for future study.
- Multi-level signaling has no additional merits when combined with equalization. With FIR pre-emphasis in combination with M-PAM, the achievable data rate even decreases for M larger than 2.

In conclusion, simple binary transmission together with equalization seems the most viable option; simple equalization schemes such as PW pre-emphasis or DFE can significantly boost the achievable data rate while the transceiver can remain simple.

Chapter 8

First demonstrator IC

8.1 Introduction

This chapter discusses the first demonstrator IC that was implemented and measured in this project. In this first chip implementation, the focus was on how to improve the data rate across global wires. From a circuit design perspective, a general solution to the limited interconnect bandwidth is the use of repeaters, which make the repeated wire delay linear with length instead of the quadratic dependency of an unrepeated wire [7]. However, the number of repeaters should be kept to a minimum as they cost area and power and make floorplanning more difficult as portions of active area all over the chip have to be reserved for large repeater circuits. Furthermore, the classical approach to repeater insertion [7], using plain buffers/inverters as repeaters has serious limitations for global communication. With plain non-clocked buffers as repeaters, delay optimization requires closely spaced repeaters and delay variations due to crosstalk and due to process variations will accumulate and limit the achievable data rate. With such a classical repeater scheme, only a small portion of the intrinsic data capacity of each line-segment is actually used.

These arguments motivated the search for more advanced solutions that can increase the data rate for a given length or can increase the unrepeated wire length for a given data rate. At the time of the first chip implementation, the most interesting advancements in on-chip transceivers presented by other research groups were found in [84] and [21]. Low-swing overdrive signaling over differential 10mm aluminum interconnects was described in [84], but with the requirement of a dedicated supply and with clocked switches along the wire (increasing the already troublesome clock-load). [21] Proposed to use 16 μm wide differential wires (20mm long) and exploit the LC regime (transmission-line behavior) of these wires, but at the expense of a significant increase in power consumption and interconnect area. Both papers achieved 1Gb/s/ch in a 0.18 μm CMOS technology.




Figure 8.1: Transceiver system overview.

With the first demonstrator IC in this project, which was implemented in $0.13\mu\text{m}$ CMOS technology, it was shown that a combination of resistive receiver termination and pulse-width pre-emphasis, can significantly boost the achievable data-data rate [33, 73]. This IC also showed that twisted wires are an effective means to cancel crosstalk [81, 82]. With a combination of the techniques, a data rate of 3Gb/s/ch was achieved over 10mm of uninterrupted wire of only twice the minimum pitch. Figure 2.1 shows an overview of the transceiver system.

In the next sections, the first demonstrator is discussed in more detail, starting with a discussion of the interconnects in section 8.2. Section 8.3 revisits pulse-width pre-emphasis and gives an analysis of the ideal pulse-width settings for this transceiver and of the robustness toward parameter variations. Section 8.4 describes the implementation of the transceiver circuits. Section 8.5 compares the transceiver to classical repeater insertion. Section 8.6 discusses the top-level of the IC. The measurement setup is discussed in section 8.7. Section 8.8 shows the experimental results and compares them to predictions. The chapter is summarized and concluded in section 8.9.

8.2 Interconnect analysis and dimensioning

As discussed earlier, the communication structure in this project is assumed to consist of point-to-point buses with all signals traveling in the same direction. For the demonstrator IC, the length of the bus is chosen to be 10mm , to represent a typical global interconnect and allow for easy comparison with other work. The chosen interconnect width of $0.4\mu\text{m}$ and spacing of $0.4\mu\text{m}$ are optimized to give the highest bandwidth per cross-sectional area (BW/Area) as was explained in section 2.7.1.

8.2.1 Interconnect Model

The bus is placed in metal 5 as it is assumed that the thick top-metal (metal 6) is reserved for clock and power routing. The bus model was simulated with a 3D EM-field solver to analyze the behavior of the interconnects and extract distributed RLC parameters. For 10mm long, $0.4\mu\text{m}$ wide wires these parameters are $R' = 0.15\text{k}\Omega/\text{mm}$, $L' = 0.25\text{nH}/\text{mm}$ and $C' = 0.23\text{pF}/\text{mm}$ ($C' = 0.27\text{pF}/\text{mm}$ for differential wires due to Miller-multiplication of the side-plate capacitance).

In the EM-field solver model, metal 4 and metal 6 plates approximate the capacitance of other perpendicular interconnects (assuming a Manhattan routing style), as a large-scale IC usually has a high wire density in all layers. In the actual demonstrator IC, ground- or Vdd-




Figure 8.2: Interconnect transfer functions and crosstalk transfer function from *s*-domain equations of two 10-mm-long single-ended interconnects ($Z_s=50\Omega$).

connected metal stripes were used in metal 4 and 6 to model the capacitance of these other interconnects.

The *s*-parameter equations from [2] were used to plot the transfer functions for single-ended interconnects with both low-ohmic and high-ohmic receiver termination (and 50Ω transmitter impedance). The result is shown in Figure 8.2. Note that skin-effect, as discussed in section 3.6.1 was not taken into account, which explains the flat transfer in the LC region instead of a continued roll-off in Figure 3.5 (another difference is the 50Ω versus idealized transmitter impedance).

Also included in the figure is the crosstalk transfer function from one wire to a direct neighbor. The three regions of the transfer function, with the first region showing the characteristic first-order behavior are also indicated.

The factor three difference in bandwidth between conventional and resistive receiver termination is also visible in Figure 8.2. The -3dB bandwidths shown in the figure are slightly better than obtained by using the first-order ramp-model from equation (3.38) [57], which predicts 85MHz and 220MHz for conventional or resistive receiver termination respectively. The *s*-parameter models predict 100MHz bandwidth for single-ended interconnects and 80MHz for differential interconnects with conventional termination. This last figure is used in subsequent analysis as best approximation of the first-order part of the wire's behavior.

8.2.2 Twisted differential interconnects

In this demonstrator IC, twisted differential interconnects were used to minimize crosstalk and increase robustness, as was discussed in section 4.4. We used a single twist in the even channels and two twists in the odd channels, as shown in Figure 4.13. The single twist in the even channel is placed at 50% of the length, the optimum position for resistive receiver




Figure 8.3: Symbol responses of a (conventionally terminated) on-chip interconnect with 1-ns symbol period using plain binary signaling (a) or PW pre-emphasis (b).

termination (with complete crosstalk cancellation if equal transmitter and receiver resistance are used) [82]. The two twists are placed at 30% and 70% of the length to minimize common-mode crosstalk [82].

8.3 Pulse-width pre-emphasis

Pulse-width pre-emphasis was used on the first demonstrator IC to increase the achievable data rate because it can significantly reduce the inter-symbol interference, as was discussed in section 7.4. That ISI reduction is indeed needed to reach gigabit data rates can be seen in Figure 8.3. With plain-binary signaling, there is clearly too much ISI in the tail of the symbol response in Figure 8.3a. The symbol response with PW pre-emphasis on the other hand shows almost no ISI.

The first-order models from section 7.4.1 can be used to estimate the required pulse-width. The analysis in that section showed that a high amount of de-emphasis is needed when the symbol times are much smaller than the channel time constant and the ideal pulse-width approaches 50% while s_{\max} (the received amplitude) approaches zero. As an example, the ideal pulse-width for 2Gb/s data rate, with a channel-corner frequency of 80MHz ($\tau_{\text{ch}}/T_s = 2\text{ns}/0.5\text{ns} = 4$) is 53% and the receiver swing is only 12% of the transmitter swing.




Figure 8.4: Eye-diagram properties for a 10-mm interconnect with conventional termination ($R_s=50\Omega$, $R_L=\infty$) with plain binary signaling (a) or PW pre-emphasis (b).

To analyze the achievable data rate with higher-order wire models, the analysis from Chapter 5 was used, as was done in section 7.4.2 for PW pre-emphasis, but now with realistic termination impedances and with a lumped RLC-model (100 lumps). Results of the analysis are shown in Figure 8.4 for the case of a 10mm differential interconnect with conventional termination, with wire parameters as measured on the prototype ($R' = 0.19\text{k}\Omega/\text{mm}$, $C' = 0.25\text{pF}/\text{mm}$, $R_{tx}=65\Omega$). Figure 8.4a shows that without pre-emphasis, the eye at the receiver side will be completely closed at rates exceeding 600 Mb/s. With the use of PW pre-emphasis, the theoretically achievable data rate increases to about 4.2Gb/s as shown in Figure 8.4b. However, at a rate of 2Gb/s, the optimal pulse-width is only 53%. At this rate, the higher-order part of the channel transfer decreases the signal swing at the receiver (V_{pp}) to only 6%, instead of the 12% predicted by the first-order model. Both the swing and the eye-opening relative to the swing rapidly decrease further for rates higher than about 2Gb/s and effects such as receiver offset will start to degrade detection, making it nearly impossible to reach the theoretical limit of about 4Gb/s.

If a low-impedance current-sensing receiver is combined with PW pre-emphasis, then the theoretical achievable data rate increases to 6.1Gb/s, a factor of 1.45 increase, similar to the predicted increase with the 3rd order models from section 3.4.2 (In [33], a rate of 7Gb/s was mentioned, but that rate was not yet corrected for the difference between predicted and measured line parameters).

At a given data rate, the higher bandwidth of the resistively terminated interconnects leaves more room for mismatch between the time constant of the symbol-shape and the time constant of the wire. In applications, room for mismatch is necessary to be robust for variations of the line-length, spread in wire parameters and spread in actual pulse-width. The numerical analysis results in Figure 8.5 illustrate how the pulse-width affects the eye-opening at a fixed data rate. For easy comparison with measurements, a data rate of 2.5Gb/s with 150 Ω receiver termination is used and it is visible that the eye remains open for large




Figure 8.5: Eye-diagram properties as a function of pre-emphasis pulse-width for a resistively terminated 10mm interconnect with a data rate of 2.5Gb/s

variations of the pulse-width. At the optimal pulse-width of about 58%, the vertical eye-opening is about 75% of the swing. The actual voltage swing at the detector will be determined by the chosen gain of the current-sensing amplifier. The results in Figure 8.5 again match well to those obtained with the 3rd order wire model (the latency matches when, besides the 3 poles from Table 4.1, the delay of $0.037/R_{\text{wire}}C_{\text{wire}}$ is also taken into account).

8.4 Transceiver Implementation

8.4.1 Transmitter

The schematic of the PW pre-emphasis transmitter that was used on the first IC is shown in Figure 8.6 on the next page, together with some signal waveforms. Conceptually, PW pre-emphasis involves the creation of a clock (Clk_{PW}) with the correct duty cycle and XOR this clock with the incoming data. On the prototype IC, the duty cycle of Clk_{PW} was controlled by an external current source (I_{bias}). The I_{bias} controls the slew rate of the falling edge of the output of an inverter, driven by a normal 50% duty cycle clock (Clk_{Tx}). The controllable slew rate is converted to a controllable falling-edge delay by the second buffer.

The resulting clock with adjustable duty cycle selects either Data or not(Data), thereby implementing the XOR operation. A latch delays the not(Data) by half a clock-cycle to increase the timing margin. With this setup, the signals at the input of the switches (transmission gates) are stable during a transition of the Clk_{PW} and the Tx_p and Tx_n path have matched delay. To drive the wire, a cascade of four scaled inverters was used and the last inverter has an effective output resistance R_{out} of about 60Ω . The size of the differential transmitter is about $300\mu\text{m}^2$. Dynamic latches, low- V_T transistors and small fan-outs (≤ 3) are used to meet the target data rate of 3Gb/s even at high temperature (100 °C) and at the slow process corner. Monte Carlo simulations showed that transistor spread mainly causes common-mode offset with little change in latency and eye-opening. With a transmitter delay of 170ps, the latency of the transmitter and channel amounts to 600ps, as shown in Figure 8.6.

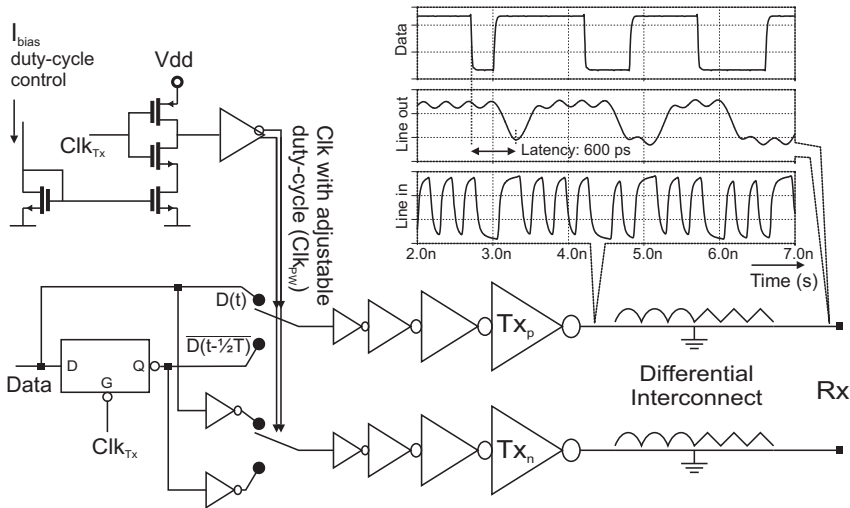


Figure 8.6: PW pre-emphasis transmitter schematic and signal waveforms

The external current provides programmability. With $I_{\text{bias}}=0$, data is transmitted conventionally without PW pre-emphasis. At a 3GHz clock an I_{bias} of 80 μA , 200 μA or 400 μA results in transmitted symbols with pulse-widths of respectively 75%, 58% or 52%. The external control over the pulse-width allows for a comparison of the analysis from the previous sections with the results. In an actual application, the I_{bias} can be fixed at design time as Figure 8.5 and the analysis in section 7.7.1 indicated that the transmission scheme is robust towards (circuit and wire) parameter deviations if the data rate is chosen sufficiently below the theoretical limit. As was discussed in section 7.7.2, an automatic calibration or adaptation algorithm could also be used to set the I_{bias} (for a bus) at run-time, but the benefit of a higher data rate would probably not outweigh the associated costs.

8.4.2 Receiver

The schematic of the receivers is shown in Figure 8.7. The input inverters use transmission gates as selectable feedback resistors, similar to Bashirullah's implementation [66, 70, 93]. In this way, either conventional capacitive termination or (active) resistive termination ($R_{\text{in}} \approx 150\Omega$) can be selected. A regenerative sense amplifier (clocked comparator) followed by a dynamic latch samples the received data and restores it to full-swing.

With the transmission gates turned on, the input inverters behave as transimpedance amplifiers, as was discussed in section 4.2.3. The input impedance is roughly equal to the $1/G_m$. The ratio between the feedback resistance R_{fb} of the transmission gates and the wire resistance R_{wire} controls the voltage gain from the transmitter to the input of the sense amplifier. A similar ratio $(R_{\text{wire}} + R_{\text{fb}})/R_{\text{wire}}$ determines the output equivalent value of the offset voltage. The output-equivalent differential offset voltage of the transimpedance amplifiers is about 7.5mV (one-sigma) and is comparable to the offset of the subsequent sense amplifier, giving a total one-sigma offset of about 10mV. The design of the receiver has been optimized for a balance between minimal offset and maximal speed. The sense amplifier, as shown in Figure 8.7, consists of a differential input pair and a cross-coupled pair with an NMOS reset transistor. The bias voltages V_{b1} and V_{b2} are generated locally

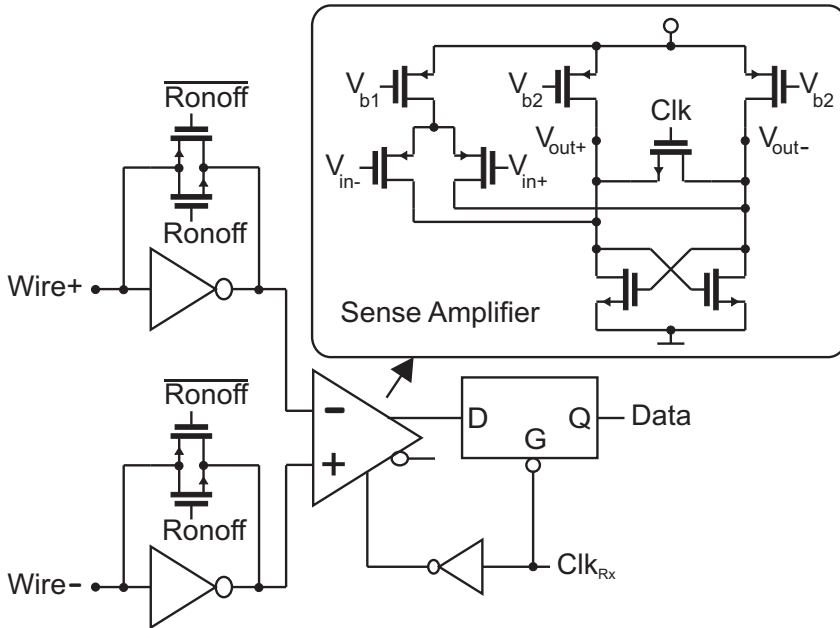


Figure 8.7: Receiver schematic with configurable low- or high-ohmic termination.

with current mirrors from a single resistor-current as their exact value is not critical. The latch after the sense amplifier converts the regenerated data to full-swing data that is stable for a full clock period. Note that this sense amplifier is fast, but does also consume static power and exhibit some hysteresis. In Chapter 9, a different type of sense amplifier is discussed.

The receiver adds only 50ps delay, making the latency of the total transceiver equal to 650ps at 3Gb/s. As with the transmitter, the receiver also uses low- V_T transistors to ensure correct operation at 3Gb/s over different process corners. Only the input inverters are normal- V_T , as a lower overdrive voltage improves the G_m versus current ratio. The total size of the prototype receiver is about $1000\mu\text{m}^2$, not optimized and including the reference circuits.

8.5 Comparison with repeaters

Before discussing the demonstrator IC on which the abovementioned transceiver was implemented, the transceiver concepts are first compared to classical repeater insertion, based on transistor-level simulations.

	Repeater system	PW Tx + resistive Rx
Wire segment length	1mm (near optimal)	10mm
Driver / Repeater NMOS width	20 μ m (optimal)	7 μ m
Energy consumption	3.1pJ/transition	2pJ/bit
Nominal delay (50% crossing)		
PW pre-emp circuit	n.a.	80ps
Transmitter driver cascade	120ps	90ps
Interconnect	680ps (10 segments)	270ps
Total	800ps	440ps
Delay variation due to neighbor-to-neighbor crosstalk	-120ps to +140ps (without shielding)	none
PVT delay spread		
Random mismatch (one-sigma)	20ps	4ps
Slow process, Temp = 90°C, Vdd = 1.1V	+280ps	+100ps
Fast process, Temp = 0°C, Vdd = 1.3V	-150ps	-50ps

Table 8.1: Comparison of first chip transceiver with conventional repeater system (simulation results)

A classical repeated single-ended interconnect system has been simulated in the same technology and Table 8.1 shows a comparison with the presented transceiver. The length of the interconnect segments and the size of the drivers in the repeated system are optimized for minimal delay with the equations from [7] (also see [54] for similar equations, but then expressed in terms of inverter delay). This optimization requires as much as ten repeaters, each with a driver size (20 μ m NMOST width) that is larger than the single driver used in this work (7 μ m NMOST width). Standard-Vt inverters were used as repeaters to reduce the power consumption and the delay penalty was small compared to low-Vt inverters (low-Vt delay = 700ps, power=3.4pJ/transition). The wire dimensions are in both cases equal to the optimized values (0.4 μ m width and 0.4 μ m spacing). This would amount to roughly the same wiring resources per channel for both systems, as the repeater system needs shields between the signal lines to avoid a severe eye degradation. When there would be no shields, then crosstalk between the neighboring wires would create a large dynamically varying delay [152] of as much of 260ps peak to peak, as visible in Table 8.1.

Note that there is an alternative method to reduce crosstalk with repeated wires and that is to alternate the placement of inverting repeaters between neighboring lines [153]. Alternating the inverters results in a partial cancellation of the crosstalk, similar to the alternating twist positions with differential wires in section 4.4. However, alternated inverters require even more positions along the bus of interconnects to be reserved for the repeaters. The average delay with alternating repeaters is also higher than with shielding according to [154], so alternating repeaters were not further investigated in this project.

Although the power consumption of the repeated system is modest (3.1pJ/transition, giving 1.6pJ/bit with 50% data activity), the many repeaters need much layout resources and the long chain of inverters creates a high static variation in delay (430ps) for different process, voltage and temperature (PVT) corners. Without additional measures, the data rate should

be lower than $1/430\text{ps} = 2.3\text{ Gb/s}$ to keep the delay variation within one clock-cycle over all corners.

Larger transistors in combination with fewer repeaters could reduce the effect of PVT and of statistical variations on delay [155, 156], but at the expense of an increase in power or an increase in nominal delay. With a receiver that samples data at the centre of the eye, half a symbol-time after the 50% crossing, the latency is already $800\text{ps}+165\text{ps} = 965\text{ps}$, while the latency of the presented transceiver is only 650ps (including the receiver). The latency of the presented transceiver is also much less affected by PVT variations.

As noted in Table 8.1, a transmitter buffer cascade (also known as an exponential horn) was also part of the tested system, with the same input load as for the PW transmitter. Each stage in the buffer was a factor $f=3$ larger than the previous (as in the PW transmitter). Theoretical optimizations for maximum bandwidth or minimum delay usually predict that $f_{opt}=e$ as is found in e.g. [7]. However, circuit simulations with actual inverters showed that the actual optimum was $f_{opt}=3.6$ for low-Vt inverters and $f=3.85$ for standard-Vt inverters in the $0.13\mu\text{m}$ CMOS technology. That this is higher than e is because an unloaded inverter also has some delay (from its own drain capacitances), which is disregarded in the theoretical optimization from [7]. Perhaps not by accident, the optimum fan-out of 3.6-3.85 is very close to the popular fan-out of 4 (FO4) delay metric, as used in e.g. [8].

In case of the repeated wire system, the buffer cascade at the transmitter can also be combined with logic that precedes the cascade, the so-called combination of ‘logical effort’ and ‘electrical effort’ [89]. The buffer cascade can even be combined with the repeaters itself, to optimize the ‘interconnect effort’. However, the latter approach loses its advantages with newer technologies according to [89] and was therefore not used in the comparison. The former approach does not really reduce delay, but enables additional logic functions to be included without adding delay.

8.5.1 Receiver clocking

In their current form, both systems need special receiver clocking strategies as the latency is higher than one clock-cycle. On the demonstrator IC, the receiver clock is supplied externally to be able to change its phase relative to the transmitter clock and measure the eye-width. In an application, one could use source-synchronous transmission and transmit clock information alongside the data bus, or one could choose to use shorter wire segments and pipeline the communication. In the latter case, the presented system would require far fewer pipeline stages than the conventional (clocked) repeater system as the presented techniques increase the achievable data rate and lower the latency for a given length of interconnect.

8.6 Demonstrator IC top-level

The pulse-width transmitter and resistive receiver were implemented on the externally configurable demonstrator IC to validate the transceiver techniques and to compare results with analysis. The chip has been fabricated in a standard 1.2V, 6M, $0.13\mu\text{m}$ CMOS process with copper interconnects. The floorplan and schematic overview of the demonstrator IC is shown in Figure 8.8 and a micrograph is shown in Figure 8.9.

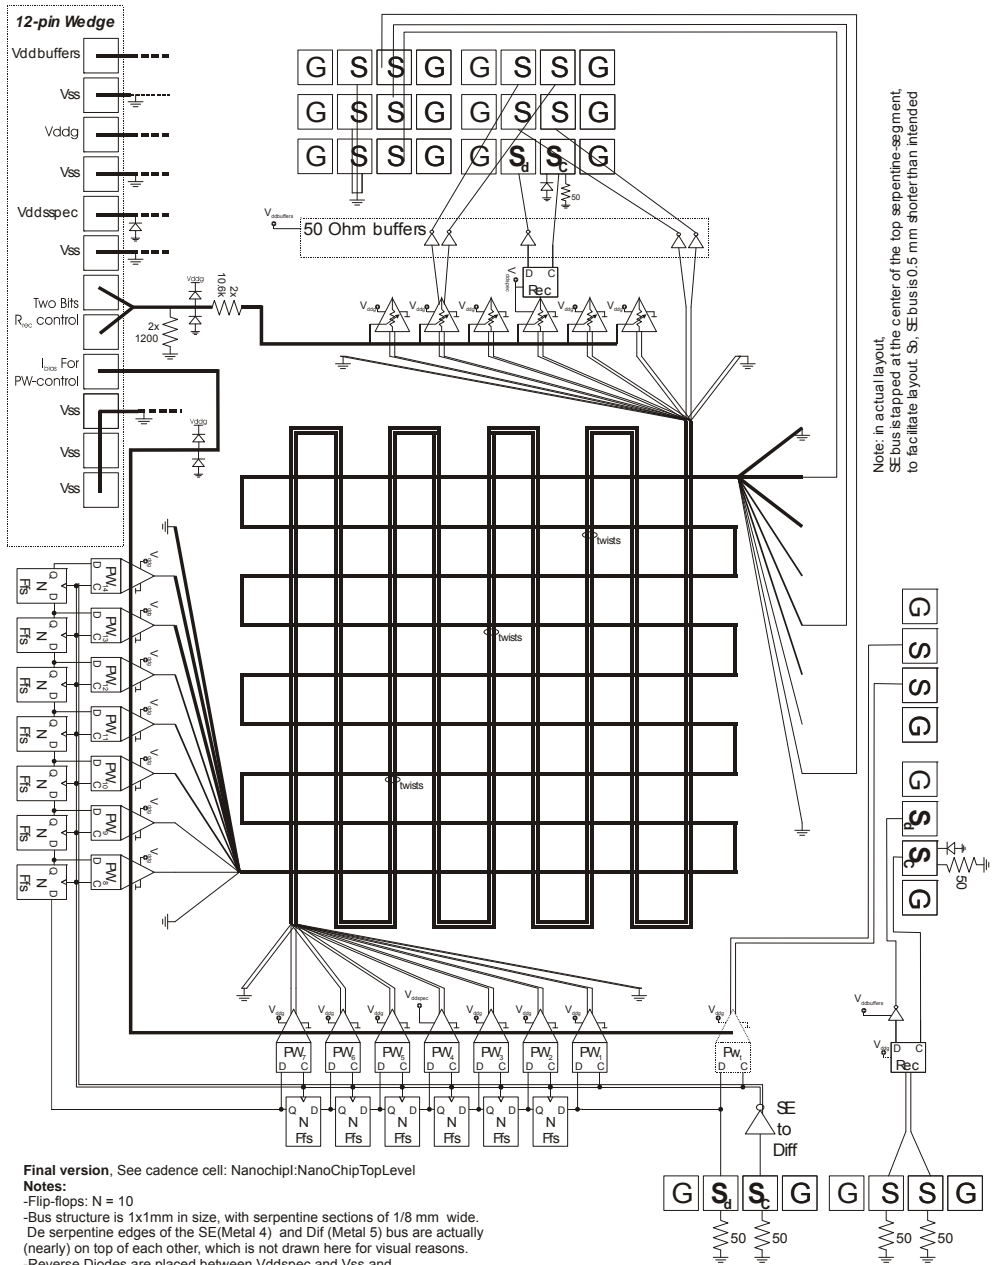


Figure 8.8: Floorplan of the first demonstrator IC.




Figure 8.9: First chip micrograph.

A 7 channel differential bus with twisted wires (width and spacing of $0.4\mu\text{m}$ each, optimized as explained in section 2.7.1) is placed in metal 5 and is completely surrounded by grounded and VDD-connected metal stripes. An additional 7 channel single-ended bus with perpendicular orientation is placed below the differential bus for additional characterization purposes (a variety of wire pitches is used in this bus) and to provide an indication of inter-layer crosstalk. An external single-channel 3.2Gb/s pattern generator/analyzer (Anritsu MP1632C) is used for the data generation and BER measurement. Large on-chip delay lines (chains of flip-flops, 10 per channel) provide all bus channels with pseudo-independent data.

Different twisting patterns and receiver configurations are used for the different channels of the differential bus, as shown in Figure 8.10. Channels 1, 4 and 6 are equipped with 50Ω output buffers and pads for measurements. The output buffers consist of properly sized

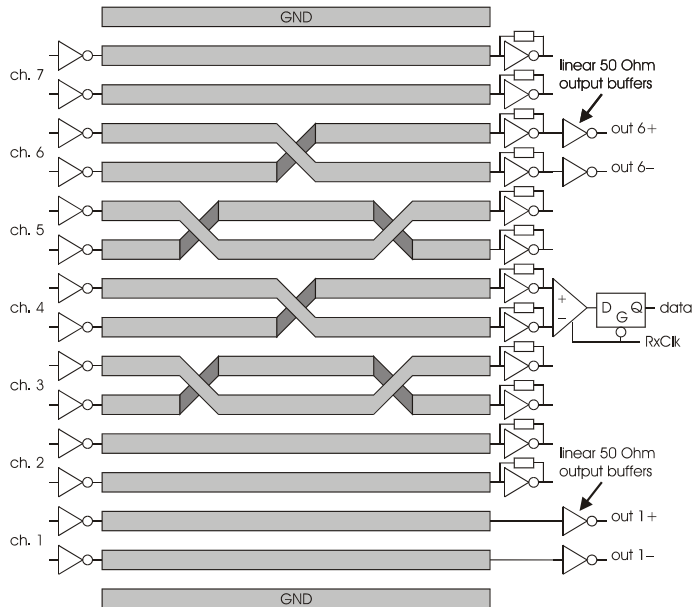


Figure 8.10: Differential bus configuration as implemented on first chip.

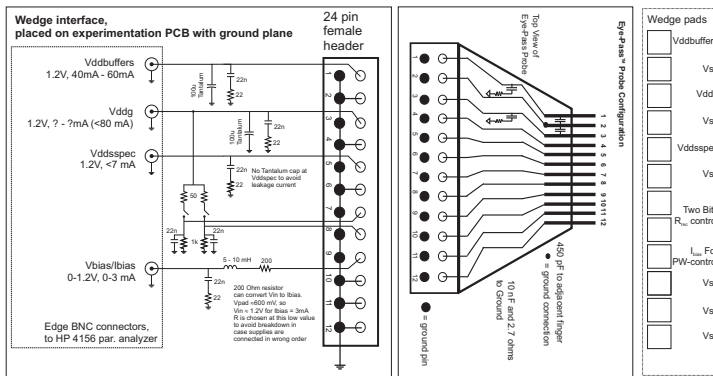


Figure 8.11: Wedge interface for first chip.

inverters with an on-chip 50Ω pull-up resistor at the output. When they are loaded by a probe with 50Ω resistance to ground, then they can accommodate a full-swing input range, with some large-signal compression. The output buffers attenuate the signal by about 6 dB (small-signal) to 9 dB (large-signal). Channel 4 is used for the BER measurements and its transceiver circuits also have their own dedicated supply to enable power measurements. The other two differential channels are used for e.g. crosstalk and eye-diagram measurements. The receiving ends of the single-ended bus interconnects are directly connected to pads to enable measurements directly on the interconnect.

8.7 Measurement setup

The chip has been measured in a probe station using 50 Ω GSSG probes for the high-speed signals a 12-pins probe (wedge) for the supplies and control signals. The GSSG probes directly connect to the measurement equipment, with the transmitter probes connected to the data pattern generator of the Anritsu MP1632C and the receiver probes either connected to the pattern analyzer of the Anritsu (for BER measurements) or to an Agilent 86100A sampling scope (for eye diagrams and other waveform measurements). The wedge has a small PCB interface, shown in Figure 8.11 and then connects to a HP4156A parameter analyzer to supply and measure the low-frequency signals. At the receiver side, each channel has its own dedicated GSSG pads to enable wide-band measurements directly on the channel. Some photo's of the measurement setup are shown in Figure 8.12 on the next page.

The setup with the delay lines between the channels not only allows for random-data BER testing in a realistic crosstalk environment, but deterministic data patterns can also be applied. The deterministic data-patterns were mostly applied for measurements with the sampling scope on channel 1, 6 and the single-ended channels. These included step response measurements (with zero I_{bias} for the transmitter to create regular binary signals) and also transfer function measurements with high-frequency square-wave patterns. The square-wave patterns were made with all-zero or all-one data patterns and a high I_{bias} for the pulse-width transmitter to create pulses with (nearly) 50% pulse-width. The transfer function was obtained by repeating the high-frequency square wave patterns at different frequencies (different clock-rates) and extracting the phase and amplitude of the base harmonic at the receiver (with suitable correlation and filtering to minimize the influence of noise sources such as the limited resolution of the sampling scope).

The phase of the high-frequency square wave can also be changed by 180 degrees by switching from an all-zero to an all-one data pattern (again with a high I_{bias} to let the pulse-width transmitter creates the square wave from this data). This enabled crosstalk transfer measurements, as the 10 cycles delay between each channel enabled transfer measurements with a neighboring channel having both the same phase as well as 180 degrees different phase. Examining the received magnitude and phase of the base harmonic both before and after a neighboring channel makes the phase transition enables extraction of the amount of crosstalk (see [82], fig. 15 for a signal example).

For the differential interconnects from channel 1 and 6, both single-ended halves were measured simultaneously with the sampling scope. Transfer and crosstalk measurements on these channels were used to verify the validity of the twisted differential wires analysis [81, 82].




Figure 8.12: Measurement setup photos.

8.8 Experimental results

8.8.1 Parameter characterizations

Before discussing the achievable data rate results, a few of the more interesting characterization measurements are shortly discussed in the list below:

- **On-chip voltage levels and probe contact resistances:** The on-chip supply voltage levels were measured to investigate the influence of probe (wedge) contact resistance. For V_{ddbuf} this was done by using an all-one data pattern with $I_{\text{bias}}=0$ and measuring on channel 1 with a high-impedance voltage meter: the all-one pattern pulls the output of channel 1 to V_{ddbuf} as there is no pull-down resistance (because of the high-impedance meter). For V_{ddg} the same measurement was done on a single-ended channel. Repeated measurements showed that the probe contact resistance varied between 5 and 10 Ω . This meant that on-chip supply levels could be as much as 0.5V lower than intended as the current drawn from V_{ddg} ranged from 0 to 50mA (depending on clock frequency and settings). For the BER measurements on channel 4, this was a lesser problem as that channel had an independent supply (with a much lower supply current of $I_{\text{ddspec}}=0\text{--}6\text{mA}$). To improve consistency between different measurements, probes were repositioned when the supply currents deviated substantially (it was unfortunately not possible to continuously measure on-chip voltage levels, as probe arms had to be moved to do this). Note that the voltage drop over the ground connection was assumed to be insignificant, because of the many ground probe contacts.
- **Wire resistance:** The resistances of the interconnects were measured on the single-ended channels, by connecting a low-impedance current meter to pull their outputs to ground and measure the current I_{wire} (with $I_{\text{bias}}=0$ and with the transmitter sending all ones). The resulting resistance ($V_{\text{ddg}}/I_{\text{wire}}$) still included the probe resistances and the output resistance of the transmitter. The former was partly removed from the equation by also measuring the on-chip V_{ddg} . The latter was estimated from simulations to be about 65 Ω . The remaining probe contact resistance from the current meter was assumed to be 10 Ω . This led to the derivation of the following wire resistances:
 - 1.6 μm wide interconnect in M4: $R'=44\Omega/\text{mm}$
 - 0.4 μm wide interconnect in M4: $R'=190\Omega/\text{mm}$

When a 16nm barrier thickness is assumed to slightly reduce the actual conductor width (which matches the value from [27]), then these measurements predict a sheet resistance of $R_{\text{sheet}} = 69 \text{ m}\Omega/\square$ which is slightly above the average value specified in the technology manual and well within the tolerance bounds. That the wire resistance of 190 Ω/mm is quite a bit higher than the 150 Ω/mm that was used for simulations during the design of the circuits, is because, at design time, the barrier was ignored and the nominal sheet-resistance was used.

Note that the resistance of the M5 differential interconnects could not be measured directly. It was therefore assumed that the resistance of the 0.4 μm wide M5 wires was the same as that of the 0.4 μm wide M4 wires.

- **Wire capacitance:** The wire capacitance was measured indirectly by fitting measured eye diagrams and achievable data rate figures onto their simulated counterparts (assuming 65Ω driver and 1900Ω wire resistance). This was done with the transceiver configured as conventional binary transceiver, to minimize the effect of secondary high-frequency poles. This resulted in the following estimated wire capacitance:
 - Differential interconnect in M5, $0.4\mu\text{m}$ wide wires: $C'=250\text{fF}/\text{mm} \pm 10\%$

This value is slightly lower than the $270\text{fF}/\text{mm}$ that was obtained with the 3D EM-field solver simulations, but the difference is within the tolerance bounds. The actual bus does have less dense surrounding metal than in simulations, because the metal plains from the EM model were replaced by metal bars with a fill density of 59% (both in metal 4 and metal 6). Rough simulations indicated that a filling of less than 100% had only minor influence on the capacitance, so in that respect, the results match quite well.

The EM-field solver simulations and the crosstalk measurements suggest that the capacitance is composed of roughly $0.05\text{ pF}/\text{mm}$ to each of the four sides; with the part of the capacitance between the differential wires doubled because of the miller multiplication.
- **Receiver input resistance:** comparison of simulations and measurements with the receiver configured with a low input impedance, confirmed that the effective receiver impedance is about 150Ω , as it was designed.

8.8.2 Signal measurements

The configurability of the transmitter and receiver enables measurements with or without PW pre-emphasis and with conventional or resistive receiver termination. Eye-diagrams for each of the four settings are shown in Figure 8.13, as measured at the output of channel 6. BER measurements (with PRBS data patterns) for the four settings were carried out at channel 4 and Table 8.2 shows the highest data rate at which bit-errors are not yet measurable ($\text{BER} < 1\text{e-}12$). At the boundary of error-free operation the BER drops sharply, as the primary bit-error sources are deterministic (ISI) or static (offset) and a BER much lower than $1\text{e-}12$ is expected at the shown data rates.

The transceiver circuits of channel 4 have a dedicated supply, and the energy consumption of the channel is also shown in Table 8.2 (measured with PRBS data patterns, giving 50% data activity). Simulated values for the energy consumption of the various parts of the transceiver are also shown. The Tx and Rx circuits consume more power than necessary for a given mode of operation, as they are designed to function in all modes and are optimized for speed.

The results show good agreement with the analysis. The 550 Mb/s achieved in the conventional case is only slightly lower than the theoretical limit of 600 Mb/s from section 8.3. Resistive termination improves the achievable data rate by nearly a factor three. The improvement of PW pre-emphasis together with conventional termination is a factor of four and is a factor of two if used in combination with resistive termination. The difference in data rate between PW pre-emphasis with conventional or with resistive termination is a factor of 1.5, close to the factor 1.4 that was predicted in section 7.4.2.




Figure 8.13: Eye-diagrams measured on the first chip with various transceiver settings. The output buffers compress the vertical scale by 6 to 9 dB First chip micrograph.

	Conventional Rx termination	Resistive Rx termination
Without PW pre-emphasis	550 Mb/s 3.4 pJ/bit (Tx: 0.2pJ/b, Wire: 1.2pJ/b, TIA: 0.7pJ/b, SA: 1.4pJ/b)	1.5 Gb/s 2.5 pJ/bit (Tx: 0.2pJ/b, Wire: 0.8pJ/b, TIA: 0.9pJ/b, SA: 0.5pJ/b)
With PW pre-emphasis	2 Gb/s 2.5 pJ/bit (Tx: 0.5pJ/b, Wire: 0.8pJ/b, TIA: 0.7pJ/b, SA: 0.5pJ/b)	3 Gb/s 2 pJ/bit (Tx: 0.5pJ/b, Wire: 0.6pJ/b, TIA: 0.5pJ/b, SA: 0.3pJ/b)

Table 8.2: Achievable data rate ($BER < 1e-12$) and energy consumption. Between brackets the simulated energy consumption values are shown for the transmitter (Tx), the wire, the transimpedance amplifier (TIA) and the sense amplifier with latch (SA)

The eye at the receiver is still open at 3.2Gb/s as visible in the bottom right of Figure 8.13, but the opening is so small (40 mVpp) that effects such as hysteresis and offset in the clocked receiver prevent reliable detection at this data rate ($BER = 10^{-8}$). At 3Gb/s, error-free operation is possible for all 10 measured samples ($I_{bias} = 400\mu A$) but without much Vdd or I_{bias} tolerance.

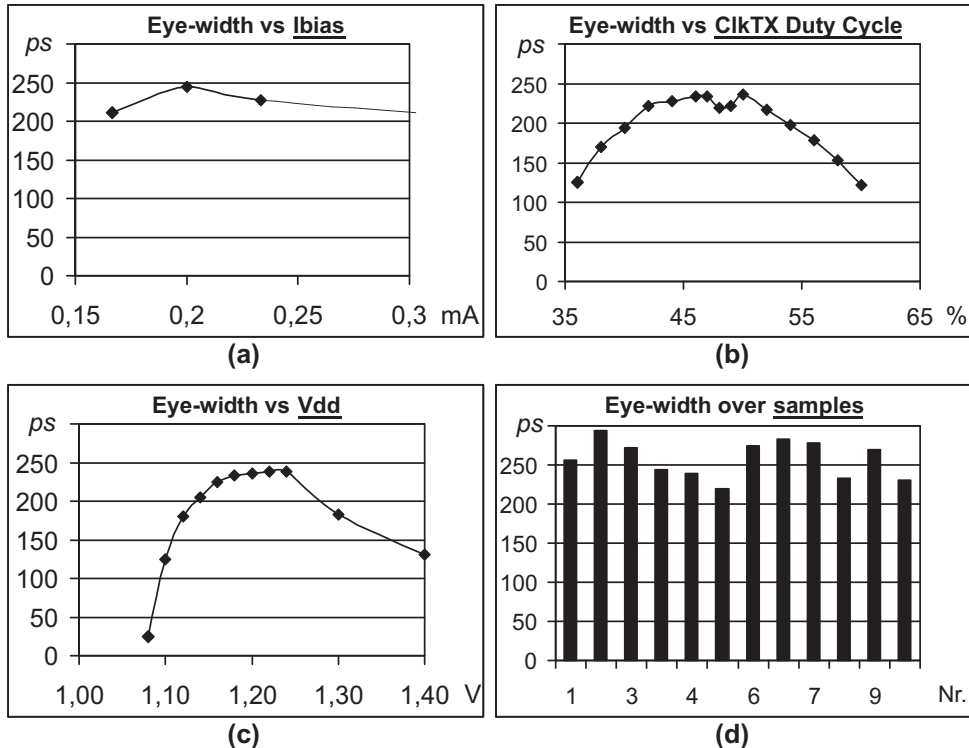


Figure 8.14: Measured eye-width versus parameters and over different samples at 2.5Gb/s.

At 2.5Gb/s, the design is robust and the BER remains immeasurable with large external parameter deviations. Figure 8.14 illustrates this robustness by plotting the measured eye-width as a function of an external parameter while keeping the other parameters at their nominal value ($V_{dd} = 1.2$, Clk_{TX} duty cycle = 50% and $I_{bias} = 200\mu A$). To measure the eye-width, a phase-shifter was used to vary the skew of the receiver clock and find the phase-shifts where the BER just becomes measurable. The optimal bias current of $200\mu A$ (giving a PW pre-emphasis duty cycle of about 58%) agrees with predictions from Figure 8.5. The measured relationship between the external Clk_{TX} duty cycle and the eye-width also behaves as expected, except for a small drop in eye-width around 50% Clk_{TX} duty cycle which can probably be attributed to measurement tolerances. The highest measured eye-width of 250ps is lower than the theoretical value of almost 400ps due to the required setup and hold times of the sense amplifier.

The influence of crosstalk on the eye-diagram is shown in Figure 8.15. This figure shows both the output of the single-ended (SE) halves and the differential output of channel 6 at a rate of 2.5 Gb/s. Each SE half of channel 6 receives crosstalk mainly from the wire-piece that runs alongside channel 7 (but the other channels in the bus and the perpendicular bus also generate some common-mode crosstalk). The eye-closure due to the crosstalk in the single-ended output is clearly visible in the figure, while the crosstalk is mitigated in the differential output. If the twist in channel 6 would not be present, then the crosstalk on both SE halves would be even higher and it would not be canceled in the differential output.

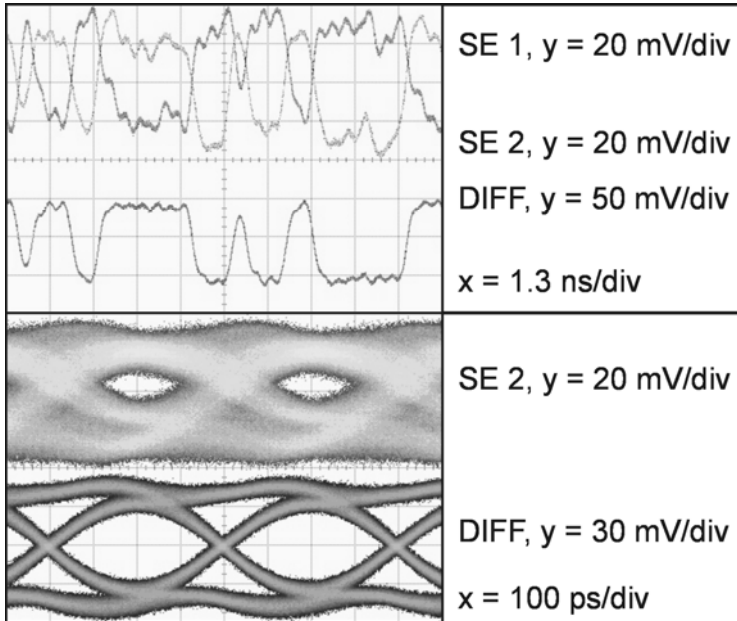


Figure 8.15: Effect of crosstalk on the single-ended and twisted differential output of channel 6 at 2.5 Gb/s. The output buffers compress the vertical scale by 6 dB.

8.9 Conclusions from first demonstrator IC

This first demonstrator IC showed the effectiveness of PW pre-emphasis, resistive receiver termination and twisted differential interconnect techniques to improve global on-chip data communication for given lengths of uninterrupted interconnect. Measurements showed that the techniques can increase the data rate from a mere 0.55Gb/s/ch with the transceiver operating in conventional mode, to as much as 3Gb/s/ch (2pJ/bit). At 2.5Gb/s/ch, the system is tolerant to parameter deviations. Analysis, such as the predicted tolerance in pulse-width, agree well with measurements.

When compared to classical repeater techniques targeted at comparable data rates, the presented transceiver can bridge an uninterrupted wire length that is a factor of ten higher and has a lower and more predictable latency (650ps versus 965ps). Pulse-width pre-emphasis and resistive termination should enable much higher data rates for shorter lengths of uninterrupted interconnect (or for interconnects placed in higher/larger metal layers), as long as the wire bandwidth is the bottleneck. The presented techniques can thus be used for repeaterless global communication or can be used to improve the trade-off between data rate and repeater spacing.

The power consumption (6mW at 3Gb/s) of the transceiver in its form as implemented on the demonstrator IC has little dependence on data activity. However, the transmitter and receiver circuits are well suited for power management as the speed-enhancing, but power consuming techniques can be easily turned on and off dynamically (for example as done in [66, 70]). Also, other improvements, such as a more energy efficient sense amplifier as presented in the next chapter, could further reduce the power consumption. For this first

demonstrator IC, the main criterion was to improve data rate, which is where especially the PW pre-emphasis excels, at the cost of power at low data activities. In those applications where power is more of a concern than data rate, the capacitive pre-emphasis transmitter that is presented in Chapter 10 is a better alternative. A capacitive pre-emphasis transmitter does lower the swing, which means that some concessions have to be done on data rate for error-free detection in the presence of receiver offset.

Chapter 9

Improved sense amplifier

9.1 Introduction

The second demonstrator IC focused on more power efficient transceiver techniques. This included a more power efficient sense amplifier with build-in DFE equalization. As a sense amplifier, or clocked comparator, is such a general circuit, a stand-alone version without the equalization was also included on the second chip, which is discussed here based on the publications [86] and in [157].

The sense amplifier that is discussed here is a variant of a ‘latch-type sense amplifier’. Latch-type sense amplifiers – also called ‘sense amplifier based flip-flops’ when the latch is followed by a second stage – , are very effective comparators. They achieve fast decisions due to a strong positive feedback and their differential input enables a low offset. Sense amplifiers (SA) are hence widely applied in e.g. logic, memories, I/O data receivers [158], A/D converters [159-162], and more recently also in on-chip transceivers [20, 33, 61, 67, 69, 72, 77, 84, 85, 121, 124].

Especially voltage-mode SA’s, as shown in Figure 9.1 on the next page, have become quite popular [158, 163-167]. This circuit is also often referred to as a ‘StrongARM’ latch and a few variants are patented by Digital Equipment Corporation (DEC) [168]. DEC was the original manufacturer of the StrongARM family of ARM-based microprocessors. The popularity of this circuit as a sense amplifier is due to its high input impedance, full swing output and absence of static power consumption. However, the stack of transistors in a conventional voltage-mode SA requires quite a large voltage headroom, which is problematic in low-voltage deep-submicron CMOS technologies. Furthermore, the speed and offset of such a circuit are very dependent on the common-mode voltage of the input V_{cm} [165], which is a problem in applications with wide common-mode ranges, for example A/D converters.

To circumvent these drawbacks, a latch-type voltage sense amplifier with a separated input and cross-coupled stage was developed in this project, the ‘double-tail’ sense amplifier [86]. There already exist many types of sense amplifiers with a separated differential input stage [169, 170], but those are circuits that do consume static power (such as current-mode logic latches). The double-tail sense amplifier is a fully dynamic circuit without any static power consumption. It can operate at a lower supply and has a more stable offset than its




Figure 9.1: Conventional latch-type voltage sense amplifier. The dotted transistors are examples of common variations.

conventional counterpart, with a significantly better offset per power ratio at high common-mode voltages. This makes it a very suitable sense-amplifier for offset-critical applications, especially when they operate with a common-mode close to the supply, as is the case for the transceiver with the capacitive transmitter that is discussed in the next chapter.

This chapter discusses the double-tail sense amplifier in more detail. The next section starts with a short discussion of the conventional sense amplifier and its drawbacks. Then, in section 9.3 the operation of the double-tail sense amplifier and what its advantages are is discussed, with a noise analysis in section 9.4. The conventional and double tail sense amplifier are compared in section 9.5. Section 9.6 discusses the measurements and section 9.7 presents some conclusions on the sense-amplifier. In Appendix A, some additional material is presented on how offset (and noise) in a sense amplifier can be analyzed and simulated or measured.

9.2 Conventional sense amplifier and its drawbacks

An extensive (numerical) analysis of the operation of the conventional sense amplifier from Figure 9.1 is given in [165]. Here we will limit ourselves with a short description of its operation (also see Figure 9.2b for signal-graphs of a functionally similar circuit).

The circuit in Figure 9.1 operates similar to other dynamic circuits with a reset or pre-charge phase and an evaluation phase. During the first part of the clock cycle, when the *Clk* is low (0V), the output nodes of the cross-coupled inverters (M1-M4) are reset to V_{dd} , using the reset transistors M7 and M8. The second part of the clock cycle is the actual sense and evaluation phase. When the *Clk* signal starts to rise, the tail (M9) of the differential pair (M5, M6) is turned on. The differential pair will discharge the *Di* and later the output (*Out*) nodes and an input-dependent voltage difference will build-up on these nodes. When the *Di* nodes have dropped about a threshold voltage (V_t) under V_{dd} , then the NMOS transistors of the cross-coupled inverters (M1, M3) turn on, marking the start of the positive feedback.

When the D_i nodes are about $2V_t$ lower than the supply, the PMOS transistors of the inverters (M2, M4) also turn on; further enhancing the positive feedback and enabling the regeneration of a small differential voltage at V_{in} to a full swing differential output.

The circuit has a large number of variations. First of all, it is of course possible to make a complementary version of this circuit, with all the NMOS and PMOS transistors swapped [171]. Second, an often found addition is a transistor between D_{i+} and D_{i-} , as shown with the dotted transistor in Figure 9.1. This transistor prevents that the output of the SA becomes floating when the input signal changes polarity after a decision has already been made [164]. The other two dotted transistors in Figure 9.1 are additional reset transistors that reset the D_i nodes [85, 158, 166, 172]. They improve operation at high common-mode input voltages (as discussed later) and also significantly reduce the hysteresis (or memory effect) of the comparator. Other variations include the use of special clocks, for example with non 50% duty cycles, or with slightly different timings for the tail and reset transistors [171] to optimize the timing of the various phases in the operation cycle.

The one thing that many variants of the dynamic sense amplifiers have in common is that the cross-coupled (latching) inverters are placed in series with the differential pair. This series (or cascode) configuration has several drawbacks in applications with limited voltage headroom. The first drawback is the fact that there is only a very short time in which the differential pair actually has gain. This is especially a problem when the input has a common-mode voltage V_{cm} close to the V_{dd} (which is often the case for example with memories, and also in transceivers for low-swing data communication, as in [62]). In that case, the differential pair will enter triode region when the D_i node voltages drop below $V_{dd}-V_t$ and the short period between the end of the reset-phase and this moment is the (sampling or sensing) interval in which the input is amplified and integrated onto the capacitances at the D_i nodes. Low amplification (due to a short integration time) of the input signal means a high sensitivity to offset from stages further in the signal chain, in this case offset originating from M1 and M3. In the conventional circuit the D_i nodes are not reset to V_{dd} , but to about one V_t below V_{dd} (through transistors M1, M3), which further reduces the integration time.

The integration time can be lengthened by also resetting D_i to V_{dd} with the additional reset transistors, as shown with the two dotted transistors in Figure 9.1. Simulations with the circuit in a $0.13\mu\text{m}$ CMOS process showed that the additional reset transistors can reduce the input-equivalent offset of M1, M3 by a factor of three (given $V_{in-com.mode} = 1.1\text{V}$ and $V_{dd} = 1.2\text{V}$).

However, the additional reset transistors solve only one of the drawbacks. The remaining drawback is the fact that there is only one current path, via tail transistor M9, which defines the current for both the differential amplifier and the latch (the cross-coupled inverters). On the one hand, one would like a small tail current to keep the differential pair in weak inversion and obtain a long integration interval and a better Gm/I ratio. On the other hand, a large tail current is desirable to enable fast discharge and regeneration in the latch. For regeneration, it is also not favorable that this tail current depends on the common-mode voltage of the input, which it does in this circuit (as M9 operates mostly in triode).

A solution to circumvent these drawbacks is to decouple the available current for the latch from the available current for the differential pair. This is accomplished with the double-tail circuit as discussed next.

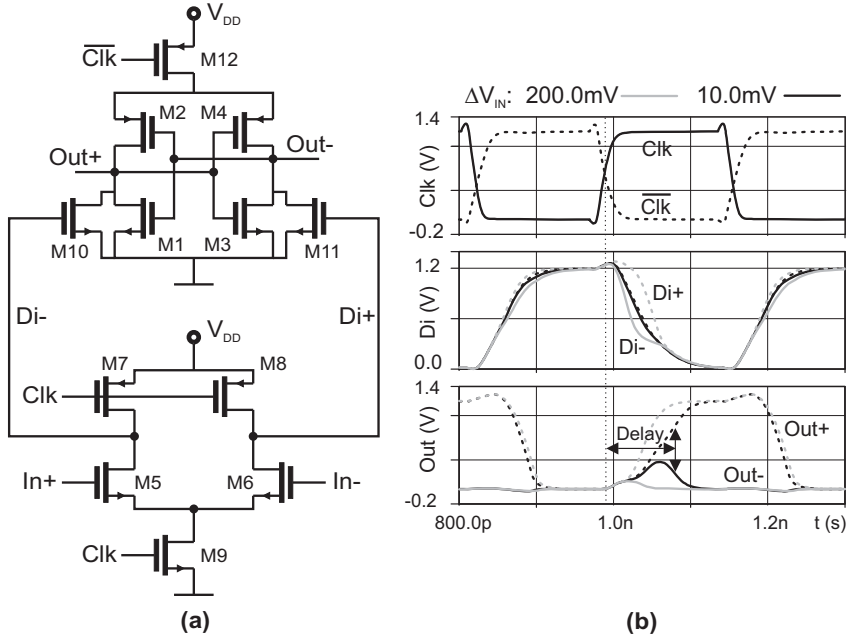


Figure 9.2: Double-tail latch-type voltage sense amplifier (a) and signal behavior (b).

9.3 Double-tail sense amplifier

The schematic of the double-tail sense amplifier is shown in Figure 9.2a. This topology has less stacking and can therefore operate at lower supply voltages. The double tail enables both a large current in the latching stage (wide M12), for fast latching independent of the V_{cm} , and a small current in the input stage (small M9), for low offset.

The signal behavior of the double-tail SA is shown in Figure 9.2b. During the reset phase ($Clk=0V$), transistors M7 and M8 pre-charge the Di nodes to V_{DD} , which in turn causes M10 and M11 to discharge the output nodes to ground (so there is no need for dedicated reset transistors at the output nodes). After the reset phase the tail transistors M9 and M12 turn on ($Clk=V_{DD}$). At the Di nodes, the common-mode voltage then drops monotonically with a rate defined by I_{M9}/C_{Di} and on top of this, an input dependent differential voltage ΔV_{Di} will build up. The intermediate stage formed by M10 and M11 passes ΔV_{Di} to the cross-coupled inverters and also provides additional shielding between in- and output, with less kickback noise [169] as a result. The cross-coupled inverters start to regenerate the voltage difference as soon as the common-mode voltage at the Di nodes is no longer high enough for M10 and M11 to clamp the outputs to ground. The timing of the various phases can be tuned with the transistor sizes. Tuning can also depend on the nominal operating point (V_{cm}).

Compared to the conventional sense amplifier, this circuit requires a few additional transistors, but the total area can be comparable, as will be shown in the next section. It also requires the availability of both a clock and a clock-not signal. Often, both a clock and a clock-not are already available in a system. If not, then a simple inverter can generate the




Figure 9.3: Linear time-variant model of a double-tail sense amplifier.

clock-not from the clock, as the clock-not is allowed to trail the clock signal without a significant impact on performance.

9.4 Sense amplifier speed, offset and noise analysis

To be able to optimize the design of the sense-amplifier, linear time-variant (LTV) models were developed. For the double-tail sense amplifier, such a model is shown in Figure 9.3. The signals in the model represent the differential signals in the actual circuit and the time-variance is controlled by the common-mode signals. The input stage acts as an integrator that is reset when the input transistors enter deep-triode. Gm_2 models the transconductance of the intermediate stage (M10 and M11) and the last four blocks represent the actual latch (the cross-coupled inverters). As mentioned above, the latch becomes active when the intermediate stage is no longer able to clamp the outputs to ground. The time constant of the positive feedback of the latch is the $\tau_{\text{latch}} = Gm_3/C_2$.

Although the stages of the actual circuit do not become active or inactive instantaneously, the model is still able to predict the behavior of the actual sense amplifier circuits quite accurately.

The delay of the sense amplifier for example, basically consists of two parts, similar as described in [165]. First, there is the fixed delay for the sampling part in which the differential pair integrates the input onto the Di nodes, without the latch being active. The second part of the delay is from the latch and this part is logarithmically dependent on the sampled input voltage, as the positive feedback creates an exponentially increasing signal.

Due to this exponential increase, it is not necessary to keep the input stage active for more than about three to four times τ_{latch} after the latch is turned on, as the contribution of the input of the latch quickly becomes insignificant compared to the internally build-up signal.

To maximize the gain in the input-stage of the sense amplifier, and hence minimize offset and noise contributions from later stages, it would be ideal to first turn on only the input integrator and leave the latch inactive until the integrator has amplified the input with a suitable factor. However, that would come at a cost of an increased delay. In this double-tail circuit it is also not a possible approach, as the timing of the different phases are all linked to the common-mode behavior of the Di nodes.

When noise is of a bigger concern than speed, then other topologies might be more suitable. These include for example the sense amplifier presented in [160, 173] and the circuit from [174]. Both these sense amplifiers also have a split between the input and the latching stage.

In the double-tail circuit, a way to maximize the gain of the input integrator is to keep the input differential pair (M5, M6) operating in (or at the edge of) weak inversion. In weak inversion the G_m/I_d is highest. This is important because the effective amplification factor equals the integrator gain (G_{m1}/C_1) times the integration time (proportional to C_1/I_d).

A more detailed numerical analysis that also takes the timing of the latch into account can be carried out with the help of the LTV models. These models can be used both for analysis of the amount of input-equivalent offset as for the amount of input-equivalent sampling noise.

For calculations of the noise in comparators and comparator-based circuits also see for example [160, 167, 175]. In [160] a concise discussion of the noise in the input-stage of the comparator is given and its contribution in a charge-redistribution ADC. In [175] a broad discussion is given on noise in comparator-based circuits, including a background of non-stationary noise sources. In [167] an excellent, detailed treatment can be found on the noise sources in the various phases of a clocked comparator itself, also based on time-domain analysis with linear time-variant models for the system.

9.4.1 Double-tail sense amplifier dimensioning for low offset

Although the rms value for the noise is high in sense amplifiers due to the very wide sampling bandwidths, it is still significantly lower than the offset (as will be confirmed in the measurements section), unless offset calibration techniques are used [158, 166, 174]. Even though offset calibration is proposed for use on-chip communication by Jose et. al. [20, 61], the overhead in complexity that calibration costs does not (yet) seem to outweigh its benefits for wide on-chip busses.

In this project we therefore focused on offset minimization and we used the LTV models to determine how the transistor parameters should be tuned to get the lowest offset at a certain total area, while maintaining high speed. The dominant cause for offset in this circuit is V_t mismatch, which is inversely related to the square root of the area of a transistor, as are most types of offset sources [176]. The total area is distributed over the transistors in such a way that the (area) derivative of the contribution to the input-equivalent offset are equal for all transistors. The result is a minimum offset for a given area .

The transistors that contribute to the offset are in order of importance: the input transistors (M5, M6), the PMOSTs from the latch (M2, M4) and the intermediate stage (M10, M11). The NMOSTs from the latch have only a very minor contribution to the offset, as the signal is already strongly amplified when these transistors become active. These NMOSTs are still important for the speed of the latter part of the regeneration phase and are hence optimized for this criterion. The reset transistors (M7, M8) also have a very low contribution to offset. They do however have an impact on the amount of hysteresis. Their effective impedance determines the amount of signal residue at the D_i nodes at the end of the reset phase. We dimensioned these reset transistors such that the input equivalent hysteresis is significantly lower than the offset, specified at the maximum clock frequency (lower than 0.5mV with a 3GHz clock).




Figure 9.4: Simulated delay and power as a function of the supply voltage ($\Delta V_{in}=50mV$, $V_{cm}=V_{DD}-0.1V$).

As all the nodes in the circuit are dynamic – with their capacitances charged and discharged in every clock-cycle – area translates directly to power and an optimal offset/area is roughly equivalent to an optimal offset/power.

Area-scaling (impedance scaling) for the complete design can subsequently be used to match the total input-equivalent offset to any desired value. Note that with this procedure, the total input-equivalent noise will scale with the same factor as the offset and their relative importance will not change.

9.5 Comparison of double-tail with conventional

To compare the conventional and the double-tail sense amplifier, both circuits were simulated in a 90nm CMOS technology with $V_{DD}=1.2V$. Both circuits were optimized with the help of the linear time-variant models and the transistor dimensions were scaled to get an equal offset standard deviation of $\sigma_{os}=10mV$ at the nominal input common-mode voltage of $V_{cm}=1.1V$ (the same conditions that are found in the transceiver on the second demonstrator IC [62]). At this high V_{cm} , the additional reset transistors at the Di nodes are a must in the conventional topology, to avoid unrealistically high offsets.

At the nominal conditions, equally high performance figures are obtained for both the double-tail as the conventional variant, with only 100ps Clk-to-output delay (including a clock-buffer) and with only 90fJ/bit consumption for 10mV offset). Note that this 100ps includes about 50ps of clock-to-sample delay (negative setup time of $\sim 50ps$). Note, that this comparable behavior of the two circuits is also due to the optimization and dimensioning for equal offset. Other dimensioning criteria can give different results [177] (but the exact dimensioning criteria used in [177] are not really clear to the author).

When the operating conditions are changed, then the two circuits start to behave differently. Figure 9.4 shows the simulated performance – in terms of (clk to output) delay and energy/cycle – as a function of the supply voltage (with V_{cm} being 0.1V lower than the supply). It is clear that the double-tail topology is faster and can operate at lower supply voltages, while it consumes approximately the same power as the conventional topology. The double-tail topology could for example operate at a supply of 0.5V at a cost of only




Figure 9.5: Simulated delay and power as a function of the common-mode voltage of the input ($\Delta V_{in}=50\text{mV}$, $V_{dd}=1.2\text{V}$).

10fJ/cycle with 1000ps delay, versus 2350ps for the conventional circuit. Note that this is for a design that is optimized to operate with $V_{DD}=1.2\text{V}$. Optimization for $V_{DD}=0.5\text{V}$ would give smaller delay, as a wider tail would be used for the input section.

Figure 9.5 shows the simulated performance as a function of the V_{cm} . Again, the double-tail topology is faster and has a wider common-mode range. The power consumption is nearly equal, except at low input common-mode voltage, where the double-tail topology is able to make faster decisions at the cost of power.

An interesting difference at high V_{cm} is the difference in offset standard deviation. At $V_{cm}=1.4\text{V}$, the offset for the conventional topology increases to $\sigma_{os}=30\text{mV}$, while the double tail offset becomes only $\sigma_{os}=15\text{mV}$, a factor two difference. At common-mode levels lower than the nominal value of 1.1V , the offset of both types of converters remain roughly equal (10% lower for the conventional circuit at $V_{cm}=0.5\text{V}$).

To obtain the above-mentioned figures, the standard deviation for the offset was estimated with monte-carlo simulations with 1000 trials. The differential input voltage ΔV_{in} of the sense-amplifier was set at a value around the expected standard deviation and the ratio of the trials with the correct positive decision (p) was subsequently used to calculate the actual offset standard deviation. The effect of hysteresis was excluded from this analysis by first using a large negative V_{in} during one clock-cycle, followed by the actual decision test. In that way all sense amplifiers start in the same negative state. The details of the sigma estimation procedure are discussed in Appendix A.




Figure 9.6: Chip micrograph with enlarged layout of sense amplifier.

9.6 Sense amplifier measurements

The double-tail sense amplifier was implemented in a 1.2V 90nm CMOS technology on the second demonstrator IC, as part of the low-swing on-chip data transceiver with capacitive pre-emphasis, which operates around a V_{cm} of 1.1V [62], as discussed in the next chapter. The V_{cm} can have large variations due to e.g. crosstalk effects. See section 10.4 for more details on this transceiver demonstrator IC and the measurement setup.

A double-tail SA with dedicated input and output pads (for probe station measurement) was placed on the same die. The layout of the double-tail SA is shown in the inset of the chip micrograph in Figure 9.6. An SR-latch (made from two NOR gates) is connected to the output of the SA to create static output signals without loss of timing information from the core of the SA. When required, more advanced ‘slave’ stages could be used [164, 178]. A simple SR-latch is for example not ideal when the sense amplifier is used at very high speeds, as it has a ‘non-overlapping’ behavior—the falling edge always comes first—which creates a significant state dependent delay. An SR-latch furthermore also has a state-dependent input capacitance which increased the hysteresis of the total sense amplifier to about 1.5mV (simulated). But, for application in the low-swing data transceiver, the SR-latch sufficed.




Figure 9.7: Measured delay as a function of the differential input voltage (a) and the common-mode input voltage (b), including a comparison with simulations.

Figure 9.7 shows the measured relative delay under different conditions (the absolute delay is not measurable due to additional delay from the output drivers). As intended, the minimal delay is found at $V_{cm}=1.1V$. At a V_{cm} of $0.6V$, there is still only 20ps increase in delay. The delay versus ΔV_{in} is 44ps/decade under nominal conditions. In comparison, measurements in [165] on a conventional topology in CMOS $0.13\mu m$ with $V_{DD}=1.5V$ show a delay versus ΔV_{in} of 100 to 170ps/dec and a 250ps increase in delay when V_{cm} is lowered to $0.6V$.

The offset in [165] is also very dependent on the V_{cm} and rises from 8.5mV to 19mV when the V_{cm} changes from 1.05V to 1.5V. For our design, measurements on 20 samples gave an offset of $\sigma_{os}=8mV$, at a V_{cm} of both 1.1V and 0.75V. If desired, area upscaling could further reduce the offset at the expense of power ($P \propto 1/\sigma_{os}^2$). As mentioned earlier, offset compensation schemes [158, 166, 174] are a good alternative if the application allows for the added complexity.

The power consumed by the SA is 113fJ/decision when ΔV_{in} has 50mV amplitude ($f_{clk}=1GHz, V_{DD}=1.2V, P=113\mu W @ 1GHz$ or $225\mu W @ 2GHz$), which drops to 92fJ/decision for full-swing inputs.

The SA's input equivalent noise was also extracted, by measuring the average number of positive decisions versus ΔV_{in} , as shown in Figure 9.8 (measured over $2.5e9$ samples). To be able to measure the intrinsic noise and avoid influence of hysteresis, decision-cycles with a very high ΔV_{in} alternate with cycles where ΔV_{in} is close to the offset. Fitting the measurements to a Gaussian cumulative distribution gave an rms noise voltage of $V_{rms}=1.5mV$, which includes noise from the measurement setup. The curve-fit matches with the average noise sigma predicted by the equations from Appendix A. The tolerance in the measurements in Figure 9.8 is quite high, which should not be the case when the noise




Figure 9.8: Measured average number of positive decisions as a function of the differential input voltage, together with a fit to a cumulative Gaussian distribution.




Figure 9.9: Bit error rate versus clock skew, at $f_{clk} = 1\text{GHz}$.

would be purely white. Therefore, it is expected that other, non-stationary effects, such as $1/f$ noise or variations in probe contact resistance also play a role.

The sense amplifier output was also connected to a spectrum analyzer to investigate the spectrum of the noise. Flicker noise with a slope of -10dB/dec was visible in the result, with a $1/f$ corner in the order of 100kHz . But it was unclear what part of the flicker noise could be attributed to the sense amplifier and what part to the output buffers (a better approach would be to post-process the bitstream from the sense amplifier instead of the voltage output, but this was not possible at the time of the measurements).

	This work (90nm)	[165] (scaled to 90nm)	[166] (scaled to 90nm)	[164] (scaled to 90nm)
Setup+Hold time	18ps	--	--	40ps
Delay/log(V _{in})	44ps/dec	>70ps/dec	--	--
Input eq. noise	1.5mV	--	--	--
Offset σ_{os}	8mV	9.5-19mV	15mV (un-calibrated)	--
Energy/decision	92fJ	--	110fJ	--

Table 9.1: Sense amplifier comparison.

Supply-induced noise – due to mismatch-related imbalances in the circuit – can also be a problem in sense-amplifiers [166]. For this SA, measurements with a sinusoidal supply variation with $V_{pp}=200\text{mV}$ ($f_{sin}=51\text{MHz}$, $f_{clk}=1\text{GHz}$) increased the sense-amplifier noise with only 2 mV ($V_{offset} = 8\text{mV}$ for the tested sample).

Setup & hold times were extracted from BER measurements around the zero crossings of full-swing input patterns, as shown in Figure 9.9. No bit errors were measured outside an interval of 18ps, so the required setup+hold time is smaller than 18ps (as input jitter is part of the 18ps). A conventional circuit in $0.18\mu\text{m}$ CMOS [164] achieves 80ps, which would still be 40ps in 90nm CMOS according to scaling theory. In the double-tail topology, the setup+hold time could be further reduced with a wider tail transistor M9, but at the expense of increased offset and noise due to a shortening of the time that M5/M6 operate in saturation. Simulations predict that the current aperture time is already fast enough to sample data patterns of 40Gb/s, provided that interleaving is used to enable a suitable long regeneration phase.

9.7 Sense amplifier conclusions

Compared to conventional dynamic sense amplifiers, the double-tail topology has an added degree of freedom that enables better optimization of the balance between speed, offset, power and common-mode voltage.

This claim is supported by comparing the performance figures with other sense amplifiers, as shown in Table 9.1. For a fair comparison, the published data from the various sense amplifier publications has been scaled to its equivalent value in a 90nm CMOS process, assuming standard scaling rules.

The double-tail sense amplifier also has a better isolation between input and output (lower kickback noise) and it can operate at lower supply voltages than its conventional counterpart.

Chapter 10

Transceiver on the second demonstrator IC

10.1 Introduction

For the second demonstrator IC, the focus was on the implementation of more power-efficient transceiver concepts. To this end, the capacitive pre-emphasis transmitter was developed for the transmitter side, and the power optimized sense amplifier with build-in DFE for the receiver side.

The capacitive pre-emphasis transmitter was originally introduced at the ISSCC 2007, independently by both Ho et al. [74] and our research group [62]. In [121] Ho et al. presented more details on their work and provided a qualitative intuitive explanation of the capacitive driver technique. The transceiver from this project was discussed in more detail in [124].

The background behind the capacitive pre-emphasis transmitter was already discussed in section 4.2.2 and the DFE and sense amplifier were discussed separately in section 7.6 and Chapter 9 respectively. In this chapter the complete transceiver is discussed in more detail, based on [62, 124].

To enable a good comparison with the first demonstrator IC, an interconnect length of 10mm was also used for the second IC. The first demonstrator IC also showed that twisted differential interconnects are efficient to cancel crosstalk, therefore no other crosstalk reduction techniques were investigated for the second transceiver. Basically, the only difference between the interconnects on the first and second transceiver are the small differences due to technology, as a 90nm CMOS process was used for the latter (versus 130nm for the first). The interconnects are hence not discussed in detail again in this chapter.

The next section in this chapter discusses the effect that the termination has on bandwidth and power for these particular interconnects. Section 10.3 discusses the implementation of the transceiver. Section 10.4 discusses the chip top-level and measurement setup and section 10.5 presents the results. The chapter is concluded in section 10.6.




Figure 10.1: Bandwidth for three different termination schemes. The wire parameters are $R_{wire}=2k\Omega$, $C_{wire}=2.8pF$.

10.2 Effect of termination on bandwidth and power

As was discussed in section 4.2.3, the power consumed in an RC-limited interconnect depends on its source (Z_S) and load impedances (Z_L). In Figure 10.1, this is shown for three different types of termination.

The conventional transceiver with inverters as both the transmitter ($Z_S=100\Omega$) and receiver ($Z_L=10fF$) has only 62MHz bandwidth and high power consumption. The current-sensing scheme (with $Z_L=190\Omega$ in Figure 10.1) increases the bandwidth up to 3 times, but with increased power at low data activities. The capacitive transmitter (with $Z_S=255fF$ in Figure 10.1), has roughly the same bandwidth improvements (not exactly in the figure, because the impedance ratios are slightly different).

In terms of power consumption, the three types of termination from Figure 10.1 show quite different behavior. As a measure for the power consumption, we use energy per bit and plot this as a function of data activity or transition probability (p_{trans}). Ideally, we would like the energy per bit to be dependent on p_{trans} with zero energy consumption if $p_{trans} = 0$ (no static power consumption). Figure 10.2 shows the energy consumption for the 3 different schemes, assuming a binary (zero-mean) Markov source as derived in 3.5.2.

In Figure 10.2 it is visible that the resistive termination scheme (current-sensing) has a large static energy consumption, as it requires a static current to maintain a non-zero voltage across a resistor. Still it is more energy efficient than a conventional scheme with high




Figure 10.2: Power consumption for the three different termination schemes from Figure 10.1 (obtained with slightly different $R_L=233\Omega$ and $C_S=311fF$).

swing along the whole line, except for very low p_{trans} where static power dominates. The capacitive transmitter scheme has the lowest energy consumption (lowest slope and no static energy consumption). It is also visible in the figure that the capacitive transmitter scheme has much lower energy consumption than the resistive termination scheme, while both increase the achievable data rate by about the same factor.

There are two reasons for this lower energy consumption. The first is the static energy consumption of the current-sensing scheme that is not present in the capacitive transmitter scheme. This is because the current-sensing scheme has a resistive path from V_s to ground, which lacks in the conventional and capacitive transmitter (see Figure 10.1). The resistive path gives a static current of about $\pm 0.3mA$ for a constant 1 or 0 ($V_{dd}/2=0.6V$ divided by the total series resistance of about $2Kohm$). This leads to at least $0.3mA \cdot 0.6V / 1.25Gbps = 0.14pJ/bit$ static energy dissipation for current-sensing. If the current-sensing amplifier is not modeled by a simple resistor but with a transimpedance amplifier, substantial bias current is also needed to realize a sufficiently low impedance. On the first demonstrator IC for example, the transimpedance amplifier consumed between 0.5 and 0.9pJ/bit (tabulated in Table 8.2).

The second reason for the attractiveness of the capacitive transmitter is the associated lower voltage swing on the interconnect, which reduces dynamic power. Although for both cases, the voltage swing at the receiver end of the interconnect is the same, the capacitive transmitter scheme has this low voltage swing along the entire interconnect, while the resistive termination scheme has a linearly increasing voltage swing towards the transmitter.

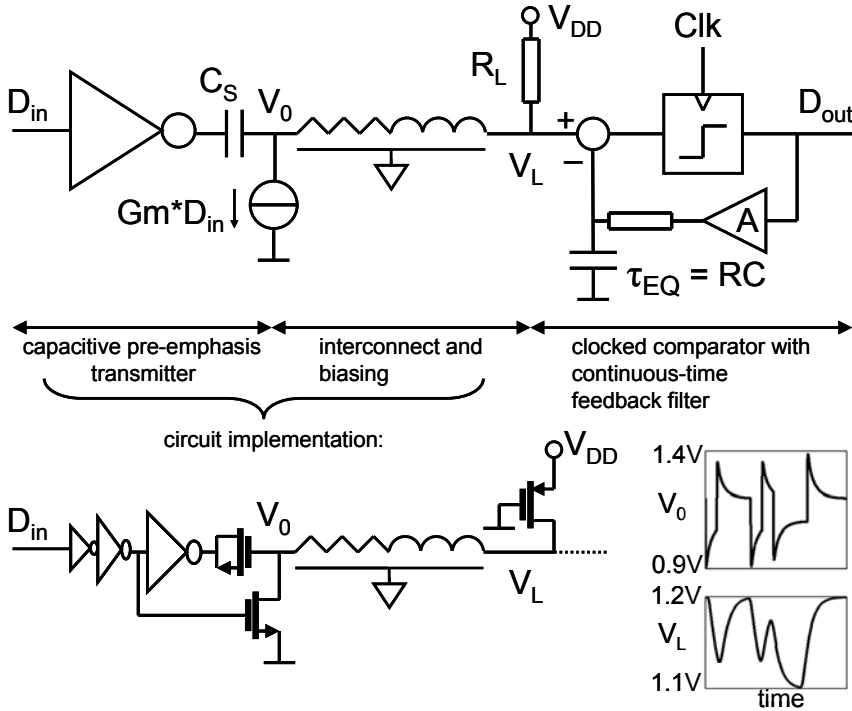


Figure 10.3: Transceiver overview.

10.3 Transceiver implementation

One half of the differential transceiver is shown in Figure 10.2 and its parts are discussed below.

10.3.1 Capacitive pre-emphasis transmitter

The series capacitance for the transmitter is made with an NMOS transistor. Due to the thin gate oxide, the area of the capacitor can be kept rather small compared to the area of the interconnect. A possible problem of the capacitive transmitter is the ill-defined DC potential on the interconnect. In order to define this DC potential for high and low V_{in} , a current source with current $G_M D_{in}$ is added at the transmitter side and a resistance R_L at the receiver side, as was also shortly discussed in section 4.2.3. If D_{in} switches between 0 and V_{DD} , V_{out} switches between V_{DD} and $V_{DD} - G_M V_{DD} R_L$. The low frequency voltage swing on the interconnect is thus $G_M R_L V_{DD}$, which is chosen at $0.08 \cdot V_{DD}$ ($V_{DD} = 1.2V$). By choosing G_M small (narrow NMOST) and R_L large (narrow long PMOST), the static energy consumption is kept small. We used $G_m \approx 5\mu S$ and $R_L \approx 16k\Omega$, which renders about $1.2V \cdot 5\mu S \cdot 16k\Omega \approx 100mV$ swing at a current which is switched between 0 and $6\mu A$. For a differential line, one current is on while the other is off, so a total continuous bias current of $6\mu A$. For gigabit communication this leads to a negligible power overhead, e.g. $<0.01pJ/b@1Gbps$), while the dynamic power consumption is in the order of $0.15pJ/b$ for 50% transition probability.

With this setup, the transfer function at low frequencies is controlled by G_M and R_L , while at high frequencies the capacitive path via C_S and C_{wire} dominates. To ensure a smooth transition region, the transfer of the two paths have to match. With a first-order approximation for the wire, it is derived in [124] that the poles and zeros of the parts of the transfer cancel each other when:

$$C_S = \frac{G_M \cdot R_L \cdot C_{wire} \cdot (1 + G_M \cdot R_{wire})}{1 - G_M \cdot (R_{wire} + R_L)} \quad (10.1)$$

When we assume that $G_M R_{wire}$ is much smaller than 1 (in other words, assume that G_M is much too small to charge the wire), then (10.1) simplifies to:

$$C_S \approx \frac{G_M \cdot R_L \cdot C_{wire}}{1 - G_M R_L} \Leftrightarrow \frac{C_S}{C_{wire}} \approx \frac{1}{\frac{1}{G_M R_L} - 1} \Leftrightarrow \frac{C_S}{C_{wire} + C_S} \approx G_M \cdot R_L \quad (10.2)$$

The right part of (10.2) is an intuitive result, which states that the magnitude transfer of the capacitive path should match magnitude transfer of the resistive path.

When we furthermore assume $C_S/C_{wire} \ll 1$ to obtain small-swing transfer (which is the same as assuming $G_M R_L \ll 1$), then the result can also be expressed as $C_S/G_M \approx R_L C_{wire}$ (as was done in [124]) which will slightly underestimate the required C_S .

Simulations showed that inequality of the time constants have only a modest effect on the eye-opening, so that process variations can be tolerated if nominally equal time constants are chosen at design time (as also discussed in section 7.7.1).

The total area of the transmitter is $226 \mu\text{m}^2$, where about $100 \mu\text{m}^2$ is used for two MOSFET-line-driver capacitors, each $C_S \approx 311 \text{fF}$. This is 5 times smaller than the metal capacitors used in [121] (40×20 metal tracks to implement a capacitive line driver, taking already about $480 \mu\text{m}^2$). It comes at the cost of a more non-linear capacitor, but eye diagram simulations show that the linearity of the capacitance is not very critical. Thus a MOS capacitance with much higher capacitance/area can be used instead of metal-metal capacitance.

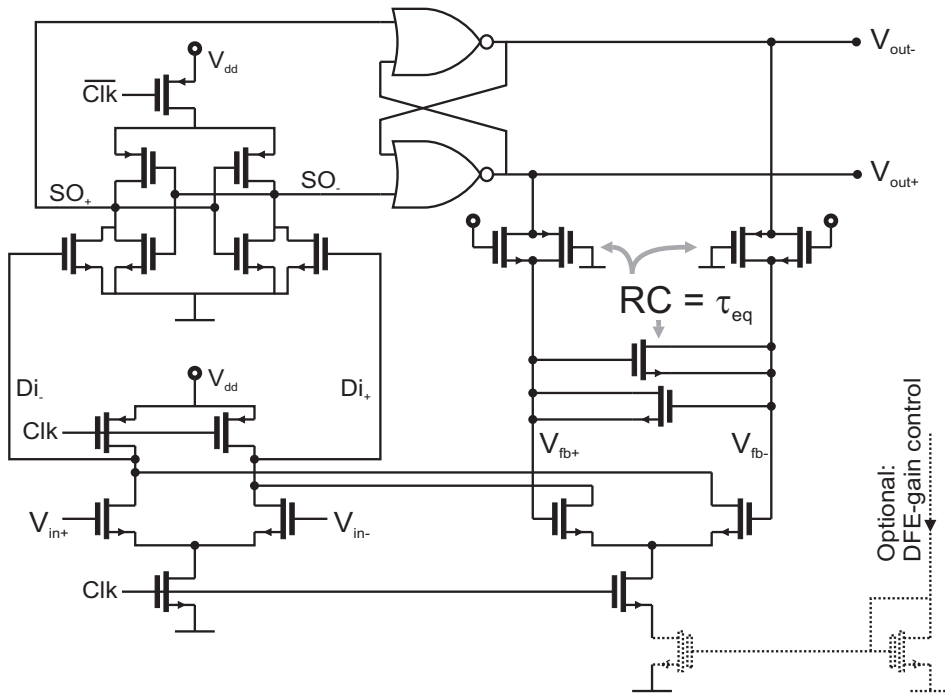


Figure 10.4: Sense amplifier with decision feedback equalization.

10.3.2 Sense amplifier with decision feedback equalization

Due to the reduced signal swing and high data rate, a sensitive receiver is needed with low offset. Simultaneously, it should have a fast decision speed and low power consumption. For a clocked comparator, it is difficult to let the power consumption scale with the data activity, the next best thing is to let its power consumption at least scale with the clock-rate. To combine the above requirements, the double-tail sense amplifier was developed, as discussed in Chapter 9. A DFE circuit is also added to the receiver and build around the sense amplifier, as is visible in Figure 10.4.

The left part of the circuit is the already discussed sense amplifier. The SR-latch behind the sense amplifier is used to convert the dynamic (pre-charged) signals at the SO nodes to static CMOS signals that are valid for a whole clock-period. The outputs of the SR-latch are directly used to drive a low-pass RC filter for the DFE. The feedback voltage from the low-pass filter is coupled back into the sense amplifier via a second differential input stage, as shown on the right of Figure 10.4. This part of the circuit implements the DFE with the continuous time feedback filter, as was discussed earlier in section 7.6.1.

The DFE gain-factor “A” that was discussed in that section and shown in Figure 7.7b is in this implementation defined by the transconductance-ratio of the feedback and main amplifier. In this circuit, the optimal gain depends on the attenuation of the capacitive divider and the amount of ISI. A ‘dynamic’ differential feedback pair is used with a switched tail current, so again no power will be consumed when the clock is inactive. The fact that the feedback output V_{fb} is full-swing, while a differential pair is usually only linear

over a small input range, poses no real problems in this circuit if it is dimensioned properly. The linear range of the feedback differential pair is maximized by giving the transistors a high overdrive voltage (meaning small W and large L). The fact that the tail MOST operates in its triode region also helps to increase the input range, as the R_{ds} of the tail transistor acts as a degeneration resistance when only one of the two transistors of the differential pair is active.

The feedback gain-factor A can be controlled by the tail-current of the feedback differential pair. As was discussed in section 7.7.1, this gain factor can be set at design-time, through proper dimensioning of the clocked tail transistor. If higher data rates are required, then the tail-current can also be controlled at run-time, for example through a current-mirror configuration, as shown with the dashed transistors in Figure 10.4.

The components that determine the time constant, the resistor and capacitor, can be implemented in various ways, but we aimed for small area consumption. That is why the resistors and capacitor producing V_{fb} (see Figure 10.4) have been implemented with MOS transistors, with pass-gates and anti-parallel gate-capacitances respectively. The gate-capacitances of a MOST have a very high capacitance per area, but are also quite non-linear due to the channel-capacitance. The use of an anti-parallel configuration reduces the non-linear effects to tolerable levels.

The total area of the receiver is $117\mu\text{m}^2$ with $32\mu\text{m}^2$ for the DFE part. The simulated power consumption is 0.12pJ/b with 0.02pJ/b for the DFE part at 2Gbps.

10.4 Demonstrator IC top-level and measurement setup

For the interconnects, we used metal 4 wires of $0.54\mu\text{m}$ width and $0.32\mu\text{m}$ spacing, optimized for the highest bandwidth per cross-sectional area. The optimization of the dimensions was done with EM-field solver simulations [2]. The optimum differs slightly from the analytical estimate from section 2.7.1 because the analytical results neglect the doubling of the mutual capacitance between the two halves of the differential interconnect and also neglect fringe capacitances.

The EM-field solver simulations (using data from the technology manual) predicted that the interconnect parameters should be $R_{\text{wire}}=1.34\text{k}\Omega$ and $C_{\text{wire}}=2.8\text{pF}$ (or $C_{\text{wire}}=2.4\text{pF}$ for a single-ended interconnect). However, the R_{wire} was later updated to $2\text{k}\Omega$, as the measurement results indicated this, as will be discussed in section 10.5 in more detail.

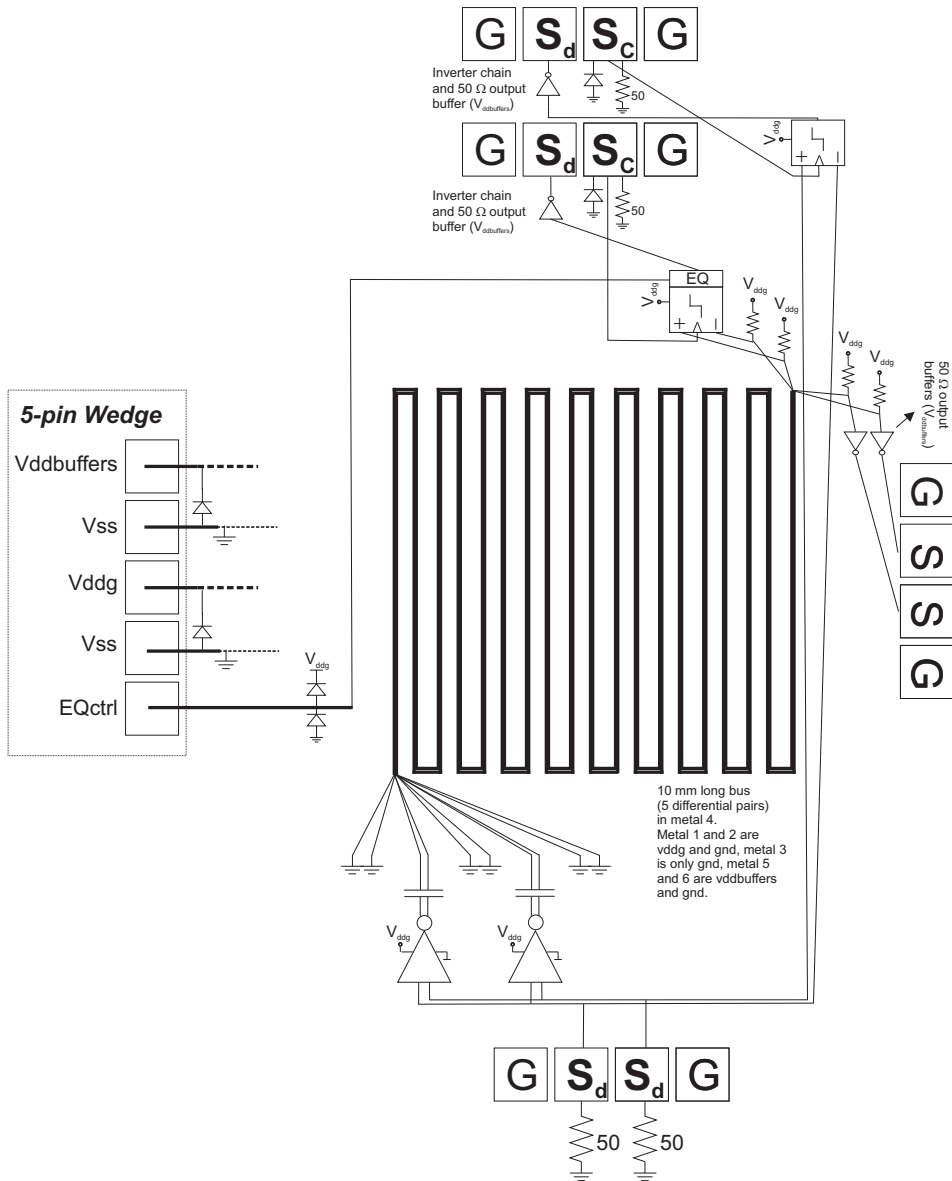


Figure 10.5: Floorplan of second chip transceiver.

The demonstrator IC was fabricated in a 90nm CMOS process. The floorplan of the entire chip is shown in Figure 10.5 and a micrograph in Figure 10.6. The floorplan and supporting circuitry around the transceivers are similar to the first demonstrator IC (see section 8.6) but then a bit simplified. It was originally the intention to add more experiments to the second demonstrator IC, including a bus with conventional repeaters and one with low-swing repeated transceivers. However, due to an unexpected advancement of the tape-out date, the demonstrator IC had to be cut down in complexity. So, in the end, only a differential




Figure 10.6: Second chip micrograph.

bus with two capacitive pre-emphasis transceivers were implemented. No twisting was applied for simplification and because the effectiveness of twisting was already demonstrated with the first chip. Instead, one inactive bus-channel was used as shield in between the two active channels. In an application, the bus should consist of twisted interconnects with staggered twist positions. At the edge of the bus, some kind of shield could be of interest to avoid crosstalk when full-swing data lines would run in parallel to a part of the low-swing wires. This would require only a relatively low area overhead for wide busses.

On the receiver side, one channel was equipped with a sense-amplifier with DFE. The receiver-side of the other channel was connected to a linear 50Ω buffer to enable eye diagram and impulse response measurements. Apart from the receiver, a stand-alone sense amplifier was also included on the demonstrator IC, as was discussed in the previous chapter.

The power consumption of the different parts of the transceivers could be measured separately, by measuring V_{ddg} and disconnecting either the receiver or transmitter side. As all circuits are dynamic, no power is consumed when they are not connected. Also note that the G_M/R_L path (Figure 10.3) is not active, because both interconnect halves are pulled to the supply when the input is removed (because the data inputs are pulled down by the termination resistances).

As was done for the first chip, the receiver clock was generated externally in order to adapt its phase to the eye position and be able to measure eye widths. The measurement equipment that was used was the same as for the first demonstrator IC, as was described in section 8.7.




Figure 10.7: Simulated versus measured eye diagrams at a rate of 1.35Gb/s, with simulated parameters $C_{\text{wire}}=2.8\text{pF}$ and R_{wire} either $1.8\text{k}\Omega$ (a), $2\text{k}\Omega$ (b) or $2.2\text{k}\Omega$ (c).

10.5 Experimental results

The characterization of the wire parameters on the second IC was a bit more difficult than on the first transceiver, as the line was not directly accessible externally. Instead, the RC product of the wire was estimated by fitting measured eye-diagrams to simulated versions as shown in Figure 10.7 (with the best fit in (b) and a 10% lower or higher time constant in (a) and (c) respectively). This showed an RC product that was quite a bit larger than intended.

Based on low frequency step response measurements discussed in the next paragraphs, it was concluded that the wire capacitance was most likely on spec ($C_{\text{wire}}=2.8\text{pF}$), but the wire resistance was higher than intended, namely $R_{\text{wire}}=2\text{k}\Omega\pm 10\%$ instead of the $1.34\text{k}\Omega$ it should have been. It did not really become clear why this was the case. It could be that the sheet resistance just had quite some deviation from its nominal value.

Apart from the high-frequency parameters, there was also quite a shift in the low-frequency G_M/R_L v.s. C_s/C_{wire} parameters, which became visible as an overshoot in the step-response measurement. Eventually, a very good match between simulations and measurements was obtained when the G_M was set at $4.3\text{e-}6\text{S}$, R_L at $16\text{e}3$ and C_{TX} at 290fF , which equal the expected parameters from the simulations and layout extractions, except for C_{TX} which was expected to be 250fF .

There is a chance that the abovementioned parameters were deduced incorrectly and that there is another set of parameters that give a good match between simulations and experimental results. However, a number of situations were tried (e.g. lower R_L than expected instead of a higher C_{TX}) and those did not give a match as good as the one obtained with the above parameters. It could also be that the C_{wire} was larger than expected instead of the R_{wire} , but that would also require an even bigger C_{TX} (a factor 1.5 larger when the R_{wire} would equal its intended value) and even then, the simulations did not give a good match with the measurements.




Figure 10.8: Measured eye-diagram at 1Gb/s and BER versus receiver clock delay. DFE is not used.




Figure 10.9: Measured eye-opening as a function of I_{EQ} for different data rates.

A measured eye-diagram for the capacitively driven line at a data rate of 1Gb/s is shown in Figure 10.8. The measured BER at the edges of the eye is also shown. The BER drops rapidly below a clock skew of -150ps and above 180ps, giving an eye-opening of 670ps.




Figure 10.10: Measured energy consumption per transmitted bit as a function of the transition probability for different data rates, with and without DFE.

Data rates up to 1.35Gb/s were achieved without decision feedback equalization (DFE) at the receiver side (DFE-gain control current $I_{EQ} = 0$). The one- σ offset of the total transceiver is 11mV, measured over 20 samples. Due to this offset, not all samples achieved 1.35Gb/s, but all samples did achieve a slightly lower data rate of 1Gb/s. If desired, area up-scaling could further reduce the offset at the expense of power ($P \propto 1/\sigma_{os}^2$). As was discussed in Chapter 9, offset compensation schemes can be a good alternative if the application allows for the added complexity, which is probably not the case for most on-chip buses. Apart from random offset, simulations over process corners indicate that the circuit is robust to PVT variations at a rate slightly lower than the maximum achievable data rate.

Data rates up to 2Gb/s were measured with DFE. Note that DFE reduces ISI, making the system less vulnerable to offset. Figure 10.9 shows that DFE improves the eye-opening for a wide range of I_{EQ} . In an application I_{EQ} can therefore be fixed at design time.

The measured energy consumption at different data rates is shown in Figure 10.10. The total power for one transceiver was derived in the following way: the power of one transmitter was calculated as half the power of the two transmitters (with the receiver disconnected). The power of the sense amplifier with DFE was calculated as the difference between the power with the receiver connected and with the receiver disconnected.

With random data at 2Gb/s, only 0.28pJ/b is dissipated. The energy dissipation of 0.12pJ/b at zero data activity is mainly due to the energy consumption in the sense amplifier, which has large transistors to get a low offset. Clock-gating can be used to eliminate its energy

consumption during inactive periods. The DFE part of the circuit requires less than 7% of the total transceiver power, while it increases the achievable data rate here with a factor 1.5.

10.6 Conclusions for Transceiver on second demonstrator IC

It was shown that a transceiver with a capacitive transmitter indeed significantly reduces the power consumption compared to the transceiver on the first demonstrator IC. Compared to the transceiver with resistive load, it achieves similar data rate, but both the static and dynamic power consumption are lower.

To further increase the achievable data rate, DFE can be used at the receiver. By using an analog feedback filter, DFE only costs little extra area and power. The transceiver achieves error free operation at 2Gb/s while consuming only 0.28pJ/bit.

This transceiver can be used to cross large on-chip distances with un-interrupted interconnects, while consuming low power. The transceiver techniques can also be used to bridge shorter distances at higher data rates (and at very high offset reliability) as will be discussed in the next chapter.

Chapter 11

Transceivers for networks on chips

11.1 Introduction

As was discussed in Chapter 2, on-chip communication is not only problematic because on-chip interconnects are becoming a speed, power and reliability bottleneck, but also because systems on chips (SoCs) start to become so complex that they require new interconnection approaches [13, 14].

Networks on chips (NoCs) have emerged as the seemingly best candidate to connect the many functional elements on present and future SoCs [13-18]. Most of the long (global) interconnects, which have the severest bandwidth limitations and crosstalk problems, are eliminated in a NoC, especially when mesh-like network configurations are used. A NoC also enables easier clock-distribution with alleviated skew requirements and less power consumption as the various processing elements can operate mesochronous [15-17] or asynchronous [18] to each other, using for example the ‘Globally Asynchronous, Locally Synchronous’ (GALS) design style.

Still, even in a NoC configuration, the network interconnects and especially the routers can consume a considerable part of the total power budget. In [17] for example, the on-chip network consumes up to 39% of the total chip power (76W when operating at 5.1GHz) [179]. 17% of the network power is consumed in the links (13W at 5.1GHz).

A NoC can therefore benefit from link-transceivers that are more advanced than the standard inverters. High-speed, low-power transceivers can for example facilitate network topologies with longer and more wires than the standard mesh topology, such as a (folded) torus or star topology, to simultaneously reduce the interconnect power and the average hop count, and hence also the latency and the associated router power [14, 15].

Low-power transceivers were proposed in the past [15, 85], but they did not improve the data rate. However, as was discussed in the previous chapter, a capacitive pre-emphasis transceiver can both increase the achievable data rate and decrease the transmission power.

In this chapter, we will adapt these techniques for NoC applications and compare the resulting transceiver with other common types of transceivers, based on our publication in the transactions on VLSI circuits [180]. Other topics that were not covered in earlier chapters are the optimization of the circuit for yield versus power and the addition of




Figure 11.1: Overview schematic of the proposed transceiver for NoCs.

synchronization circuitry. Yield is an important issue given PVT variations, random mismatch, crosstalk and the fact that many transceivers will be present on a NoC.

A schematic overview of the proposed NoC transceiver is shown in Figure 11.1. The transmitter uses a series capacitance to lower the swing on the interconnect, increase its bandwidth and lower the power dissipation, as was tested on the second demonstrator IC. The interconnects consist of twisted differential pairs to be robust towards disturbances such as supply noise and crosstalk and the double-tail sense amplifier clocks the data at the receiving end and regenerates it to full swing. A difference with the transceiver tested on the second demonstrator IC is that the sense amplifier does not include decision feedback equalization (to simplify the transceiver) and that a clock or strobe channel is present alongside the data-channels to enable source-synchronous operation.

This chapter is organized as follows. The following section discusses data links for networks on chip and the drawbacks of conventional transceivers. Section 11.3 describes the improved low-swing transmitters and section 11.4 discusses the accompanying receivers. Section 11.5 includes synchronization in the discussion and describes the entire transceiver. The chapter ends with the conclusions in section 11.6.

11.2 Data communication on a NoC

11.2.1 Interconnects for Networks on a Chip

In this discussion about transceivers for NoCs, we will focus on interconnects that span one or two processing tiles. A wire length of 2mm is assumed throughout the discussions, but the same techniques apply to a variety of lengths. The transceiver presented in the previous chapter focused on much longer (10mm) wires and contains some additional equalization circuitry to boost the data rate. Wires of 2mm have a much higher intrinsic bandwidth (the RC product scales with the length squared [8]), so we will focus here on slightly simpler transceivers and leave out the receiver equalization.

We also assume that the interconnects are used unidirectional, as bidirectional use of the interconnects complicates the design of fast and power-efficient transceivers. Bidirectional communication can be implemented with a second set of interconnects, as is often done in NoCs.

To maximize the throughput between two routers, it makes sense to use wide data paths [14] with many densely packed interconnects. In section 2.7 it was shown that the cross-sectional dimensions of interconnects should be chosen roughly equal to optimize the bandwidth per cross-sectional area (BW/Area). A bus with these optimized interconnects will have the highest achievable throughput for a certain bus area. Wires in the thick (reverse-scaled) top-metal layers will have lower resistance and have higher bandwidths so it makes sense to use the top metal layers for the link when the data rate per wire is a limiting factor [13, 14]. However, the BW/Area is roughly the same as for thinner metal layers, so one could choose to also use the lower metal layers for the link. In this last case, certain areas of the chip could be dedicated to the link interconnects to enable high throughput in a well defined link environment.

Twisted differential interconnects are used, because despite the doubling of the wires/channel, they can still increase the BW/Area, especially when combined with the capacitive pre-emphasis transmitter. This is because differential interconnects enable more robust transceivers that hardly suffer from crosstalk and can thus operate at higher speeds and at a lower swing [85], as was discussed earlier in section 4.4.1. A comparison between single-ended and differential NoC transceivers is also included in the discussion below.

In the 1.2V 6M 90nm CMOS process that was used for the second demonstrator IC, metal-4 wires with a width of 0.54 μm and a spacing of 0.32 μm have the highest BW/area under the assumption that the wires are surrounded by other wires in all directions. Under these conditions, the interconnect parameters are:

$$R_{\text{wire}} = 200 \Omega/\text{mm} \quad , \quad C_{\text{wire}} = 280 \text{ fF}/\text{mm} \quad (11.1)$$

or $C_{\text{wire}}=240\text{fF}/\text{mm}$ for single-ended interconnects [62].

With these dimensions, one differential channel will have a pitch of 1.72 μm . A link with for example a length of $l=2\text{mm}$ and a width $w=64\text{bits}$ in both directions will occupy an area of $2wl \cdot 1.72\mu\text{m} = 0.44\text{mm}^2$ when placed in one metal layer, which could still easily fit above a 2x2mm tile. When five metal layers would be available to connect routers in a mesh topology with e.g. $N \times N = 5 \times 5$ tiles of 2x2mm each, then the total link area would become




Figure 11.2: Conventional transceiver schematic.

$2N(N-1) \cdot 0.44/5 = 3.5\text{mm}^2$, only 4% of the tile area of 100mm^2 . The total wire-length would then be: $2 \cdot 2N(N-1) \cdot 2wl = 2 \cdot 2 \cdot 5 \cdot 4 \cdot 2 \cdot 64 \cdot 2\text{mm} = 20.48\text{m}$.

11.2.2 Conventional data transmission

In conventional digital IC design practice, interconnects that are used for chip-wide data communication are simply treated as part of the normal digital design flow, perhaps with a few additional steps such as the (automated) placement of repeaters, to minimize the delay per interconnect length [7].

An example of a ‘conventional transceiver’ for data communication on a NoC is shown in Figure 11.2. It does not have repeaters because delay optimal repeater insertion comes at the price of about 90% increase in power consumption (the C_{gates} add 60% to the total capacitance [8] and the C_{drains} add another 30%). Furthermore, for these relatively short wires, repeaters reduce the delay only marginally [8] as the dominant time constant of the interconnect itself is still only $\frac{1}{2}R_{\text{wire}}C_{\text{wire}} = 96\text{ps}$. To be able to approach this intrinsic wire speed, the transmitter from Figure 11.2 does need to use a buffer-cascade with a large and power-hungry driver. Later on in section 11.3, it will be shown that it is also possible to use a smaller and more power efficient low-swing capacitive transmitter.

In classical synchronous systems, the maximum delay of a combinatorial logic stage is limited to the clock period – or vice versa: the clock-rate is limited by the stage with the maximum delay – and this constraint is usually also imposed on the data transceivers. But such a constraint is not necessary for a communication channel as is often demonstrated in wireline communication where several bits can be in flight along the channel at any given time. The channel bandwidth is the real limiting factor for the data rate. For on-chip transceivers, it is also easy to achieve data rates higher than $1/T_{\text{delay}}$ provided that proper clocking schemes are used, such as pipelined or source synchronous schemes, as will be demonstrated in section 11.5.

Without additional layout measures, a conventional transceiver is not very suitable as a high-speed transceiver, because its delay can vary widely due to crosstalk [8]. Figure 11.3 shows the effect of capacitive crosstalk between neighboring data wires in a bus. The average delay of the transmitter and the 2mm of interconnect amounts to 205ps, but the delay speeds up to 160ps when neighboring aggressors make a transition in the same direction and the delay increases to 262ps when the neighboring aggressors switch in the opposite direction. Crosstalk not only creates this varying delay (reduced eye-width), but it also decreases the voltage noise margin (reduced eye-height) as is visible in the figure, and as was also shown in section 5.3.3. Above a certain data rate, crosstalk from specific aggressor data patterns can even prohibit proper detection of data bits, as visible for the bit in the victim signal at $t=1.9\text{ns}$. Quantitatively, crosstalk between neighboring wires in one metal layer can decrease the achievable data rate by more than a factor two, as was




Figure 11.3: Signals at 5Gb/s for three neighboring channels from a conventional transceiver.

analyzed in section 6.2.1. Crosstalk problems become even worse when the surrounding metal layers are also used as data paths.

A conventional transceiver is also not very power efficient as the transmitter needs to fully charge and discharge large wire and driver capacitances. The setup shown in Figure 11.2 consumes 775fJ per upward transition (consumed mainly to charge the wire) and 65fJ per downward transition (consumed to charge the driver capacitances), which averages to 420fJ/transition. As an example for what this would cost on an entire chip, assume the same situation as earlier with 2mm long 64bits wide links in both directions, used in a mesh of 5x5 tiles. Furthermore assume a clock-frequency of 5GHz for the links, with an average switching-activity of about $p_{act} = 25\%$ (heavy traffic). Then the total link power becomes $P = 2N(N-1) \cdot 2w \cdot p_{act} \cdot E_{trans} \cdot f_{clk} = 2 \cdot 5 \cdot 4 \cdot 2 \cdot 64 \cdot 0.25 \cdot 420 \cdot 10^{-15} \cdot 5 \cdot 10^9 = 2.7W$, which is not acceptable for low-power applications such as mobile baseband processors [18]. The reported link power for the 80-tile NoC from [17] is even higher: 13W at 5.1GHz.

11.2.3 Link improvements

It is well recognized that low-swing signaling can reduce the interconnect power consumption [85, 181], but at the cost of a reduced noise margin. The degradation of data integrity due to supply- and substrate-noise increases as the swing goes down. Crosstalk also becomes an even more severe problem, especially when a full-swing aggressor interconnect is routed in the vicinity of a low-swing victim.

Fortunately, the regular nature of the top-level wiring in a NoC and the re-usability of the interconnection links justify a slightly higher design-effort to better optimize the wires [14]. In this way, routing of full-swing wires next to low-swing wires can be avoided, as well as the routing of far-end wire parts next to near-end ones. Application of these simple rules leaves only the crosstalk between the different wires from the same bus, with the neighbor-to-neighbor crosstalk as dominant part.

Application of twisted differential wires can effectively mitigate neighbor-to-neighbor crosstalk, needing only one twist in every even wire pair and two twists in every odd pair, as was discussed in section 4.4 and indicated in Figure 11.1. When a multi-layer bus is used,




Figure 11.4: Low-swing transceiver with multiple V_{DD} 's.

then additional twists may be used to further reduce crosstalk (section 4.4.7). The optimal positions of these twists depend on the type of wire termination. With equal impedances for transmitter and receiver, intra-bus crosstalk is perfectly canceled and the optimal twist positions are symmetric around the midpoint [82].

In the presented transceiver, with the capacitive termination at both transmitter and receiver-side, practically all crosstalk is canceled as will be shown later.

11.3 Low-swing transmitters

The energy-cost for a rising edge with swing V equals the well-known $E=CV^2$, as was discussed in section 4.5.1. Half of this energy is dissipated during charging. The other half is stored in the interconnect and dissipated at a later time when the interconnect is discharged (the resistance of the interconnect prevents efficient charge-recycling techniques as discussed in section 4.5.3). To reduce the link power it hence makes sense to reduce the swing. If only a single supply voltage is available and active circuits are used to reduce the swing, there is no quadratic but linearly relation with the swing ($E=CV_{DD}V_{swing}$). When a dedicated supply voltage is available to generate the low-swing signal, then the power is again quadratically dependent on the swing. Many low-swing techniques with a dedicated supply voltage (either generated on- or off-chip) for the transmitter have therefore been introduced in the past [15, 84, 85, 182].

The need for a dedicated supply voltage is quite a drawback, but the use of multiple supply grids becomes more accepted now that SoC-designs start to use multiple supplies (multiple voltage islands). SoCs use for example a high voltage (V_{DDH}) for the high performance (logic) parts and a slightly lower voltage (V_{DDL}) for the slower parts of the chip. Low-swing interconnect drivers could switch between these two supplies to generate the low-swing signal, with equal power efficiency as the dedicated supply variant, but without the need for yet another supply grid. An example schematic of such a low-swing transceiver is shown in Figure 11.4.

This variant still has several drawbacks. A first drawback is the fact that the noise-margin is directly related to the amount of supply-noise and a short droop in one of the two supplies can easily introduce a bit-error. Tight coupling between the two supplies, to lower the differential ($V_{DDH}-V_{DDL}$) noise, could reduce this problem, but at the expense of area overhead for example for coupling capacitors. A second drawback, which is found in most low-swing transmitters, are the large transistors that are needed to drive the interconnects

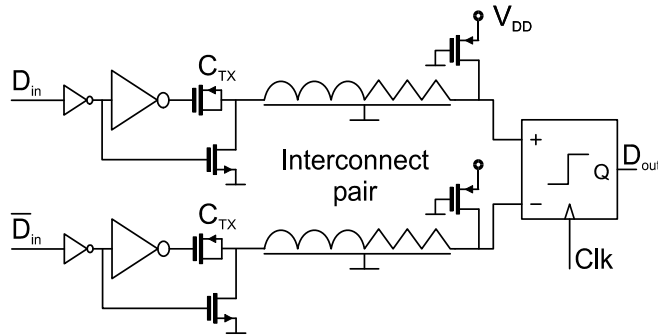


Figure 11.5: Proposed low-swing capacitive pre-emphasis transceiver.

with sufficient speed. Driving these large transistors costs a lot of power and hence decreases the efficiency.

To circumvent these drawbacks and simultaneously increase the achievable data rate, we propose to use capacitive pre-emphasis transmitters, as was discussed in more detail in the previous chapter. The capacitive transmitter uses a series capacitance (C_{TX}) to drive the interconnect, as shown in the overview schematic in Figure 11.1. This capacitance, together with the wire capacitance, acts as a capacitive divider which reduces the swing by a factor of $C_{TX}/(C_{wire}+C_{TX})$. The capacitive transmitter also increases the bandwidth of the interconnect, as C_{TX} emphasizes each transition with an overshoot. Compared to the low-swing transmitters that switch between supplies, the capacitive transmitter is much less sensitive to supply noise, as the capacitor divider also attenuates this noise. It does furthermore not require a special supply voltage and the lower theoretical efficiency ($E=CV_{swing}V_{DD}$) is more than compensated by the reduction in energy overhead at the driver side.

To illustrate these claims, the capacitive transmitter and the multiple-Vdd circuit were simulated and compared. The implementation of the capacitive transmitter that was used for the comparison is shown in Figure 11.5. It uses a MOST as C_{TX} , as the high capacitance-density of the gate-oxide makes it very suitable as transmitter capacitance [62]. For the 2mm interconnects, a MOST with $W=L=2.7\mu\text{m}$ gives a swing reduction to 10% of the supply voltage. A PMOST channel-capacitance is used with the gate connected to the driver to avoid loading the driver with the junction capacitances. An NMOST (current-source) at the Tx-side and a PMOST (resistive) load at the Rx-side define the low-frequency behavior and DC-operating point as discussed in section 10.3 and these are narrow and long transistors to minimize the static current.

Some signal waveforms of both circuits are shown in Figure 11.6 on the next page, which clearly illustrate the pre-emphasis effect of the capacitive transmitter. Numerical results are shown in Table 11.1, which also includes the simulation results of the conventional full-swing transceiver from Figure 11.2.




Figure 11.6: Signals at 5Gb/s for the multiple V_{DD} transmitter (a) and the capacitive transmitter (b).

	Conventional full-swing	Multi- V_{DD} low-swing	Capacitive low-swing
Technology	1.2V, 6 metal, 90nm CMOS		
Interconnects	2mm, $R_{wire}=400\Omega$		
	Shielded single ended $C_{wire}=480fF$	Twisted differential $C_{wire}=560fF$	
Supply	1.2V	$V_{DDH} = 1.2V$ $V_{DDL} = 1.08V$	1.2V
Voltage swing	1.2V	120mV	120mV
Driver size	$W_n=8\mu m$ $W_p=20\mu m$	$W_p=20\mu m$	$W_n=1.6\mu m$ $W_p=4\mu m$
Energy/trans			
Total	420fJ	135fJ	105fJ
Wire (theory)	346fJ	8fJ	80fJ
Tx overhead	74fJ	127fJ	25fJ
Static power	$\sim 9nW$ (leakage)	$\sim 10nW$ (leakage)	$6\mu W$
Data rate	fully shielded		
50% eye opening	5Gb/s	5Gb/s	9Gb/s
zero eye opening	9Gb/s	9Gb/s	12Gb/s
Delay (50%)			
Transmitter			
Nominal /Slow	105ps/150ps	90ps/130ps	60ps/80ps
($T=25^\circ C/T=100^\circ C$)			
Interconnect	100ps	115ps	80ps

Table 11.1: Comparison of the different transmitters.

Both low-swing circuits have the same voltage swing and the driver sizes were chosen such that the circuits could reach 5Gb/s with an eye-diagram that would be at least 50% open. This means that a relatively large driver is needed for the multiple- V_{DD} circuit, which creates a significant overhead of 127fJ/transition; 16 times more than the energy that is theoretically consumed. The capacitive transmitter has only 25fJ overhead on top of its theoretical energy as the series capacitance reduces the capacitive load seen by the driver and hence enables a smaller driver-size.

In total, the capacitive transmitter is the most power-efficient (total of 105fJ/transition). The smaller driver chain also has less delay and the pre-emphasis effect provides a higher achievable data rate of 9Gb/s with 50% vertical eye opening versus 5Gb/s for the other two circuits. The conventional full-swing transmitter can only achieve this 5Gb/s when every signal wire is fully shielded from any neighbors, to mitigate crosstalk. Compared to the conventional transmitter, the capacitive transmitter operates at four times lower power consumption, despite the fact that it uses two active wires per channel instead of one.

The table also shows that the delay of the capacitive transmitter increases with 20ps (33%) at the slow process corner and 100°C temperature. The delay of the conventional alternatives increases by a larger margin of 42%/44%.

The swing (vertical eye-opening) of the capacitive transmitter is affected by process variations, mainly because the N- and PMOST that define the magnitude of the low-frequency transfer spread with respect to each other (the capacitance ratio C_{TX}/C_{wire} is more stable). This effect can reduce the swing in the worst-case corner to 95mV. Compared to the other low-swing transmitter, which has to cope with supply variations that can easily amount to ± 100 mV, this is still quite stable behavior.

The transistors that define the low-frequency behavior in the capacitive transmitter also cause a bit of static power consumption. However, the dynamic power easily dominates the static part for data rates above 90MHz (assuming random data). When the link is not used, it is easy to stop the static power consumption by setting both the D_{in} and the D_{in} -not high, to break the current-path from the transmitter NMOSTs through the wire to the PMOST loads at the receiver.

When the link is in use, the receiver PMOSTs operate in triode and act as large resistances, connected to the (local) V_{DD} . Note that this configuration makes the capacitive transceiver well suited to cross (bridge) voltage domains, which can be an advantage in SoCs that operate with multiple voltage islands. This capability is both due to the differential nature and due to the fact that the DC operating point (common-mode voltage) is determined locally at the receiving end, which is good for robust operation of the sense amplifier. This in contrast to the multiple-supply transceiver which has its common-mode defined at the transmitting end.

The PMOST resistances are connected to the highest available reference: the (local) V_{DD} , which is not only simple, but is also beneficial for the channel-capacitance density of the C_{TX} -PMOST which is highest when it reaches strong inversion. Connecting the (PMOST) resistances to the supply does however require that the receiving sense amplifier can cope with an input common-mode voltage that is close to V_{DD} . TO this end, a double-tail sense amplifier as was discussed in section 9.3 is ideal.

11.4 Receiver and optimal swing

The double-tail sense amplifier is used as receiver because its is fast and can operate over a wide common-mode and supply voltage range. The offset of the double-tail sense amplifier is also stable and does not increase significantly for high input common-mode levels, which is attractive for this application.

The schematic of the circuit was already shown earlier in Figure 9.2 together with its signal behavior. To create static output signals, an SR-latch can be added at the output of the circuit or two sense amplifiers can be interleaved as shown in the next section.

Offset is the bottleneck for the sense amplifier in this application (the measured rms noise is a factor five lower than the σ_{os}). Therefore, the transistor dimensions of the double-tail sense amplifier were optimized relative to each other to get the lowest offset standard deviation (σ_{os}) per unit of power cost. Width scaling (or impedance or area scaling) can subsequently be applied to all the transistors together to match the offset standard deviation to the desired specification ($P \propto 1/\sigma_{os}^2$) [176] while maintaining the original speed characteristics.

The offset specification depends on the signal swing (eye opening) and the required yield and reliability. With a swing that equals for example six times the offset standard deviation ($6\sigma_{os}$), the probability that a sense-amplifier will introduce bit-errors due to its offset is only $\sim 2 \cdot 10^{-9}$. With Q being the cumulative Gaussian distribution function (also see Appendix A) and Y being the yield-factor in terms of sigma (six in this case), this value is easily calculated as:

$$p_{defect} = 1 - yield = 2Q(-Y) \quad (11.2)$$

For the earlier introduced 25-tile NoC example with $2 \times 5 \times 4 \times 2 \times 64 = 5120$ sense amplifiers on a chip, the chance for offset related bit-errors would still be only 10ppm. A double-tail sense amplifier that has an offset standard deviation of 10mV (according to 1000 Monte-Carlo simulations) consumes about 90fJ/bit. This sense amplifier can be scaled-down to get an offset of 20mV, when a 6σ yield per sense amplifier is desired at 120mV swing. The energy times offset-variance ($E\sigma^2$) remains constant, so the corresponding energy consumption will be $90/(20/10)^2 = 22.5$ fJ/bit.

The values for the swing and yield-factor above are not chosen randomly but actually define a power optimum, due to the trade-off between transmitter and receiver power. The energy that is consumed in the transmitter, including the interconnect, has a more or less fixed overhead part and a part that is proportional to the swing:

$$E_{TX} = p_{activity} (E_{overhead} + C_{wire} V_{DD} V_{swing}) \quad J/bit \quad (11.3)$$

Where $p_{activity}$ is the data activity (transition probability). The energy consumption of the sense amplifier is inversely proportional to the square of the offset and the required yield parameter Y relates offset to swing:

$$E_{RX} = \frac{E\sigma^2}{\sigma_{allowed}^2} = \frac{E\sigma^2}{(V_{swing}/Y)^2} \quad J/bit \quad (11.4)$$




Figure 11.7: Energy consumption versus swing.

With the substitution of $E\sigma^2=90\text{fJ}\cdot(10\text{mV})^2$, $Y=6$, $p=0.5$ (random bits), and the data from Table 11.1, a graph can be plotted of these two equations and their sum, as shown in Figure 11.7. The figure clearly emphasizes the advantage of low-swing signaling. At large signal swings, the lowered sense amplifier power can by far not compensate for the increase in line power and full-swing signaling would cost over 5 times more power than signaling with the optimal swing. For the given parameters, this optimum is indeed about 120mV (125mV to be exact).

The optimum is also analytically solvable by taking the sum of E_{TX} and E_{RX} , differentiate, and solve for zero:

$$\frac{d(E_{TX} + E_{RX})}{dV_{swing}} = 0 \Leftrightarrow p_{activity} C_{wire} V_{DD} - 2 \frac{E\sigma^2 Y^2}{V_{swing}^3} = 0 \Rightarrow$$

$$V_{swing\ opt} = \sqrt[3]{2 \frac{E\sigma^2 Y^2}{P_{activity} C_{wire} V_{DD}}} \quad (11.5)$$

The equation shows that the optimum is only weakly dependent (with a third-order root) on properties such as C_{wire} and $p_{activity}$, so the optimum will not change much for different wire lengths or different data activities. We can make the reasonable assumption that the energy consumed in the sense amplifier is, at a given offset, quadratically proportional to the supply ($E\sigma^2 \propto V_{DD}^2$). Under that assumption, the optimal swing is proportional to the third-order root of the V_{DD} and a change in supply voltage will also have only a small influence on the optimum.

A change in technology should even have no influence on the optimum swing when we assume feature size (s) scaling with classical Dennard scaling rules (see section 2.5 or [1]). First, the $C_{wire}/length$ does not change much over different technologies [8], but in a NoC, the size of the tiles and thereby the lengths of the wires probably do scale, so $C_{wire} \propto s$.

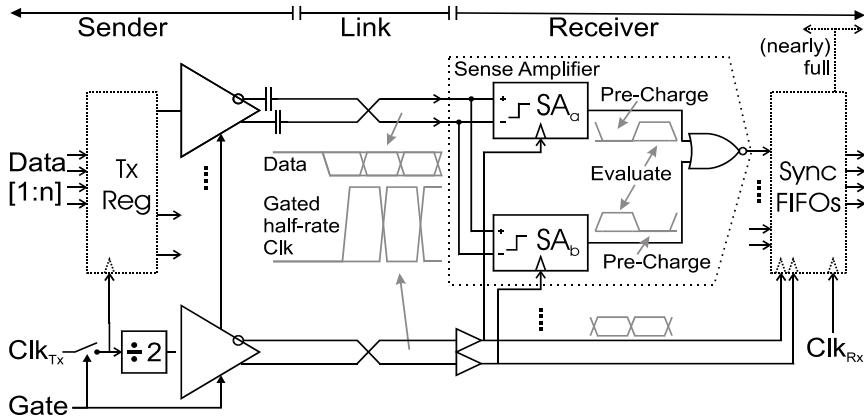


Figure 11.8: Complete Transceiver.

Second, V_{DD} (ideally) scales with s . Third, the energy scales with $E \propto C_{MOS} V_{DD}^2$ and with $C_{MOS} \propto Area^2/t_{ox} \propto s$ this becomes $E \propto s^3$. Fourth, the offset σ^2 scales with $1/s$ as $\sigma^2 \propto t_{ox}/Area^2$ [176]. Put altogether in (11.5), these four factors cancel each other out.

These observations are in line with the results from [181], where an optimum swing is calculated for the case when the receiver would be a ‘linear’ amplifier instead of a latching sense amplifier. Despite the use of a quite different calculation approach and a different technology, a similar optimal swing is found there.

At the optimal swing of 120mV, the equations predict 53fJ/bit energy consumption for the transmitter and interconnect and 22fJ/bit for the sense amplifier. The actual sense amplifier circuit that is simulated in the complete transceiver is scaled for this optimum and consumes 24fJ/bit. This is 10% more than predicted because the minimum width in the technology limits the down-sizing of some transistors and because the actual sense amplifier consists of two interleaved instances which creates a slight power overhead of 1fJ.

11.5 Complete transceiver

11.5.1 Transceiver with synchronization

The previous section discussed the circuits for the data link, but did intentionally not yet mention how the clock is supplied to the receiver, as the data transceiver can operate with many different clocking-schemes, depending on the clocking strategy of the application (the SoC).

In a synchronous NoC, the receiver can simply be clocked with a local copy of the global clock, provided that the link latency does not exceed a clock period. In a completely asynchronous NoC without any clock signals, handshake signals could be used to provide the sense amplifier with a ‘clock’. But for most NoCs, the transceiver clocking strategy that is likely to be most suitable is a source-synchronous scheme in which the transmitter sends a copy of its local clock (or ‘strobe’ or ‘sync’ signal) alongside the data [15-17]. It is a very

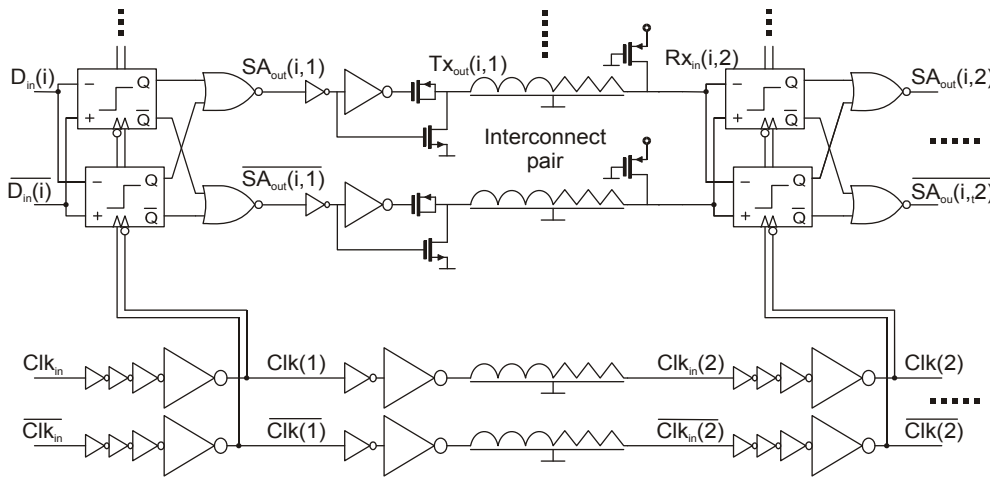


Figure 11.9: Cascade of direct forwarding transceivers.

simple and fast technique that is applicable to both synchronous, mesochronous and GALS systems, as long as each router has a local clock available.

This option will be investigated further in this section and a source-synchronous transceiver schematic is shown in Figure 11.8. At the left side the data words (flits) from the transmitting router enter the transceiver where they are optionally buffered in a transmitter register. The capacitive transmitters transmit the data over the link. Parallel to the data-bus, a gated half-rate clock is also transmitted (or in other words, data transfer is ‘double-pumped’ or at ‘double-data rate’). The sense-amplifier at the receiver consists of two interleaved parts which act on the opposite edges of the clock to enable proper sampling with a half-rate clock. Simple NOR-gates are used to combine the two outputs and create a static output signal.

The clock is transmitted at half-rate because a full-rate clock would be more heavily attenuated by the wire transfer. Full-swing drivers are used for the transmission of the clock to provide as much voltage-swing as possible. Attenuation of the clock can not be compensated by clocked sense amplifiers and conventional amplifiers (cascades of inverters) are used at the receiving end.

The clock is also gated to stop transmission when there is no data (e.g. in between packets). Both halves of the transmitter are also set high during absence of data, to eliminate static current as mentioned earlier. When the clock is stopped, both the halves of the differential clock signals will become low, to signal the receiver that there is no data. When this happens, both halves of the sense amplifier are also reset low, which enables automatic elimination of static current in following transceiver stages in case transceivers are cascaded, as discussed below.

11.5.2 Cascaded transceivers

The synchronizing FIFO that is shown in Figure 11.8 is normally present to re-align the data with the local clock, and is often combined with queues to buffer the incoming data




Figure 11.10: Direct-forwarding transceiver signals at 5Gb/s.

[179]. However, in certain router schemes, one can also omit the re-alignment at intermediate routers and directly forward the data to the next link, thereby greatly reducing the latency of the hops. Direct forwarding - also known as wave-pipelining - can for example be useful in a circuit-switched network [183] where the crossbars that connect the links are pre-configured and there is no need to re-align the data to the local router-clock at each hop, but only at the destination. Source-synchronous transceivers with direct-forwarding can also be interesting for more fine-grained systems that use static routing, such as FPGA's.

To test the concept of direct-forwarding and its wave-pipelined clock, a number of transceivers are cascaded and simulated (omitting the switch fabric for simplicity), as shown in Figure 11.9. Each transceiver in the chain resembles the schematic from Figure 11.8, but without the synchronizing FIFO and with the interleaved sense amplifiers also performing the function of input register. Chains of inverters are used in the clock-path to drive the clock-interconnects. The number of inverters is chosen such that the delay of the clock-path is larger than the delay of the data path: $t_d(\text{Clk1 to Clk2}) > t_d(\text{Clk1 to SA}_{\text{Out}1}) + t_d(\text{SA}_{\text{Out}1 to TX}_{\text{out}1}) + t_d(\text{TX}_{\text{out}1 to RX}_{\text{in}2}) + t_{\text{setup}}(\text{SA})$. The closer these two delays match, the shorter the latency will be, but at the cost of a reduced timing margin.

Some simulated time signals are shown in Figure 11.10. As visible in the figure, the transmission and especially the startup of the clock is in this setup a speed-limiting factor, as the interconnects already cause quite some attenuation of the 2.5GHz clock. At rates higher than 5Gb/s/channel, the accumulation of clock disturbances over multiple stages prevents proper reception during the startup-transient. Simulations with clock-wires in a two times larger metal layer (such that they have four times lower resistance) showed that the entire system is capable to run at 9Gb/s. The purpose of Figure 11.10 is to show that




Figure 11.11: Line output signals of three channels in a twisted bus.

even when the clock wires have to fit in the same area as a single-data channel, it is still possible to reach 5Gb/s.

In the current setup, which uses moderately aggressive timing between data and clock, the latency is 300ps for a single stage (independent of the data rate), so it would cost 1500ps to cross 10mm of interconnect over five stages, which is only slightly larger than the latency of transceivers that use un-interrupted interconnects of 10mm [33, 62] (which were implemented on demonstrator IC's, as discussed earlier).

The energy consumed in a single stage is 129fJ/transition which amounts to 75fJ/bit for random data ($24\text{fJ} + p_{\text{activity}} \times 105\text{fJ}$). In comparison, [15] needs 350fJ/bit to cross 5mm at 1.6GHz, while 2.5 stages from this design can do it for 188fJ/bit. The pseudo-differential low-swing transceiver from [85] needs 1.92pJ/transition to drive a wire that has a capacitance of 1pF, which would correspond to two stages from this design, which need only 256fJ/transition. The transceiver on our second demonstrator IC uses a similar data transceiver which is optimized to cross 10mm of uninterrupted wire. Five stages from this design need 35% more energy per bit, but the multiple stages (clocked repeaters) enable a much higher data rate (5Gb/s versus 2Gb/s) and a higher yield (with respect to offset and PVT variations).

The power consumed in the clock is left out of the comparison above. In this design, the power needed for transmission of the forwarded clock is shared across all the data channels in the bus. The transmission of the clock consumes 1.3pJ/transition when its inverter cascade is loaded by 64 sense amplifiers, which amounts to 20fJ/bit/channel.

The source-synchronous nature of this transceiver helps to make it resilient towards process spread. Simulations with the slow process corner at a temperature of 100°C show an increase in delay of 65ps per stage. At the fast process corner at -25°C, the delay per stage is 45ps lower than in the nominal situation. At both corners, the transceiver chain still operates correctly at 5Gb/s as the change in clock-path delay is equal to the change in data-path delay within 5ps.

The simulations described above were, for simplicity reasons carried out with only one data channel with simple one-dimensional lumped models for the interconnects. To test the effect of crosstalk, a simulation with a bus with twisted interconnects was also carried out. Simulation results of the interconnect outputs of three neighboring channels are shown in Figure 11.11. Hardly any crosstalk is visible in the outputs (compare to the single-ended bus signals in Figure 11.3), which illustrates the effectiveness of the twists.

The fact that part of the wire capacitance is mutual between the wires in the bus does create a change the common-mode transfer of the bus, but the dip in the common-mode visible in Figure 11.11 is a startup transient that does not cause any difficulty for the sense amplifiers

11.6 Conclusions on NoC transceivers

In this chapter, we have shown that the combination of a low-swing capacitive pre-emphasis transmitter, a bus with properly twisted differential wires, a double-tail sense amplifier and a source-synchronous clocking scheme is very suitable for communication in a NoC.

Compared to other low-swing transceivers, the capacitive transceiver does: 1) not need a second supply; 2) can operate at higher speeds; 3) has a higher power efficiency and 4) has a better immunity to supply noise. The capacitively coupled transmitter also makes the transceiver suitable to cross different voltage domains.

The transceiver circuits are compatible with standard digital CMOS circuits and are easily scalable to future technologies. Analysis predicts that the power-optimal swing is about 120mV, also in future technologies

At this swing, the power consumption of the presented differential transmitter is four times lower than the power consumption of a conventional full-swing single-ended transmitter, while the obtainable data rate is 80% higher. When we include the power of the sense amplifier and assume (optimistically) that a full-swing transmitter needs no dedicated receiver, then the presented transceiver is still a factor 3.3 more power efficient. For the 25-tile NoC example with 5Ghz clock and 25% average switching activity, this would mean that the total link power would drop down to 0.8W, instead of the original 2.7W.

With multiple transceiver stages cascaded in a wave-pipelined fashion, the transceiver can also compete with our global-interconnect transceivers as it enables high data rates (5Gb/s versus 3Gb/s in [33] or 2Gb/s in [62]) at a high reliability (6σ for random offset and correct operation over process and temperature corners) and with simple build-in synchronization. As such, the transceiver is also suitable for the long link distances that are for example found in networks with a torus or star topology.

Chapter 12

Conclusions and recommendations

12.1 Conclusions

This project started with the long known premise that interconnects have a scaling problem, in the sense that they become slower when their cross-section is reduced. However, as was discussed in this thesis, this problem can be overcome by a combination of technology, architecture and circuit improvements. On the technology side, reverse scaling combined with a moderate increase in the number of layers for every new technology generation can sustain the increase in bandwidth demand, but only when the wires scale down in length at roughly the same rate as the scaling of the devices (as discussed in section 2.6.3). Techniques that enable a down-scaling of the wire lengths include 3D integration and architectures such as networks on chip. But an increase in the number of interconnects and interconnect layers will result in an increase in interconnect power. And even in new architectures, some global interconnects for data communication will still remain necessary. Therefore, circuit improvements that increase the speed and lower the power for on-chip communication are also needed.

It was shown in Chapter 3 that global interconnects can be characterized with a distributed RC behavior. Even for those interconnects that are so thick or short that skin-effect becomes a dominant factor, the transfer still resembles the transfer of distributed RC. It was also shown that the optimum aggregate data rate is reached when the cross-sectional dimensions and spacings of the interconnects are equal in all directions (assuming the dielectrics are the same in all directions). On-chip transmission line configurations such as microstrips or co-planar waveguides are thus not very beneficial, as they require very wide interconnects [20, 21] which will reduce the bandwidth per area, while skin-effect will limit the performance [19, 45, 46].

For optimum data rate it is also important to use point-to-point interconnects (as was shown in section 3.8.3), which coincides well with the architecture trends to move towards router-based networks instead of multi-drop buses. For point-to-point buses with regular layouts (such as in a mesh based NoC's), it is not difficult to improve crosstalk behavior with proper layouts. As was discussed in section 4.4, a very effective measure to cancel neighbor-to-neighbor crosstalk is to use twisted differential interconnects [82]. The apparent increase in wire resources and power brought by differential signaling is more

Property	Units	[108] Ch'02	[184] Zh'05	[68] Ka'05	[70] Ba'06	[20] Jo'06	[61] Jo'07	[121] Ho'08	[64] Zh'09	[72] Ki'10	[33] Sc'06	[124] Me'10
Technology node	Nm	180	180	130	350	180	180	180	250	90	130	90
Single / Differential		Diff	Single	Single	Single	Diff.	Diff.	Diff.	Diff.	Diff.	Diff.	Diff.
Supply Voltage	V	1.8	1.8	1.2	2.5	1.8	1.8	1.8	2.5	1.2	1.2	1.2
Line length l	mm	20	10	10	17.5	3	14	10*	5	10	10	10
Line width w	μm	2-16	4.5	0.6	2	2-(4+4)	2-8	2-0.3	2-0.4	2-0.6	2-0.4	2-0.54
Line spacing s	μm	2-1*	1*	0.63	1*	2-4	2-8	2-0.3	0.4+2	2-0.4	2-0.4	2-0.32
Metal height h	μm	2	0.5*	0.35	0.5*	0.53	0.53	0.5*	0.5*	0.33*	0.35	0.33
Oxide thickness d	μm	1.9	0.5*	0.36	0.5*	0.5*	0.5*	0.5*	0.5*	0.27*	0.46	0.27
Cross-sectional Area $A_c = (w+s)(h+d)$	μm ²	133	5.5	0.87	3	24.7	33	1.2	3.2	1.2	1.3	1.03
Achieved data rate f_D	Gbps	1	2	0.2	1	8	3	1	2	6	3	2
Energy per bit E_b	pJ/b	16.1	2.3	1.7	5.8	0.29	2	0.84	1.16	0.63	2.0	0.28
Normalized Speed $= f_D \cdot l^2 / A_c$	Gbps·mm ² / μm ²	3.0	36	23	102	2.9	18	53	15.6	500	231	194
Normalized Energy $= E_b / \text{length}$	fJ/b /mm	805	230	170	331	97	143	105	232	63	200	28

* parameter not given in the paper; estimated values based on typical technology data

Table 12.1: Comparison of different transceivers, the last two columns denote this work.

than compensated by the ability to move towards lower signal swings and by the fact that crosstalk shields are no longer needed (not even in multi-layer buses, as long as they are uni-directional).

To be able to predict achievable data rates over interconnects, the notion was used that the dominant error sources that limit the achievable data rate – inter-symbol interference and inter-channel interference (crosstalk) – are inherently deterministic. It is thus possible to analyze their magnitude without having to revert to statistical simulations and methods for this analysis can be based on the symbol response of the interconnects, as was discussed in Chapter 5.

The analysis of line termination and the analysis of different modulation and equalization techniques showed that the simplest concepts have the most merits. On the termination side, of the various types that were investigated, capacitive transmitter termination seems the most promising for many applications. It is simple to implement, does not require much area and simultaneously improves the wire bandwidth and decreases power consumption. On the communication side, plain binary signaling in combination with proper line termination and simple first-order equalization increases the achievable data rate significantly, while the more complex techniques such as multi-level signaling or band-pass signaling showed little benefits. This can also be seen in the table in Appendix B.

That the simple transceiver concepts keep their merits in the implementation was validated with the demonstrator IC's. The first demonstrator IC showed that a combination of PW pre-emphasis with low-ohmic receiver termination can boost the achievable data rate over 10mm of interconnect to 3Gb/s at an energy cost of 2pJ/bit, while a conventional




Figure 12.1: Comparison of different transceivers, based on Table 12.1.

transceiver achieves only 550Mb/s [33]. The second demonstrator IC showed that a capacitive pre-emphasis transmitter in combination with DFE with a continuous-time feedback filter at the receiver can significantly reduce the power consumption down to 0.28pJ/bit, at a data rate of 2Gb/s, again over 10mm interconnect [124]. This second transceiver, but then without the DFE, also forms the basis for the transceiver for network on chip (NoC) applications that was discussed in Chapter 11. Simulations showed that five repeated sections of the NoC transceiver should be able to achieve to 5Gb/s over 10mm of interconnect at a power consumption of 375fJ/bit [180]. When the clock-wires for the source synchronous repeated transmission in this transceiver would be routed in a thicker metal layer, then the data rate would go up further to 9Gb/s.

A comparison of our measured results with other on-chip communication transceivers is given in Table 12.1 and shown in Figure 12.1. The comparison in the figure is made based on a normalized energy per bit per length versus a normalized data rate per cross-section times distance squared. This was done earlier in [2, 124]. Here the results are updated with two recent publications [64, 72] (and the 50mV swing results from [121] are now used, which gives a better efficiency than at 200mV).

Our earlier work with the pulse-width pre-emphasis is on the right side of the graph, only surpassed in normalized data rate by the recent work from Kim et. al. [72]. At the bottom of the graph is our capacitive pre-emphasis transceiver, which achieves the best power efficiency at a small expense in data rate.

The table and the figure only show results verified on silicon. When we would also include the simulation results for the NoC transceiver then that transceiver would end up at the

favorable right bottom side of the graph with a normalized data rate of 485 Gbps-mm²/μm² and a normalized energy of 38fJ/b/mm. Except for its high power efficiency and high speeds, the NoC transceiver is also dimensioned for high yield (6σ for random offset) and simple application (without parameters that need to be calibrated or matched to the wire). Also, the transceiver is well suited for modern systems on chip that can have multiple voltage domains and can have different clock phases for different parts of the chip, the so-called ‘globally asynchronous, locally synchronous’ (GALS) systems. The differential transceiver can bridge voltage domains and the source-synchronous scheme facilitates simple re-clocking at the receiver. The transceiver clock can also be stopped to enter a power-down mode without static power consumption (from which the transceiver can wake up within one clock cycle).

Of the presented transceivers, the NoC transceiver thus seems best suited for future application in commercial mass-produced IC’s. As this was a research project, the steps towards industrialization are left to other parties. We hope we have paved the way and made our contribution to the field by introducing some high-speed, power efficient, reliable and easy to use transceiver concepts.

12.2 Original contributions

This section discuss the original contributions in this thesis. The contributions are split into three lists, corresponding to the three major parts of the thesis.

Original contributions for interconnect characterization and optimization

- In section 3.6 it was shown with the interconnect transfer analysis that skin-effect has a very similar impact on the transfer function as distributed RC behavior (they are both described by the diffusion equation). The practical implication is that we can use a distributed RC-model to describe the interconnect behavior without much modeling error, regardless of length or thickness.
- In section 4.2.4, it was shown that receiver termination of an interconnect with an RL series can boost the achievable data rate even further than with the low-ohmic receiver or capacitive transmitter termination (where the latter was an original contribution from this project [2]). RL termination is however not as power efficient as capacitive transmitter termination, nor very straightforward to implement, and was thus not implemented on silicon.
- Crosstalk can be mitigated by placing one or two twists at the correct position, as was shown in [2]. In section 4.4.7 in this thesis, it was shown how more twists can be applied to also cancel crosstalk in multi-layer buses.
- The power model from section 4.5.2 is able to accurately predict interconnect power for statistical data signals. It is especially suitable for those situations where the symbol times are shorter than the wire time constant (where the classical CV² model is no longer sufficient).

Original contributions for communication methods and analysis

- In Chapter 5, a method was introduced to predict eye diagram properties for various communication schemes, based on symbol response analysis. The method has some

resemblance to the peak-distortion analysis [95] and to the model used to analyze binary signaling in [96], with the difference that the model that is presented here is more generally applicable (and can for example also predict the effect of crosstalk).

- In Chapter 6 and Chapter 7, the model was applied to predict achievable rates for on-chip channels with various signaling schemes. Closed-form analytical solutions were presented to predict the achievable data rate with first-order channels and first-order equalization.
- Pulse-width pre-emphasis, as discussed in section 7.4, was introduced and subsequently developed further together with J.R. Schrader [44].
- Section 7.6 discussed decision feedback equalization with a continuous-time analog filter. To the best of the authors knowledge, this is original work, but it is quite conceivable that it is used somewhere else. It has recently also been applied in I/O transceivers by another research group [139, 185].

Original contributions for on-chip transceiver circuits

- The ‘double-tail’ sense amplifier, as discussed in Chapter 9, was introduced to improve data detection with small signal swings at high common-mode voltage levels. Also, methods to estimate a sense amplifier’s noise or offset standard deviation are presented in Appendix A— including derivations for the tolerance in this estimate.
- It was shown in Chapter 10 that DFE with an analog filter can be combined with a sense-amplifier with only a marginal increase in power consumption (20fJ/bit at 2GHz).
- It was shown in Chapter 11 that capacitive transmitters and double-tail sense amplifiers are very good candidates to create low-swing transceivers for networks on chips (with lower power consumption than low-swing transmitters that use a dedicated supply) and it was shown what the optimal swing is for such a transceiver. Also, it was shown that these transceivers can be used with source synchronous clock distribution and optionally with direct forwarding (wave pipelining).

12.3 Recommendations for further study

In this project, different equalization techniques have mostly been treated separately (except for the combination of equalization with line termination techniques). This is also because we have limited ourselves to simple transceivers, using equalization at one side and line termination at the other side of the wire. It was even argued that it is a good idea to further simplify the transceiver and use simple capacitive transmitters and sense amplifiers at regular intervals along the wire, as was done for the NoC transceivers.

However, for those wires that have to span the entire chip, have stringent latency requirements (so they can not be broken into multiple segments), and can not be moved to metal layers with large cross-sections (to increase their bandwidth), more elaborate transceiver schemes can be useful. To this end, we have shown that significant bandwidth improvements can be achieved with FIR or PW pre-emphasis or with DFE.

For even higher data rates, a combination of pre-emphasis and DFE can be used. This has recently been demonstrated by Kim et. al. [72], as was also mentioned in the conclusions.

They showed that a combination of FIR pre-emphasis and a resistively terminated receiver with a one-tap (unrolled) DFE can achieve data rates up to 6Gb/s over 10mm of interconnect with 0.63pJ/b consumption (using a re-arranged current-summing transmitter, that avoids any local currents that cancel in the sum).

Because a resistively terminated receiver has a similar bandwidth as a capacitively terminated transmitter, it can be assumed that similar data rates can be achieved when a capacitive transmitter with FIR pre-emphasis [74, 78, 121] is used instead of the current-switching transmitter, with the advantage that the power consumption decreases. Whether this is indeed the case could be a topic for further experimentation.

When this combination of a capacitive transmitter, pre-emphasis and DFE still not meets the data rate demand, then one can consider a variant of the system from Kim et. al. [72], but then with resistive-inductive (RL) receiver termination instead of resistive termination. In section 4.2.4 it was shown that this type of termination should give another factor three bandwidth increase, at the cost of a more complex receiver (with more power costs than for a capacitive transmitter, which is why it was never implemented on silicon in this project).

But, the higher the desired data rate, the more it also becomes important to match the time constants of the termination (in case of RL-termination) and of the equalization to the time constants of the channel. Possible techniques to do this adaptively were already discussed in section 7.7.2, but not yet implemented in this project.

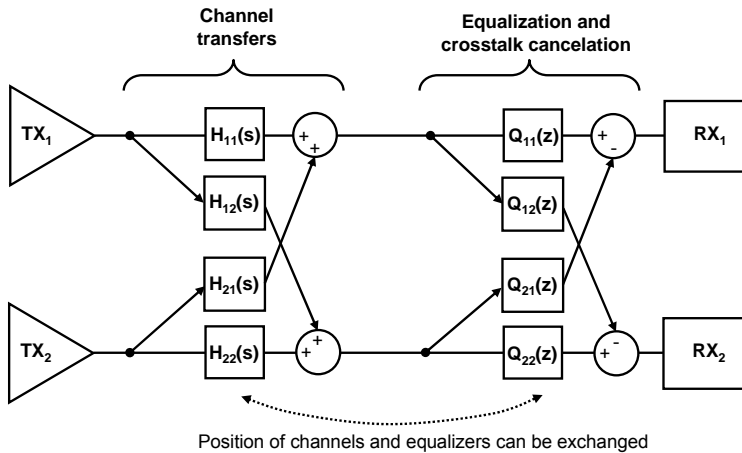


Figure 12.2: Possible method for crosstalk cancellation through equalization

12.3.1 Recommendations on side-topics

Apart from the central theme of this thesis, for which recommendations for further research are mentioned above, a few side-topics were also investigated in this project. The ones that deserve further attention are listed below.

Crosstalk cancellation through equalization

Next to the differential twisted wires, it was also briefly investigated if crosstalk cancellation through equalization could be used to mitigate crosstalk without the need for differential wires. The initial analysis with a MIMO system, as shown in Figure 7.1, indicated that the crosstalk cancellation filters (Q_{12} or Q_{21}) combine well with the normal equalization filters (Q_{11} or Q_{22}). When there would be no equalization ($Q_{11}=Q_{22}=1$), then the crosstalk filters Q_{12} and Q_{21} should be high-pass filters with their dominant pole close to the dominant pole of the channel (also see channel crosstalk transfer in Figure 4.10). But when equalization is used, then this pole is cancelled by the equalizing filters Q_{11} or Q_{22} , so it is also not needed for the crosstalk filter. This simplifies (reduces the order needed for) the crosstalk filter. This topic was not further investigated in this project, as the twisted differential wires were effective to cancel crosstalk and enabled robust low-swing signaling.

In the backplane communication literature, good results were obtained with crosstalk cancellation at the transmitter side (with the goal to reduce crosstalk induced jitter) [186, 187]. In on-chip communication, a combination of bus encoding and a simple crosstalk equalizer (similar to DFE) is presented in [142]. But, it seems that there is still room for improvement, with equalizers and crosstalk filters that are better matched to the interconnect transfer, which could be an alternative to differential twisted wires in those cases where minimum interconnect area is of prime importance.




Figure 12.3: Spatial CDMA with 3 channels over 4 wires, showing channel 1 in (a), channel 2 in (b) and channel 3 in (c).

Spatial CDMA

Another topic that was shortly investigated was ‘spatial CDMA’. Here the intention is to reduce the wire overhead and still be able to use differential receivers and benefit from common-mode noise rejection. The concept is to superimpose a set of $N-1$ DC balanced signals over N wires in such a way that the signals are orthogonal to each other and can be disentangled at the receiver side. Hadamard matrices (see section 4.4.7) can be used to define the polarities for the signals, similar to CDMA signaling (as was used in section 2.5.3) but then with the code spread over different wires instead of over time, hence the term spatial CDMA. Figure 12.3 shows an example where the 2nd to last row of a H^4 matrix are used to define the polarities for channel 1 to 3 (the first row of the Hadamard matrix can not be used, as it is not DC balanced). Feasibility of the concept was shown with circuit simulations (by Thomas Schaink).

As the attention shifted to other topics, a few questions about spatial CDMA remain unanswered. One question is which type of transmitter is most suitable. Current summing transmitters, as shown in Figure 12.3, are simple (and their power efficiency can be increased with techniques as used in [72]). But perhaps capacitive transmitters are usable as well, to enable higher bandwidths and reduce power consumption. The effect of ISI and crosstalk on spatial CDMA performance is another topic that deserves further attention. One effect of crosstalk is that it will partly destroy the orthogonality between the channels. It seems likely however that this form of crosstalk can be pre-compensated by tuning the transmitter strengths (by solving e.g. the matrix H^{TX} from $H^{TX} * H^{Xtalk} = H^4$). But how ISI and crosstalk from previously transmitted symbols can be compensated was not yet investigated.

In the literature, some similar topics have been investigated. In [188], a backplane transceiver is for example presented, where three DC balanced data signals are transmitted over four wires, but the signals are encoded a bit differently, using the so-called ‘phantom line’ approach that was developed earlier for telephone systems. Incremental signaling [189], is another method to transmit $N-1$ signals over N interconnects, where each subsequent channel uses the previous channel as its reference. How these techniques compare to spatial CDMA could be a subject of further investigation.

Partial response signaling

Partial response signaling is a filtering method that operates on the data-stream (most often with very simple filters such as $1+z^{-1}$, which is called duo-binary), to modify the spectral shape of the data and reduce ISI [95].

In high-speed transceivers, it is sometimes applied in combination with other equalization methods. In [190] for example, a duo-binary partial response transmitter is used in a 12Gb/s backplane transceiver.

For on-chip transceivers, methods that can be classified as partial response signaling have also been proposed. An extension to the capacitive transmitter with FIR pre-emphasis is for example presented in [78], where it is combined with an AC-coupled receiver to create pulsed return to zero (RZ) signals (similar to [77]). Pulsed RZ signaling is a form of partial response signaling, similar to modified duo-binary [95].

Partial response signaling was only briefly looked at in this project, but at first sight, it did not seem that the decrease in ISI would outweigh the drawback for the detector: instead of one detector threshold, multiple thresholds are needed to detect partial response signaling (e.g. a comparator with hysteresis for the case in [78]), which reduces the noise margin compared to a single detector transceiver. But more detailed analysis is needed to quantitatively assess possible drawbacks or merits.

Shannon bound

In the project, we briefly used Shannon’s channel capacity theorem [95] to estimate the boundaries for the achievable data rates over on-chip channels. This was done under the assumption that the signal to noise ratio (SNR) at the receiver is proportional to the channel magnitude transfer (effectively assuming that the transmitted power is flat over frequencies and that there is a flat noise density N_0 at the receiver). The capacity formula was subsequently applied piecewise over small frequency regions (from f_i to f_{i+1}), assuming that the channel transfer is flat within the region:

$$C = BW \log_2 \left(1 + \frac{P_{sig}}{BW \cdot N_0} \right) \rightarrow C_i \approx (f_{i+1} - f_i) \log_2 \left(1 + SNR \frac{H(f_i) + H(f_{i+1})}{2(f_{i+1} - f_i)} \right) \quad (12.1)$$

The results were numerically evaluated up to the point where the capacity $C = \sum(C_i)$ no longer increased with increasing frequency. This was applied to the 10mm interconnect with low-ohmic termination as found on the first demonstrator IC. With an assumption of 100mV noise at the receiver ($SNR = P_{sig}/P_{noise} = 21.6\text{dB}$ with $P_{sig} = (1.2\text{V})^2$, $P_{noise} = (0.1\text{V})^2$) the resulting channel capacity was estimated to be 6Gb/s. This is only a factor two higher than what was achieved with the pulse-width pre-emphasis (see Chapter 8).

The difficulty with the estimate is the assumption in the SNR because of the uncertainty in the receiver noise floor. When the SNR is for example increased to 40dB, then the capacity rises to 20Gb/s. Perhaps it is possible to find a receiver noise density curve with more physical meaning than the 100mV noise floor that was assumed above, but this was not done in this project.

List of publications

1. D. Schinkel, R. P. de Boer, A. J. Annema and A. J. M. van Tuijl, "A 1-V 15 μ W High-Precision Temperature Switch," *Proc. 27th European Solid-State Circuits Conf.*, pp. 104-107, Sept. 2001.
2. D. Schinkel, A.J.M. van Tuijl and A. J. Annema, "Reducing quantization noise with recursive sigma-delta modulators," *Proc. IEEE Int. Symp. On Circuits and Systems*, pp. I-1084-1087, May 2004.
3. D. Schinkel, R. P. de Boer, A. J. Annema and A. J. M. van Tuijl, "A 1-V 15 μ W High-Accuracy Temperature Switch," *Kluwer Int. J. Analog Int. Circ. Sig. Processing*, Vol. 41, pp. 13-20, Oct. 2004.
4. D. Schinkel, E. Mensink, E. A. M. Klumperink, A. J. M. van Tuijl, B. Nauta, "A 3Gb/s/ch Transceiver for RC-limited On-Chip Interconnects," *IEEE ISSCC Dig. Tech. Papers*, pp. 386-387, Feb. 2005.
5. E. Mensink, D. Schinkel, E.A.M. Klumperink, A.J.M. van Tuijl, B. Nauta, "Optimally-Placed Twists in Global On-Chip Differential Interconnects," *Proc. 31th ESSCIRC*, pp. 475-478, Sept. 2005.
6. T. S. Doorn, A.J.M. van Tuijl, D. Schinkel, A.J. Annema, M. Berkhout, B. Nauta, "An audio FIR-DAC in a BCD process for high power Class-D amplifiers," *Proc. 31th ESSCIRC*, pp. 459-462, Sept. 2005.
7. D. Schinkel, E. Mensink, et. al., "A 3-Gb/s/ch Transceiver for 10-mm Uninterrupted RC-limited Global On-Chip Interconnects," *IEEE Journal of Solid-State Circuits*, Vol. 41, pp. 297- 306, Jan. 2006.
8. D. Schinkel, E. Mensink, et. al., "Double-Tail Latch-Type Voltage Sense Amplifier With 18ps Setup+Hold Time," *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, pp. 314-315, Feb. 2007.
9. E. Mensink, D. Schinkel, et. al., "A 0.28pJ/b 2Gb/s/ch transceiver in 90nm CMOS for 10mm On-Chip Interconnects," *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, pp. 414-415, Feb. 2007.
10. E. Mensink, D. Schinkel, et. al. "Optimal Positions of Twists in Global On-Chip Differential Interconnects," *IEEE Trans. on VLSI Sytems*, pp. 438-446, April 2007.
11. M. van Elzakker, A.J.M. van Tuijl, P.F.J. Geraedts, D. Schinkel, E.A.M. Klumperink, B. Nauta, "A 4.4 fJ/conversion-step, 10 bit, 1 MS/s charge redistribution ADC with 1.9uW power," *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, pp. 244-245+610, Feb. 2008.

12. D. Schinkel, E. Mensink, E. A. M. Klumperink, A. J. M. van Tuijl, B. Nauta, "Low-Power, High-Speed Transceivers for Network-on-Chip Communication," *IEEE Trans. on VLSI Systems*, Vol. 17, nr 1, pp 12-21, Jan. 2009.
13. E. Mensink, D. Schinkel, E. A. M. Klumperink, A. J. M. van Tuijl, B. Nauta, "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, Feb. 2010
14. M. van Elzakker, A.J.M. van Tuijl, P.F.J. Geraedts, D. Schinkel, E.A.M. Klumperink, B. Nauta, "A 10-bit Charge-Redistribution ADC consuming 1.9 μ W at 1MS/s", *IEEE Journal of Solid-State Circuits*, Vol.45, no.5, pp.1007-1015, May 2010

Next to the list above, a number of publications have been presented at the *Workshop on Circuits, Systems and Signal Processing, ProRisc* in Eindhoven, The Netherlands:

15. D. Schinkel, R. P. de Boer, A. J. Annema and A. J. M. van Tuijl, "A 1-V 15 μ W High-Precision Temperature Switch," *Proc. of the 13th ProRISC*, Nov. 2002
16. E. Mensink, D. Schinkel, et. al., "Interconnects and On-Chip Data Communication Techniques," *Proc. of the 15th ProRISC*, pp. 556-561, Nov.2004
17. D. Schinkel, A.J. M. van Tuijl, A. J. Annema, "Using Recursive Multibit Sigma-Delta Modulators to Reduce the Quantization Noise Power," *Proc. of the 15th ProRISC*, pp. 589-593, Nov. 2004
18. D. Schinkel, E. Mensink, E. A. M. Klumperink, A. J. M. van Tuijl and B. Nauta, "A Transceiver for High-Speed Global On-Chip Data Communication," *Proc. of the 16th ProRISC*, pp. 279-283, Nov. 2005
19. E. Mensink, D. Schinkel, E. A. M. Klumperink, A. J. M. van Tuijl and B. Nauta, "Global On-Chip Differential Interconnects with Optimally-Placed Twists", *Proc. of the 16th ProRISC*, pp. 253-258, Nov. 2005
20. E. Mensink, D. Schinkel, E. A. M. Klumperink, A. J. M. van Tuijl and B. Nauta, "A Power Efficient 2Gb/s Transceiver in 90nm CMOS for 10mm On-Chip Interconnect," *Proc. of the 18th ProRISC*, pp. 60-63, Nov. 2007
21. D. Schinkel, E. Mensink, E. A. M. Klumperink, A. J. M. van Tuijl and B. Nauta, "A Low-Offset Double-Tail Latch-Type Voltage Sense Amplifier," *Proc. of the 18th ProRISC*, pp. 89-94, Nov. 2007

Patents

1. D. Schinkel and P. A. C. M. Nuijten, "Data Converter," European patent EP1556953, July 2005; US patent US7,034,726, April 2006.
2. D. Schinkel, A. J. M van Tuijl and P. A. C. M. Nuijten, "Volume Control Device For Digital Signals," Application nrs. WO2004086615, US2006182186, European patent EP1611678, Jan 2006.

About the author



Daniël Schinkel (S'03-M'08) was born in Finsterwolde, the Netherlands, in 1978. He received the M.Sc. degree in electrical engineering (with honors) from the University of Twente, the Netherlands, in 2003. During his studies he worked on various occasions as a trainee at the Mixed-Signal Circuits and Systems Department of Philips Research, Eindhoven the Netherlands. This work resulted in a number of publications and two patents.

From 2003 to 2007 he worked as a PhD student at the same university, at the IC-design group headed by Bram Nauta. The subject of his PhD work was high-speed on-chip communication. In the same period, he also occasionally worked as a freelance consultant on the subject of sigma-delta converters.

He is one of the founders of Axiom IC, an IC-design company that started in 2007 and focuses on the design of state-of-the-art analog and mixed signal circuits.

His technical interests include analog and mixed-signal circuit design, sigma-delta data converters, class-D power amplifiers and high-speed communication circuits. He holds two patents and has authored or co-authored about 20 papers.

Appendix A

Standard deviation estimation in comparators

This appendix describes methods to estimate the offset or noise standard deviation (σ) in comparator circuits. Extraction of this standard deviation from simulations or measurements is not as simple as with normal linear circuits. Shorted comparator inputs will for example not produce an output signal from which the input equivalent noise or offset can be estimated. Instead, a small input signal has to be applied.

When we assume that we have a comparator that has thermal noise or offset with a Gaussian distribution (and ensure that hysteresis is excluded from the analysis, as will be discussed below), then we can estimate the standard deviation. This can be done by applying a small input voltage V_{in} and observe the average number of correct comparator decisions as a ratio to the total number of decisions. The standard deviation follows from this ratio (p) together with the inverse of the cumulative Gaussian distribution Q :

$$\sigma = \frac{V_{in}}{Q_{inverse}(p)} \quad (\text{A.1})$$

Not all mathematical computer tools have the inverse of the Gaussian cumulative distribution (also called the 'normal quantile' or 'probit' function) available. A more commonly available function is the inverse of the complementary error-function (e.g. `erfcinv` in Matlab), which is a scaled version of $Q_{inverse}$:

$$\sigma = \frac{V_{in}}{-\sqrt{2} \cdot \text{erfcinv}(2p)} \quad (\text{A.2})$$

A.1 Accuracy of standard deviation estimation

The offset standard deviation can be derived from a number of monte-carlo simulation trials (as was done in section 9.5) or from a number of measured samples (as was done in section 9.6). The rms-noise can be derived by measuring the average of a number of comparator decisions from a transient noise simulation or a measurement. In both cases, it is important that hysteresis effects do not influence the outcome. This can be ensured when the initial state of the comparator prior to the measurement is always the same. The prior state can either be a well defined decision or a perfect equilibrium (possible in simulations). In the

former case, one has to know the amount of hysteresis, and remove it from the equations, by taking:

$$V_{in} = V_{in}' - V_{hysteresis} \quad (\text{A.3})$$

A question that arises is how accurate the estimated standard deviation actually is and how the tolerances in the estimate can be minimized. This is analyzed below.

When a number of trials are taken, then they should on average produce a number of positive decisions k equal to the trial-count n multiplied by the probability of a positive decision p . From the observed k , the probability p can be estimated: $p_{est} = k/n$ and equation (A.1) or (A.2) can then be used to calculate the corresponding standard deviation of the offset (or noise). However, the number of positive decisions in the trials is in itself also a random variable with a certain variance (the sample variance [91]).

When we assume that each sample is an independent decision (so no effects of e.g. hysteresis), then the trials can be modeled as Bernoulli trials which have a Binomial distribution [91, 92]. The mean of the Binomial distribution is np (the average value of k) and its variance is [92]:

$$\text{var}(k) = np(1 - p) \quad (\text{A.4})$$

So the variance in k depends on the probability p , which is a function of the input voltage through the inverse of (A.1) or (A.2):

$$p = Q\left(\frac{V_{in}}{\sigma}\right) = \frac{1}{2} \text{erfc}\left(-\frac{V_{in}}{\sigma\sqrt{2}}\right) \quad (\text{A.5})$$

Thus, (A.4) can be rewritten as:

$$\text{var}(k) = n \left(\frac{1}{2} \text{erfc}\left(-\frac{V_{in}}{\sigma\sqrt{2}}\right) - \frac{1}{4} \text{erfc}\left(-\frac{V_{in}}{\sigma\sqrt{2}}\right)^2 \right) \quad (\text{A.6})$$

To determine how the variance in k translates to a variance in the estimated standard deviation, we can approximate equation (A.1) or (A.2) by a linear function, which should give accurate results when the variance in the estimated sigma is small:

$$\sigma_{est} = f(p_{est}) \approx a \cdot p_{est} + b \quad (\text{A.7})$$

For a linear function, the variance of the output has a simple relation to the variance of the input:

$$\text{var}(ap_{est} + b) = a^2 \text{var}(p_{est}) \quad (\text{A.8})$$

First-order Taylor expansion of (A.2) can be used to determine a . According to [191], the derivative of an *erfcinv* is:

$$\frac{d \text{erfcinv}(x)}{dx} = -\frac{1}{2} \sqrt{\pi} e^{\text{erfcinv}(x)^2} \quad (\text{A.9})$$

So with some mathematical derivations it follows from (A.2) and (A.9) that:




Figure A.1: Tolerance in the estimated σ as a function of the normalized comparator input voltage, for three different number of trials.

$$a = \frac{d \sigma_{est}(p_{est})}{d p_{est}} = \frac{-V_{in} \sqrt{\pi} e^{\text{erfcinv}(2p)^2}}{\sqrt{2} \text{erfcinv}(2p)^2} \quad (\text{A.10})$$

With (A.5), this can be rewritten as:

$$a = \frac{d \sigma_{est}(p_{est})}{d p_{est}} = -\sqrt{\pi} e^{\frac{V_{in}^2}{2\sigma^2}} \frac{\sqrt{2}\sigma^2}{V_{in}} \quad (\text{A.11})$$

The variance in p_{est} itself follows from (A.6):

$$\text{var}(p_{est}) = \frac{\text{var}(k)}{n^2} = \frac{1}{n} \left(\frac{1}{2} \text{erfc} \left(-\frac{V_{in}}{\sigma\sqrt{2}} \right) - \frac{1}{4} \text{erfc} \left(-\frac{V_{in}}{\sigma\sqrt{2}} \right)^2 \right) \quad (\text{A.12})$$

By combining (A.7) to (A.12), the variance in the estimated sigma can be expressed as:

$$\text{var}(\sigma_{est}) \approx \frac{\pi}{n} \left(\text{erfc} \left(-\frac{V_{in}}{\sigma\sqrt{2}} \right) - \frac{1}{2} \text{erfc} \left(-\frac{V_{in}}{\sigma\sqrt{2}} \right)^2 \right) \left(\frac{\sigma}{V_{in}} \right)^2 e^{\frac{V_{in}^2}{\sigma^2}} \sigma^2 \quad (\text{A.13})$$

Not surprisingly, the variance in the estimate is proportional to σ^2 and inversely proportional to n (in other words, the accuracy of the estimate increases proportionally to the square root of n).

In Figure A.1, the result of (A.13) is plotted for three different sample sizes n . Instead of $\text{var}(\sigma_{est})$ the square-root of the variance (the deviation) is plotted, divided by σ itself, as a measure of the relative tolerance in the estimate. The results show that the optimum choice of input signal is $V_{in}=1.58\sigma$. At this optimum input signal, the tolerance in the estimated sigma is 4% with 1000 trials, 13% with 100 trials and 40% with 10 trials.

These results were verified with simulations (that numerically evaluate (A.2) with many sets of random samples). At $V_{in}=1.58\cdot\sigma$ and $n=1000$, the simulations also show a tolerance of 4%. With the optimal input voltage of σ_{est} , the shape of the histogram of resembles a Gaussian function. When the input is changed to lower or higher values then the accuracy drops and the shape of the histogram also becomes more asymmetrical (with an elongated tail at the right or at the left side respectively).

At $n=100$, the simulations give a tolerance of 14%, slightly larger than the predicted 13%. At $n=10$, the simulations give a tolerance of 79%, much higher than the predicted 40%. This is because the linear approximation from (A.7) is no longer accurate enough for such small sample sizes. Interestingly, the simulations also show that the mean of σ_{est} becomes much lower than the actual σ when $n=10$ (as in some cases, the samples will incorrectly predict zero σ_{est} when there are 5 positive and 5 negative decisions). Ten comparator decisions are simply not enough to get an accurate estimate, which is not really surprising. For most situations, it should be no problem to obtain much more comparator decisions, e.g. the 1000 that were used for the circuit simulations in section 9.5. When small simulation times are desired, then hundred trials might be considered. For $n=100$, the simulations show that the V_{in} that gives the smallest variance drops slightly to $V_{in}=1.4\cdot\sigma$, at which the tolerance is 13.6%.

So, the results above can be used to set the optimal V_{in} for standard deviation estimation, and provide the tolerance in the result. Before the actual trials, the standard deviation σ has to be guessed to set the optimal V_{in} , or a short initial test can be done to get a rough indication. Fortunately, the optimum is quite shallow and a V_{in} between 1 to 2 times σ is quite adequate (at least for proper sample sizes of e.g. $n=1000$).

A.2 Decision averaging versus impedance scaling

There is another topic on which statistical analysis as used above, can shed some light, which is the difference between scaling a comparator or averaging over multiple decisions. The classical approach to reduce the offset or noise of a comparator (or other circuits) is impedance scaling. One can scale the circuit by a factor N which reduces the offset and noise sigma by a factor \sqrt{N} , at the cost of an N -fold increase in area and power per conversion.

Alternatively, one can use N original comparators and average the outcome which should also decrease noise and offset. This method also increases power per effective conversion by a factor N , but can be more flexible as one can adapt the number of averaged decisions depending on noise or offset requirements (e.g. depending on the data rate in transceiver applications). When one is only interested in noise reduction and not in offset (e.g. because offset is reduced with calibration), then one could use only one (original) comparator and repeatedly sample the same input and average the outcome, which would also save area.

The interesting question is how much this averaging approach reduces offset and/or noise and how it compares to impedance scaling. Given that the binomial distribution predicts the outcomes of n decisions, as was discussed in section A.1, we are interested in the chance that the majority vote is in the correct direction, given a certain probability per decision p . This can be solved with the cumulative binomial distribution function F :




Figure A.2: Difference between scaling a comparator and averaging over multiple decisions.

$$P(k > \text{floor}(\frac{1}{2}n)) = 1 - F(\text{floor}(\frac{1}{2}n)) \quad (\text{A.14})$$

And F is given by [92]:

$$F(l) = \sum_{k=0}^l P(k) \quad (\text{A.15})$$

$$P(k) = \binom{n}{k} p^k (1-p)^{n-k}$$

Note that when n is even, then the outcome $P(k = \text{floor}(n/2))$ is a tie, which is in the formula above not counted as a positive result. In an application, to avoid ties, it makes sense to only average over odd numbers of decisions.

We can compare the results of this formula to the probability of a correct decision with only one trial, which is equal to the cumulative Gaussian distribution from (A.5). This is done in Figure A.2. The majority of 5 decisions is compared to one decision from both the ‘original’ and a 5-times larger comparator (which has an $\sqrt{5}$ lower σ).

The results show an interesting, but not entirely unexpected result, namely that it is more efficient (in terms of power versus probability for a correct decision) to use N parallel comparators, instead of averaging over N decisions. The shape of the two probability-curves is still similar (although not perfectly equal), but the variance of the majority of N -decision is higher, which can be attributed to the notion that information is lost when random signals (noise) are averaged after binary quantization instead of before.

Additional evaluation of the equations shows that an average of 5 decisions yields roughly the same probabilities as scaling a comparator size by 3.5 (at least in the region

of $\sigma < V_{in} < \sigma$). This makes the latter option a factor 1.43 more power efficient. This factor rises a bit for larger n , to for example 1.54 for $n=21$.

The difference is not very large, but is sufficient to conclude that it will usually be better to use impedance scaling instead of decision averaging, except for those applications where the increased flexibility of the averaging approach outweighs its disadvantage in power efficiency. In this project impedance scaling was thus used to obtain low offset.

Appendix B

Overview of achievable data rates

This appendix gives an overview of the achievable data rates as analyzed in Chapter 6 and Chapter 7. The data rate is normalized to the RC product of a single wire to enable easy reuse. For reference note that the measured RC products in this project are:

$R_{\text{wire}}C_{\text{wire}} = 3.8\text{ns}$ for a 10mm 0.4 μm wide M5 wire in 130nm CMOS and

$R_{\text{wire}}C_{\text{wire}} = 4.8\text{ns}$ for a 10mm 0.54 μm wide M4 wire in 90nm CMOS

Wire type		Signaling type	Achievable data rate (bit/ $R_{\text{wire}}C_{\text{wire}}$)			
			$R_s=0,$ $R_L=\infty$	$R_s=0,$ $R_L=R_{\text{wire}}/10$	$C_s=C_{\text{wire}}/10$ $C_L=0$	
single-ended	Unshielded <i>Xtalk</i> from 2 neighbors	plain binary	1.8	4.6	4.0	
	Unshielded <i>Xtalk</i> from 4 neighbors		1.7	4.4	3.3	
	Shielded		3.1	8.8	8.8	
Twisted differential	2.4		6.8	7.2 / 7.0 ¹⁾		
Shielded single-ended (scale by 1/1.26-1/1.29 for twisted differential)		4-PAM	3.2 ²⁾			
		8-PAM	3.3 ²⁾			
		Band -pass	2-PAM	~ 9 ²⁾		
			4-PSK	~ 6 ²⁾		
		CDMA	$\ll 2$			
		binary + EQ	FIR	18.1	26.1	26.1
			PW	22.8	32.7	32.7
			DFE	25.5	37	37
		4-PAM + FIR EQ	14.6	22.2	22.2	
		4-PAM + PW EQ	25.5 ²⁾	37.0 ²⁾	37.0 ²⁾	
4-PAM + DFE	27.2 ²⁾	39.7 ²⁾	39.7 ²⁾			

1) $C_s=1.25*0.1C_{\text{wire}}$ to compensate for the increase in effective wire capacitance

2) Benefit over binary signaling only at very small eye-openings

3) Limit with only 1% (10mV) absolute eye-opening left (because there is no zero crossing for theoretical limit).

Table B.1: Overview of achievable data rates for various wire configurations and signaling types.

References

- [1] R. H. Dennard, F. H. Gaensslen, V. L. Rideout, E. Bassous, and A. R. LeBlanc, "Design of ion-implanted MOSFET's with very small physical dimensions," *Solid-State Circuits, IEEE Journal of*, vol. 9, pp. 256-268, Oct 1974.
- [2] E. Mensink, "High-Speed Global On-Chip Interconnects and Transceivers," University of Twente, The Netherlands, PhD Thesis, Enschede, 2007.
- [3] ITRS, "International Technology Roadmap for Semiconductors," Edition 2009 and updates 2010.
- [4] W. S. Song and L. A. Glasser, "Power distribution techniques for VLSI circuits," *Solid-State Circuits, IEEE Journal of*, vol. 21, pp. 150-156, Feb. 1986.
- [5] P. Zarkesh-Ha and J. D. Meindl, "Optimum on-chip power distribution networks for gigascale integration (GSI)," *Interconnect Technology Conference, 2001. Proceedings of the IEEE 2001 International*, pp. 125-127, June 2001.
- [6] J. Nurmi, H. Tenhunen, J. Isoaho, and A. Jantsch, *Interconnect-Centric Design for Advanced SoC and NoC*: Kluwer Academic Publishers, 2004.
- [7] H. Bakoglu, *Circuits, Interconnections and Packaging for VLSI*: Reading, MA: Addison-Wesley, 1990.
- [8] R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," *Proceedings of the IEEE*, vol. 89, pp. 490-504, April 2001.
- [9] J. A. Davis, V. K. De, and J. D. Meindl, "A stochastic wire-length distribution for gigascale integration (GSI). I. Derivation and validation," *Electron Devices, IEEE Transactions on*, vol. 45, pp. 580-589, March 1998.
- [10] T. N. Theis, "The future of interconnection technology," *IBM Journal of Research and Development*, vol. 44, May 2000.
- [11] B. S. Landman and R. L. Russo, "On a pin versus block relationship for partitions of logic paths," *IEEE Transactions on Computers*, vol. C-20, pp. 1469-1479, Dec. 1971.
- [12] D. Pham, S. Asano, M. Bolliger, M. N. Day, H. P. Hofstee, et al., "The design and implementation of a first-generation CELL processor," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 184-185, Feb. 2005.
- [13] L. Benini and G. De Micheli, "Networks on chips: a new SoC paradigm," *IEEE Computer*, vol. 35, pp. 70-78, Jan. 2002.
- [14] W. J. Dally and B. Towles, "Route packets, not wires: on-chip interconnection networks," *Proc. 38th Design Automation Conf.*, pp. 684-689, June 2001.

- [15] K. Lee, S.-J. Lee, S.-E. Kim, H.-M. Choi, D. Kim, et al., "A 51mW 1.6GHz on-chip network for low-power heterogeneous SoC platform," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 152-153, Feb. 2004.
- [16] S.-J. Lee, K. Lee, S.-J. Song, and H.-J. Yoo, "Packet-switched on-chip interconnection network for system-on-chip applications," *Circuits and Systems II, IEEE Trans. on*, vol. 52, pp. 308-312, June 2005.
- [17] S. Vangal, J. Howard, G. Ruhl, S. Dighe, H. Wilson, et al., "An 80-Tile 1.28TFLOPS Network-on-Chip in 65nm CMOS," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 98-99, Feb. 2007.
- [18] D. Lattard, E. Beigné, C. Bernard, C. Bour, F. Clermidy, et al., "A Telecom Baseband Circuit based on an Asynchronous Network-on-Chip," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 258-259, Feb. 2007.
- [19] B. Kleveland, Q. Xiaoning, L. Madden, T. Furusawa, R. W. Dutton, et al., "High-frequency characterization of on-chip digital interconnects," *Solid-State Circuits, IEEE Journal of*, vol. 37, pp. 716-725, June 2002.
- [20] A. P. Jose, G. Patounakis, and K. L. Shepard, "Pulsed current-mode signaling for nearly speed-of-light intrachip communication," *Solid-State Circuits, IEEE Journal of*, vol. 41, pp. 772-780, April 2006.
- [21] R. T. Chang, N. Talwalkar, C. P. Yue, and S. S. Wong, "Near speed-of-light signaling over on-chip electrical interconnects," *Solid-State Circuits, IEEE Journal of*, vol. 38, pp. 834-838, May 2003.
- [22] A. V. Mezhiba and E. G. Friedman, "Inductive properties of high-performance power distribution grids," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 10, pp. 762-776, 2002.
- [23] M. T. Bohr, "Interconnect scaling-the real limiter to high performance ULSI," *Proceedings of the Int. Electron Devices Meeting*, pp. 241-244, Dec. 1995.
- [24] J. A. Davis, R. Venkatesan, A. Kaloyeros, M. Beylansky, S. J. Souri, et al., "Interconnect limits on gigascale integration (GSI) in the 21st century," *Proceedings of the IEEE*, vol. 89, pp. 305-324, March 2001.
- [25] A. Deutsch, P. W. Coteus, G. V. Kopsay, H. H. Smith, C. W. Surovic, et al., "On-chip wiring design challenges for gigahertz operation," *Proceedings of the IEEE*, vol. 89, pp. 529-555, April 2001.
- [26] K. C. Saraswat and F. Mohammadi, "Effect of Scaling of Interconnections on the Time Delay of VLSI Circuits," *Solid-State Circuits, IEEE Journal of*, vol. 17, pp. 275-280, April 1982.
- [27] ITRS, "International Technology Roadmap for Semiconductors," Edition 2001.
- [28] D. Edelstein, J. Heidenreich, R. Goldblatt, W. Cote, C. Uzoh, et al., "Full copper wiring in a sub-0.25um CMOS ULSI technology," *Electron Devices Meeting, 1997. Technical Digest., International*, pp. 773-776, Dec. 1997.
- [29] ITRS, "International Technology Roadmap for Semiconductors," Edition 2005
- [30] M. F. Chang, V. P. Roychowdhury, Z. Liyang, S. Hyunchol, and Q. Yongxi, "RF/wireless interconnect for inter- and intra-chip communications," *Proceedings of the IEEE*, vol. 89, pp. 456-466, April 2001.
- [31] K. Cadien, M. Reshotko, B. Block, A. Bowen, D. Kencke, et al., "Challenges for On-Chip Optical Interconnects," *Proceedings of SPIE*, vol. 5730, pp. 133-143, March 2005.
- [32] S. Radovanovic, "Integrated photodiodes for Gb/s data-rates in standard CMOS technology," University of Twente, The Netherlands, PhD Thesis, Enschede, 2004.

-
- [33] D. Schinkel, E. Mensink, E. A. M. Klumperink, E. van Tuijl, and B. Nauta, "A 3-Gb/s/ch transceiver for 10-mm uninterrupted RC-limited global on-chip interconnects," *Solid-State Circuits, IEEE Journal of*, vol. 41, pp. 297-306, Jan. 2006.
- [34] H. Shah, P. Shiu, B. Bell, M. Aldredge, N. Sopory, et al., "Repeater insertion and wire sizing optimization for throughput-centric VLSI global interconnects," *Computer Aided Design (ICCAD), IEEE/ACM Intern. Conf. on*, pp. 280-284, Nov. 2002.
- [35] D. Pamunuwa, L. R. Zheng, and H. Tenhunen, "Maximizing throughput over parallel wire structures in the deep submicrometer regime," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 11, pp. 224-243, April 2003.
- [36] X.-C. Li, J.-F. Mao, H.-F. Huang, and Y. Liu, "Global interconnect width and spacing optimization for latency, bandwidth and power dissipation," *Electron Devices, IEEE Transactions on*, vol. 52, pp. 2272-2279, Oct. 2005.
- [37] ITRS, "International Technology Roadmap for Semiconductors," Update 2006.
- [38] D. K. Cheng, *Field and Wave Electromagnetics*, 2nd ed.: Addison-Wesley Publishing Company, 1989.
- [39] J. R. Schrader, E. A. M. Klumperink, J. L. Visschers, and B. Nauta, "Pulse-width modulation pre-emphasis applied in a wireline transmitter, achieving 33 dB loss compensation at 5-Gb/s in 0.13-/spl mu/m CMOS," *Solid-State Circuits, IEEE Journal of*, vol. 41, pp. 990-999, April 2006.
- [40] K. Banerjee and A. Mehrotra, "Analysis of on-chip inductance effects for distributed RLC interconnects," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 21, pp. 904-915, Aug. 2002.
- [41] W. J. Dally and J. W. Poulton, *Digital Systems Engineering*: Cambridge University Press, 1998.
- [42] A. Deutsch, G. V. Kopcsay, P. J. Restle, H. H. Smith, G. Katopis, et al., "When are transmission-line effects important for on-chip interconnections?," *Microwave Theory and Techniques, IEEE Transactions on*, vol. 45, pp. 1836-1846, Oct. 1997.
- [43] D. W. Kerst and J. C. Sprott, "Electrical Circuit Modeling of Conductors with Skin Effect" *Applied Physics, Journal of*, vol. 60, pp. 475-481, July 1986.
- [44] J. R. Schrader, "Wireline Equalization using Pulse-Width Modulation," University of Twente, The Netherlands, PhD Thesis, Enschede, 2007.
- [45] Y. Cao, X. Huang, D. Sylvester, K. Tsu-Jae, and H. Chenming, "Impact of on-chip interconnect frequency-dependent R(f)L(f) on digital and RF design," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 13, pp. 158-162, 2005.
- [46] T. Mido and K. Asada, "An analysis on VLSI interconnection considering skin effect," *Design Automation Conference 1998. Proceedings of the ASP-DAC '98. Asia and South Pacific*, pp. 403-408, Feb. 1998.
- [47] L. T. Pillage and R. A. Rohrer, "Asymptotic waveform evaluation for timing analysis," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 9, pp. 352-366, April, 1990.
- [48] C. V. Kashyap, C. J. Alpert, F. Liu, and A. Devgan, "Closed-form expressions for extending step delay and slew metrics to ramp inputs for RC trees," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 23, pp. 509-516, April 2004.

- [49] R. Mita, G. Palumbo, and M. Poli, "Propagation Delay of an RC-Chain With a Ramp Input," *Circuits and Systems II, IEEE Trans. on*, vol. 54, pp. 66-70, Jan. 2007.
- [50] S. Kim and S. S. Wong, "Closed-Form RC and RLC Delay Models Considering Input Rise Time," *Circuits and Systems I, IEEE Trans. on*, vol. 54, pp. 2001-2010, Sept. 2007.
- [51] W. C. Elmore, "The transient response of damped linear network with particular regard to wideband amplifiers," *J. Applied Physics*, vol. 19, pp. 55-63, Jan. 1948.
- [52] P. Penfield and J. Rubinstein, "Signal Delay in RC Tree Networks," *Proc. Design Automation Conf.*, pp. 613-617, June 1981.
- [53] J. Rubinstein, P. Penfield, and M. A. Horowitz, "Signal Delay in RC Tree Networks," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 2, pp. 202-211, July 1983.
- [54] R. Ho, "On-chip wires: scaling and efficiency," Doctor of Philosophy, Department of Electrical Engineering, Stanford University, Aug. 2003.
- [55] A. B. Kahng and S. Muddu, "An analytical delay model for RLC interconnects," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 16, pp. 1507-1514, Dec 1997.
- [56] Y. I. Ismail, E. G. Friedman, and J. L. Neves, "Equivalent Elmore delay for RLC trees," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 19, pp. 83-97, Jan 2000.
- [57] E. Seevinck, P. J. van Beers, and H. Ontrop, "Current-mode techniques for high-speed VLSI circuits with application to current sense amplifier for CMOS SRAM's," *Solid-State Circuits, IEEE Journal of*, vol. 26, pp. 525-536, April 1991.
- [58] D. G. Manolakis, V. K. Ingle, and S. M. Kogon, *Statistical and Adaptive Signal Processing - Spectral Estimation, Signal Modeling, Adaptive Filtering, and Array Processing*. Norwood MA, USA: Artech House, 2005.
- [59] W.-K. Chen, *The Circuits and Filters Handbook*, 2 ed., 2003.
- [60] R. Bronson, *Differential equations*: McGraw-Hill, 2003.
- [61] A. P. Jose and K. L. Shepard, "Distributed Loss-Compensation Techniques for Energy-Efficient Low-Latency On-Chip Communication," *Solid-State Circuits, IEEE Journal of*, vol. 42, pp. 1415-1424, June 2007.
- [62] E. Mensink, D. Schinkel, E. Klumperink, E. van Tuijl, and B. Nauta, "A 0.28pJ/b 2Gb/s/ch Transceiver in 90nm CMOS for 10mm On-Chip Interconnects," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 414-415, Feb. 2007.
- [63] L. Zhang, J. Wilson, R. Bashirullah, L. Lei, X. Jian, et al., "A 32Gb/s On-chip Bus with Driver Pre-emphasis Signaling," *Custom Integrated Circuits Conference, Proc. of the IEEE*, pp. 265-268, Sept 2006.
- [64] L. Zhang, J. M. Wilson, R. Bashirullah, L. Lei, X. Jian, et al., "A 32-Gb/s On-Chip Bus With Driver Pre-Emphasis Signaling," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 17, pp. 1267-1274, Sept 2009.
- [65] A. Katoch, E. Seevinck, and H. Veendrick, "Fast signal propagation for point to point on-chip long interconnects using current sensing," *European Solid-State Circuits Conf. (ESSCIRC), Proc. of the*, pp. 195-198, Sept. 2002.
- [66] R. Bashirullah, L. Wentai, R. Cavin, and D. Edwards, "A 16Gb/s adaptive bandwidth on-chip bus based on hybrid current/voltage mode signaling," *VLSI Circuits, Digest of Tech. Papers, Symp. on*, pp. 392-393, June 2004.

-
- [67] A. Maheshwari and W. Bureson, "Differential current-sensing for on-chip interconnects," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 12, pp. 1321-1329, Dec 2004.
- [68] A. Katoch, H. Veendrick, and E. Seevinck, "High speed current-mode signaling circuits for on-chip interconnects," *Circuits and Systems (ISCAS), IEEE Intern. Symp. on*, pp. 4138-4141, May 2005.
- [69] N. Tzartzanis and W. W. Walker, "Differential current-mode sensing for efficient on-chip global signaling," *Solid-State Circuits, IEEE Journal of*, vol. 40, pp. 2141-2147, Nov 2005.
- [70] R. Bashirullah, L. Wentai, R. Cavin, III, and D. Edwards, "A 16 Gb/s adaptive bandwidth on-chip bus based on hybrid current/voltage mode signaling," *Solid-State Circuits, IEEE Journal of*, vol. 41, pp. 461-473, Feb. 2006.
- [71] B. Kim and V. Stojanovic, "A 4Gb/s/ch 356fJ/b 10mm equalized on-chip interconnect with nonlinear charge-injecting transmit filter and transimpedance receiver in 90nm CMOS," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 66-67,67a, Feb. 2009.
- [72] B. Kim and V. Stojanovic, "An Energy-Efficient Equalized Transceiver for RC-Dominant Channels," *Solid-State Circuits, IEEE Journal of*, vol. 45, pp. 1186-1197, June 2010.
- [73] D. Schinkel, E. Mensink, E. Klumperink, E. van Tuijl, and B. Nauta, "A 3Gb/s/ch transceiver for RC-limited on-chip interconnects," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 386-387,606, Feb. 2005.
- [74] R. Ho, T. Ono, F. Liu, R. Hopkins, A. AChow, et al., "High-Speed and Low-Energy Capacitively-Driven On-Chip Wires," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 412-413, Feb. 2007.
- [75] B. Razavi, *RF Microelectronics*: Prentice Hall PTR, 1997.
- [76] Wikipedia. <http://en.wikipedia.org/wiki/Gyrator>. Available: <http://en.wikipedia.org/wiki/Gyrator>
- [77] J. Bae, J.-Y. Kim, and H.-J. Yoo, "A 0.6pJ/b 3Gb/s/ch transceiver in 0.18 um CMOS for 10mm on-chip interconnects," *Circuits and Systems (ISCAS), IEEE Intern. Symp. on*, pp. 2861-2864, May 2008.
- [78] J.-s. Seo, R. Ho, J. Lexau, M. Dayringer, D. Sylvester, et al., "High-bandwidth and low-energy on-chip signaling with adaptive pre-emphasis in 90nm CMOS," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 182-183, Feb. 2010.
- [79] A. P. Jose and K. L. Shepard, "Distributed Loss Compensation for Low-latency On-chip Interconnects," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 516-517, Feb. 2006.
- [80] A. Morgenshtein, I. Cidon, A. Kolodny, and R. Ginosar, "Comparative Analysis of Serial and Parallel Links in Networks-on-Chip," *Proc. SoC'04 Conf, Finland*, pp. 185-188, Nov. 2004.
- [81] E. Mensink, D. Schinkel, E. Klumperink, E. van Tuijl, and B. Nauta, "Optimally-placed twists in global on-chip differential interconnects," *European Solid-State Circuits Conf. (ESSCIRC), Proc. of the*, pp. 475-478, Sept. 2005.
- [82] E. Mensink, D. Schinkel, E. A. M. Klumperink, E. Van Tuijl, and B. Nauta, "Optimal Positions of Twists in Global On-Chip Differential Interconnects," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 15, pp. 438-446, April 2007.

- [83] H. Hidaka, K. Fujishima, Y. Matsuda, M. Asakura, and T. Yoshihara, "Twisted bit-line architectures for multi-megabit DRAMs," *Solid-State Circuits, IEEE Journal of*, vol. 24-27, p. 21, Feb. 1989.
- [84] R. Ho, K. Mai, and M. Horowitz, "Efficient on-chip global interconnects," *VLSI Circuits, Digest of Tech. Papers, Symp. on*, pp. 271-274, June, 2003.
- [85] H. Zhang, V. George, and J. M. Rabaey, "Low-swing on-chip signaling techniques: effectiveness and robustness," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 8, pp. 264-272, June 2000.
- [86] D. Schinkel, E. Mensink, E. Klumperink, E. van Tuijl, and B. Nauta, "A Double-Tail Latch-Type Voltage Sense Amplifier with 18ps Setup+Hold Time," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 314-315, Feb. 2007.
- [87] Wikipedia. http://en.wikipedia.org/wiki/Modal_analysis. Available: http://en.wikipedia.org/wiki/Modal_analysis
- [88] A. Hedayat and W. D. Wallis, "Hadamard Matrices and Their Applications," *The Annals of Statistics*, vol. 6, pp. 1184-1238, Nov. 1978.
- [89] S. Srinivasaraghavan and W. Burlison, "Interconnect effort - a unification of repeater insertion and logical effort," *VLSI, Proc. IEEE Computer Society Annual Symposium on*, pp. 55-61, Feb. 2003.
- [90] H. Kaul and D. Sylvester, "Low-power on-chip communication based on transition-aware global signaling (TAGS)," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 12, pp. 464-476, May 2004.
- [91] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 3 ed. New York: McGraw-Hill, 1991.
- [92] L. W. Couch, *Digital and Analog Communication Systems*, Fifth ed.: Prentice Hall, 1997.
- [93] R. Bashirullah, L. Wentai, R. Cavin, and D. Edwards, "A hybrid current/voltage mode on-chip signaling scheme with adaptive bandwidth capability," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 12, pp. 876-880, Aug. 2004.
- [94] P. Z. Peebles, *Digital Communication Systems*: Prentice Hall, 1986.
- [95] J. G. Proakis, *Digital Communications*, 4 ed.: McGraw-Hill, 2000.
- [96] B. Casper, M. Haycock, and R. Mooney, "An accurate and efficient analysis method for multi-Gb/s chip-to-chip signaling schemes," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, pp. 54-57, June 2002.
- [97] R. Farjad-Rad, C. K. K. Yang, M. A. Horowitz, and T. H. Lee, "A 0.4-um CMOS 10-Gb/s 4-PAM pre-emphasis serial link transmitter," *Solid-State Circuits, IEEE Journal of*, vol. 34, pp. 580-585, May 1999.
- [98] J. L. Zerbe, C. W. Werner, V. Stojanovic, F. Chen, J. Wei, et al., "Equalization and clock recovery for a 2.5-10-Gb/s 2-PAM/4-PAM backplane transceiver cell," *Solid-State Circuits, IEEE Journal of* vol. 38, pp. 2121 - 2130, Dec. 2003.
- [99] R. Payne, B. Bhakta, S. Ramaswamy, W. Song, J. Powers, et al., "A 6.25Gb/s binary adaptive DFE with first post-cursor tap cancellation for serial backplane communications," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 68-69, 585, Feb. 2005.
- [100] M. Sorna, T. Beukerna, K. Selander, S. Zier, B. Ji, et al., "A 6.4Gb/s CMOS SerDes core with feedforward and decision-feedback equalization," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 62-63, 585, Feb. 2005.

-
- [101] Wikipedia. http://en.wikipedia.org/wiki/Error_vector_magnitude. Available: http://en.wikipedia.org/wiki/Error_vector_magnitude
- [102] P. K. Hanumolu, B. Casper, R. Mooney, W. Gu-Yeon, and M. Un-Ku, "Analysis of PLL clock jitter in high-speed serial links," *Circuits and Systems II, IEEE Trans. on*, vol. 50, pp. 879-886, Nov. 2003.
- [103] C. Hogge, Jr., "A self correcting clock recovery circuit," *Lightwave Technology, Journal of*, vol. 3, pp. 1312-1314, Dec 1985.
- [104] J. D. H. Alexander, "Clock recovery from random binary signals," *Electronics Letters*, vol. 11, pp. 541-542, Oct 1975.
- [105] R. Farjad-Rad, C. K. K. Yang, M. A. Horowitz, and T. H. Lee, "A 0.3- μ m CMOS 8-Gb/s 4-PAM serial link transceiver," *Solid-State Circuits, IEEE Journal of*, vol. 35, pp. 757-764, May 2000.
- [106] J. T. Stonick, W. Gu-Yeon, J. L. Sonntag, and D. K. Weindler, "An adaptive PAM-4 5-Gb/s backplane transceiver in 0.25- μ m CMOS," *Solid-State Circuits, IEEE Journal of*, vol. 38, pp. 436-443, March 2003.
- [107] H. Johnson. (Feb. 2000). *Multilevel signaling*. Available: <http://www.sigcon.com/Pubs/misc/mls.htm>.
- [108] R. T. Chang, C. P. Yue, and S. S. Wong, "Near speed-of-light on-chip electrical interconnect," *VLSI Circuits, Digest of Tech. Papers, Symp. on*, pp. 18-21, June 2002.
- [109] E. Mensink, D. Schinkel, E. A. M. Klumperink, and E. Van Tuijl, "Interconnects and On-Chip Data Communication Techniques," *Circuits, Systems and Signal Processing (ProRISC), Annual Workshop on*, Nov. 2004.
- [110] A. Amirkhany, A. Abbasfar, J. Savoj, M. Jeeradit, B. Garlepp, et al., "A 24Gb/s Software Programmable Multi-Channel Transmitter," *VLSI Circuits, Digest of Tech. Papers, Symp. on*, pp. 38-39, June 2007.
- [111] T. Beukema, M. Sorna, K. Selander, S. Zier, B. L. Ji, et al., "A 6.4-Gb/s CMOS SerDes core with feed-forward and decision-feedback equalization," *Solid-State Circuits, IEEE Journal of*, vol. 40, pp. 2633-2645, Dec 2005.
- [112] K. Krishna, D. A. Yokoyama-Martin, A. Caffee, C. Jones, M. Loikkanen, et al., "A multigigabit backplane transceiver core in 0.13- μ m CMOS with a power-efficient equalization architecture," *Solid-State Circuits, IEEE Journal of*, vol. 40, pp. 2658-2666, Dec 2005.
- [113] R. Payne, P. Landman, B. Bhakta, S. Ramaswamy, W. Song, et al., "A 6.25-Gb/s binary transceiver in 0.13- μ m CMOS for serial data transmission across high loss legacy backplane channels," *Solid-State Circuits, IEEE Journal of*, vol. 40, pp. 2646-2657, Dec 2005.
- [114] J. F. Bulzacchelli, M. Meghelli, S. V. Rylov, W. Rhee, A. V. Rylyakov, et al., "A 10-Gb/s 5-Tap DFE/4-Tap FFE Transceiver in 90-nm CMOS Technology," *Solid-State Circuits, IEEE Journal of*, vol. 41, pp. 2885-2900, Dec. 2006.
- [115] J. Ren, H. Lee, Q. Lin, B. Leibowitz, E. H. Chen, et al., "Precursor ISI Reduction in High-Speed I/O," *VLSI Circuits, Digest of Tech. Papers, Symp. on*, pp. 134-135, June 2007.
- [116] K. Fukuda, H. Yamashita, F. Yuki, M. Yagyū, R. Nemoto, et al., "An 8Gb/s Transceiver with 3x-Oversampling 2-Threshold Eye-Tracking CDR Circuit for -36.8dB-loss Backplane," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 98-598, Feb 2008.

- [117] H. Lee, K.-Y. K. Chang, J.-H. Chun, T. Wu, Y. Frans, et al., "A 16 Gb/s/Link, 64 GB/s Bidirectional Asymmetric Memory Interface," *Solid-State Circuits, IEEE Journal of*, vol. 44, pp. 1235-1247, April 2009.
- [118] J. F. Buckwalter, M. Meghelli, D. J. Friedman, and A. Hajimiri, "Phase and amplitude pre-emphasis techniques for low-power serial links," *Solid-State Circuits, IEEE Journal of*, vol. 41, pp. 1391-1399, June 2006.
- [119] K. L. J. Wong, E. H. Chen, and C. K. K. Yang, "Edge and Data Adaptive Equalization of Serial-Link Transceivers," *Solid-State Circuits, IEEE Journal of*, vol. 43, pp. 2157-2169, Sept 2008.
- [120] C. Pelard, E. Gebara, A. J. Kim, M. G. Vrazel, F. Bien, et al., "Realization of multigigabit channel equalization and crosstalk cancellation integrated circuits," *Solid-State Circuits, IEEE Journal of*, vol. 39, pp. 1659-1670, Oct. 2004.
- [121] R. Ho, T. Ono, R. D. Hopkins, A. Chow, J. Schauer, et al., "High Speed and Low Energy Capacitively Driven On-Chip Wires," *Solid-State Circuits, IEEE Journal of*, vol. 43, pp. 52-60, Jan 2008.
- [122] L. Luo, J. M. Wilson, S. E. Mick, X. Jian, Z. Liang, et al., "3 gb/s AC coupled chip-to-chip communication using a low swing pulse receiver," *Solid-State Circuits, IEEE Journal of*, vol. 41, pp. 287-296, Jan 2006.
- [123] G. Balamurugan, J. Kennedy, G. Banerjee, J. E. Jaussi, M. Mansuri, et al., "A Scalable 5-15 Gbps, 14-75 mW Low-Power I/O Transceiver in 65 nm CMOS," *Solid-State Circuits, IEEE Journal of*, vol. 43, pp. 1010-1019, April 2008.
- [124] E. Mensink, D. Schinkel, E. A. M. Klumperink, E. van Tuijl, and B. Nauta, "Power Efficient Gigabit Communication Over Capacitively Driven RC-Limited On-Chip Interconnects," *Solid-State Circuits, IEEE Journal of*, vol. 45, pp. 447-457, Feb. 2010.
- [125] M. H. Shakiba, "A 2.5 Gb/s adaptive cable equalizer," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 396-397, Feb 1999.
- [126] Y. Kudoh, M. Fukaiishi, and M. Mizuno, "A 0.13- μ m CMOS 5-Gb/s 10-m 28AWG cable transceiver with no-feedback-loop continuous-time post-equalizer," *Solid-State Circuits, IEEE Journal of*, vol. 38, pp. 741-746, May 2003.
- [127] R. Farjad-Rad, N. Hiok-Taiq, M. J. Edward Lee, R. Senthinathan, W. J. Dally, et al., "0.622-8.0 Gbps 150 mW serial IO macrocell with fully flexible preemphasis and equalization," *VLSI Circuits, Digest of Tech. Papers, Symp. on*, pp. 63-66, June 2003.
- [128] J.-S. Choi, M.-S. Hwang, and D.-K. Jeong, "A 0.18 μ m CMOS 3.5-gb/s continuous-time adaptive cable equalizer using enhanced low-frequency gain control method," *Solid-State Circuits, IEEE Journal of*, vol. 39, pp. 419-425, March 2004.
- [129] S. Gondi, L. Jri, D. Takeuchi, and B. Razavi, "A 10Gb/s CMOS adaptive equalizer for backplane applications," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 328-601 Vol. 1, Feb. 2005.
- [130] G. E. Zhang and M. M. Green, "A 10 Gb/s BiCMOS adaptive cable equalizer," *Solid-State Circuits, IEEE Journal of*, vol. 40, pp. 2132-2140, Nov 2005.
- [131] S. Gondi and B. Razavi, "Equalization and Clock and Data Recovery Techniques for 10-Gb/s CMOS Serial-Link Receivers," *Solid-State Circuits, IEEE Journal of*, vol. 42, pp. 1999-2011, Sept 2007.

-
- [132] F. Gerfers, G. W. den Besten, P. V. Petkov, J. E. Conder, and A. J. Koellmann, "A 0.2-2 Gb/s 6x OSR Receiver Using a Digitally Self-Adaptive Equalizer," *Solid-State Circuits, IEEE Journal of*, vol. 43, pp. 1436-1448, June 2008.
- [133] C.-F. Liao and S.-I. Liu, "A 40 Gb/s CMOS Serial-Link Receiver With Adaptive Equalization and Clock/Data Recovery," *Solid-State Circuits, IEEE Journal of*, vol. 43, pp. 2492-2502, Nov. 2008.
- [134] Y.-S. Sohn, S.-J. Bae, H.-J. Park, C.-H. Kim, and S.-I. Cho, "A 2.2 Gbps CMOS look-ahead DFE receiver for multidrop channel with pin-to-pin time skew compensation," *Custom Integrated Circuits Conference, Proc. of the IEEE*, pp. 473-476, Sept 2003.
- [135] V. Stojanovic, A. Ho, B. W. Garlepp, F. Chen, J. Wei, et al., "Autonomous dual-mode (PAM2/4) serial link transceiver with adaptive equalization and data recovery," *Solid-State Circuits, IEEE Journal of*, vol. 40, pp. 1012-1026, April 2005.
- [136] E. H. Chen, R. Jihong, B. Leibowitz, L. Hae-Chang, L. Qi, et al., "Near-Optimal Equalizer and Timing Adaptation for I/O Links Using a BER-Based Metric," *Solid-State Circuits, IEEE Journal of*, vol. 43, pp. 2144-2156, Sept 2008.
- [137] H. Wang and J. Lee, "A 21-Gb/s 87-mW Transceiver With FFE/DFE/Analog Equalizer in 65-nm CMOS Technology," *Solid-State Circuits, IEEE Journal of*, vol. 45, pp. 909-920, April 2010.
- [138] M. Pozzoni, S. Erba, P. Viola, M. Pisati, E. Depaoli, et al., "A Multi-Standard 1.5 to 10 Gb/s Latch-Based 3-Tap DFE Receiver With a SSC Tolerant CDR for Serial Backplane Communication," *Solid-State Circuits, IEEE Journal of*, vol. 44, pp. 1306-1315, April 2009.
- [139] Y. Liu, B. Kim, T. O. Dickson, J. F. Bulzacchelli, and D. J. Friedman, "A 10Gb/s compact low-power serial I/O with DFE-IIR equalization in 65nm CMOS," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 182-183,183a, Feb 2009.
- [140] S. Kasturia and J. H. Winters, "Techniques for high-speed implementation of nonlinear cancellation," *Selected Areas in Communications, IEEE Journal on*, vol. 9, pp. 711-717, June 1991.
- [141] V. Stojanovic, A. Ho, B. Garlepp, F. Chen, J. Wei, et al., "Adaptive equalization and data recovery in a dual-mode (PAM2/4) serial link transceiver," *VLSI Circuits, Digest of Tech. Papers, Symp. on*, pp. 348-351, June 2004.
- [142] S. R. Sridhara, G. Balamurugan, and N. R. Shanbhag, "Joint Equalization and Coding for On-Chip Bus Communication," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 16, pp. 314-318, March 2008.
- [143] S. Jae-Yoon, N. Jang-Jin, S. Young-Soo, P. Hong-June, K. Chang-Hyun, et al., "A CMOS transceiver for DRAM bus system with a demultiplexed equalization scheme," *Solid-State Circuits, IEEE Journal of*, vol. 37, pp. 245-250, Feb. 2002.
- [144] S.-J. Bae, H.-J. Chi, H.-R. Kim, and H.-J. Park, "A 3Gb/s 8b single-ended transceiver for 4-drop DRAM interface with digital calibration of equalization skew and offset coefficients," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 520-614, Feb. 2005.
- [145] J. E. Jaussi, G. Balamurugan, D. R. Johnson, B. Casper, A. Martin, et al., "8-Gb/s source-synchronous I/O link with adaptive receiver equalization, offset cancellation, and clock de-skew," *Solid-State Circuits, IEEE Journal of*, vol. 40, pp. 80-88, Jan 2005.

- [146] E. H. Chen, R. Jihong, J. Zerbe, B. Leibowitz, L. Haechang, et al., "BER-based Adaptation of I/O Link Equalizers," *VLSI Circuits, Digest of Tech. Papers, Symp. on*, pp. 36-37, June 2007.
- [147] A. X. Widmer and P. A. Franaszek, "A DC-balanced, partitioned-block, 8B/10B transmission code," *IBM Journal of Research and Development*, vol. 27, pp. 440-451, Sept. 1983.
- [148] C. A. Belfiore and J. H. Park, Jr., "Decision feedback equalization," *Proceedings of the IEEE*, vol. 67, pp. 1143-1156, Aug. 1979.
- [149] J. E. C. Brown, P. J. Hurst, and L. Der, "A 35 Mb/s mixed-signal decision-feedback equalizer for disk drives in 2um CMOS," *Solid-State Circuits, IEEE Journal of*, vol. 31, pp. 1258-1266, Sept 1996.
- [150] N. Sithimahachaiikul, J. P. Keane, and P. J. Hurst, "An adaptive DFE using an IIR feedback equalizer for 100Base-TX Ethernet," *Circuits and Systems, 2004. NEWCAS 2004. The 2nd Annual IEEE Northeast Workshop on*, pp. 173-176, June 2004.
- [151] P. M. Crespo and M. L. Honig, "Pole-zero decision feedback equalization with a rapidly converging adaptive IIR algorithm," *Selected Areas in Communications, IEEE Journal on*, vol. 9, pp. 817-829, Aug. 1991.
- [152] H. Tenhunen and D. Pamunuwa, "On dynamic delay and repeater insertion," *Circuits and Systems, IEEE Int. Symp. on*, vol. 1, pp. I-97-I-100 vol.1, Aug 2002.
- [153] A. B. Kahng, S. Muddu, E. Sarto, and R. Sharma, "Interconnect tuning strategies for high-performance ICs," *Design, Automation and Test in Europe, Proceedings*, pp. 471-478, Feb 1998.
- [154] R. Hossain, F. Viglione, and M. Cavalli, "Designing fast on-chip interconnects for deep submicrometer technologies," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 11, pp. 276-280, April 2003.
- [155] G. S. Garcea, N. P. v. d. Meijs, and R. H. J. M. Otten, "Buffer Planning for Global Wires Under Statistical Process Variations," *Circuits, Systems and Signal Processing (ProRISC), Annual Workshop on*, pp. 64-69, Nov. 2003.
- [156] G. S. Garcea, N. P. van der Meijs, K. J. van der Kolk, and R. H. J. M. Otten, "Statistically aware buffer planning," *Design, Automation and Test in Europe Conference and Exhibition, Proceedings*, vol. 2, pp. 1402-1403 Vol.2, Feb. 2004.
- [157] D. Schinkel, E. Mensink, E. A. M. Klumperink, E. Van Tuijl, and B. Nauta, "A Low-Offset Double-Tail Latch-Type Voltage Sense Amplifier," *Circuits, Systems and Signal Processing (ProRISC), Annual Workshop on*, Nov. 2007.
- [158] M. J. E. Lee, W. J. Dally, and P. Chiang, "Low-power area-efficient high-speed I/O circuit techniques," *Solid-State Circuits, IEEE Journal of*, vol. 35, pp. 1591-1599, Nov. 2000.
- [159] W. Ellersick, Y. Chih-Kong Ken, M. Horowitz, and W. Dally, "GAD: A 12-GS/s CMOS 4-bit A/D converter for an equalized multi-level link," *VLSI Circuits, Digest of Tech. Papers, Symp. on*, pp. 49-52, June 1999.
- [160] M. v. Elzakker, E. v. Tuijl, P. Geraedts, D. Schinkel, E. Klumperink, et al., "A 10-bit Charge-Redistribution ADC Consuming 1.9uW at 1 MS/s," *Solid-State Circuits, IEEE Journal of*, vol. 45, pp. 1007-1015, May 2010.
- [161] G. v. d. Plas, S. Decoutere, and S. Donnay, "A 0.16pJ/Conversion-Step 2.5mW 1.25GS/s 4b ADC in a 90nm Digital CMOS Process," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, p. 2310, Feb 2006.

-
- [162] G. v. d. Plas and B. Verbruggen, "A 150 MS/s 133uW 7 bit ADC in 90 nm Digital CMOS," *Solid-State Circuits, IEEE Journal of*, vol. 43, pp. 2631-2640, Dec. 2008.
- [163] T. Kobayashi, K. Nogami, T. Shirotori, and Y. Fujimoto, "A current-controlled latch sense amplifier and a static power-saving input buffer for low-power architecture," *Solid-State Circuits, IEEE Journal of*, vol. 28, pp. 523-527, April 1993.
- [164] B. Nikolic, V. G. Oklobdzija, V. Stojanovic, J. Wenyan, C. James Kar-Shing, et al., "Improved sense-amplifier-based flip-flop: design and measurements," *Solid-State Circuits, IEEE Journal of*, vol. 35, pp. 876-884, June 2000.
- [165] B. Wicht, T. Nirschl, and D. Schmitt-Landsiedel, "Yield and speed optimization of a latch-type voltage sense amplifier," *Solid-State Circuits, IEEE Journal of*, vol. 39, pp. 1148-1158, July 2004.
- [166] K. L. J. Wong and C. K. K. Yang, "Offset compensation in comparators with minimum input-referred supply noise," *Solid-State Circuits, IEEE Journal of*, vol. 39, pp. 837-840, May 2004.
- [167] P. Nuzzo, F. De Bernardinis, P. Terreni, and G. Van der Plas, "Noise Analysis of Regenerative Comparators for Reconfigurable ADC Architectures," *Circuits and Systems I, IEEE Trans. on*, vol. 55, pp. 1441-1454, July 2008.
- [168] W. C. Madden and W. J. Bowhill, "High Input Impedance Strobed CMOS Differential Sense Amplifier," March 1990.
- [169] P. M. Figueiredo and J. C. Vital, "Kickback noise reduction techniques for CMOS latched comparators," *Circuits and Systems II, IEEE Trans. on*, vol. 53, pp. 541-545, July 2006.
- [170] Y. Okaniwa, H. Tamura, M. Kibune, D. Yamazaki, C. Tsz-Shing, et al., "A 40-Gb/s CMOS clocked comparator with bandwidth modulation technique," *Solid-State Circuits, IEEE Journal of*, vol. 40, pp. 1680-1687, Aug. 2005.
- [171] B. Goll and H. Zimmermann, "A low-power 2-GSample/s comparator in 120 nm CMOS technology," *European Solid-State Circuits Conf. (ESSCIRC), Proc. of the*, pp. 507-510, Sept. 2005.
- [172] M. Matsui, H. Hara, Y. Uetani, K. Lee-Sup, T. Nagamatsu, et al., "A 200 MHz 13 mm² 2-D DCT macrocell using sense-amplifying pipeline flip-flop scheme," *Solid-State Circuits, IEEE Journal of*, vol. 29, pp. 1482-1490, Dec. 1994.
- [173] M. v. Elzakker, E. v. Tuijl, P. Geraedts, D. Schinkel, E. Klumperink, et al., "A 1.9uW 4.4fJ/Conversion-step 10b 1MS/s Charge-Redistribution ADC," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 244-610, Feb 2008.
- [174] M. Miyahara, Y. Asada, P. Daehwa, and A. Matsuzawa, "A low-noise self-calibrating dynamic comparator for high-speed ADCs," *Asian Solid-State Circuits Conf. (A-SSCC), Dig. Tech. Papers*, pp. 269-272, Nov 2008.
- [175] T. Sepke, P. Holloway, C. G. Sodini, and H. S. Lee, "Noise Analysis for Comparator-Based Circuits," *Circuits and Systems I, IEEE Trans. on*, vol. 56, pp. 541-553, March 2009.
- [176] M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," *Solid-State Circuits, IEEE Journal of*, vol. 24, pp. 1433-1439, Oct. 1989.
- [177] T. Jiang and P. Y. Chiang, "Sense amplifier power and delay characterization for operation under low-V_{dd} and low-voltage clock swing," *Circuits and Systems, IEEE Int. Symp. on*, pp. 181-184, May 2009.

- [178] H. Zhang and P. Mazumder, "Design of a new sense amplifier flip-flop with improved power-delay-product," *Circuits and Systems, IEEE Int. Symp. on*, pp. 1262-1265 Vol. 2, May 2005.
- [179] S. Vangal, A. Singh, J. Howard, S. Dighe, N. Borkar, et al., "A 5.1GHz 0.34mm² Router for Network-on-Chip Applications," *VLSI Circuits, Digest of Tech. Papers, Symp. on*, pp. 42-43, June 2007.
- [180] D. Schinkel, E. Mensink, E. Klumperink, E. van Tuijl, and B. Nauta, "Low-Power, High-Speed Transceivers for Network-on-Chip Communication," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 17, pp. 12-21, Jan 2009.
- [181] C. Svensson, "Optimum voltage swing on on-chip and off-chip interconnect," *Solid-State Circuits, IEEE Journal of*, vol. 36, pp. 1108-1112, July 2001.
- [182] F. Worm, P. Ienne, P. Thiran, and G. De Micheli, "A robust self-calibrating transmission scheme for on-chip networks," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 13, pp. 126-139, Jan. 2005.
- [183] P. T. Wolkotte, G. J. M. Smit, G. K. Rauwerda, and L. T. Smit, "An Energy-Efficient Reconfigurable Circuit-Switched Network-on-Chip," *IEEE Proc. Int. Symp. Parallel and Distributed Processing*, pp. 155a-155a, April 2005.
- [184] L. Zhang, J. Wilson, R. Bashirullah, L. Lei, X. Jian, et al., "Driver pre-emphasis techniques for on-chip global buses," *Low Power Electronics and Design (ISLPED), Proc. of the Intern. Symp. on*, pp. 186-191, Aug. 2005.
- [185] B. Kim, Y. Liu, T. O. Dickson, J. F. Bulzacchelli, and D. J. Friedman, "A 10-Gb/s Compact Low-Power Serial I/O With DFE-IIR Equalization in 65-nm CMOS," *Solid-State Circuits, IEEE Journal of*, vol. 44, pp. 3526-3538, Dec. 2009.
- [186] J. F. Buckwalter and A. Hajimiri, "Cancellation of crosstalk-induced jitter," *Solid-State Circuits, IEEE Journal of*, vol. 41, pp. 621-632, March 2006.
- [187] H.-K. Jung, K. Lee, J.-S. Kim, J.-J. Lee, J.-Y. Sim, et al., "A 4 Gb/s 3-bit Parallel Transmitter With the Crosstalk-Induced Jitter Compensation Using TX Data Timing Control," *Solid-State Circuits, IEEE Journal of*, vol. 44, pp. 2891-2900, Nov 2009.
- [188] S.-W. Choi, H.-B. Lee, and H. J. Park, "A three-data differential signaling over four conductors with pre-emphasis and equalization: a CMOS current mode implementation," *Solid-State Circuits, IEEE Journal of*, vol. 41, pp. 633-641, March 2006.
- [189] A. Carusone, K. Farzan, and D. A. Johns, "Differential signaling with a reduced number of signal paths," *Circuits and Systems II, IEEE Trans. on*, vol. 48, pp. 294-300, March 2001.
- [190] K. Yamaguchi, K. Sunaga, S. Kaeriyama, T. Nedachi, M. Takamiya, et al., "12Gb/s duobinary signaling with x2 oversampled edge equalization," *Int. Solid State Circuits Conf. (ISSCC), Dig. Tech. Papers*, pp. 70-585, Feb 2005.
- [191] Wolfram-Mathworld. <http://mathworld.wolfram.com/InverseErf.html>. Available: <http://mathworld.wolfram.com/InverseErf.html>