

Development of a Monolithic Very Large Scale Optoelectronic Integrated Circuit Technology

by

Joseph F. Ahadian

B.S.E.E., University of California, San Diego (1993)

S.M.E.E., Massachusetts Institute of Technology (1995)

Submitted to the
Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2000

© 2000 Massachusetts Institute of Technology
All Rights Reserved.

Author _____

Department of Electrical Engineering and Computer Science
December 17, 1999


Certified by _____

Clifton G. Fonstad, Jr.
Professor of Electrical Engineering
Thesis Supervisor

Accepted by _____

Arthur C. Smith

Chair, Department Committee on Graduate Students
Department of Electrical Engineering and Computer Science



ENG

Development of a Monolithic Very Large Scale Optoelectronic Integrated Circuit Technology

by

Joseph F. Ahadian

Submitted to the Department of Electrical Engineering and Computer Science
on December 17, 1999, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Optical interconnects have been proposed for use in high-speed digital systems as a means of overcoming the performance limitations of electrical interconnects at length scales ranging from one millimeter to one hundred meters. To achieve this goal, an optoelectronic very large scale integration (OE-VLSI) technology is needed which closely couples large numbers of optoelectronic devices, such as light emitters and photodetectors, with complex electronics. This thesis has been concerned with the development of an optoelectronic integration technology known as Epitaxy-on-Electronics (EoE). EoE produces monolithic optoelectronic integrated circuits (OEICs) by combining conventional epitaxial growth and fabrication techniques with commercial GaAs VLSI electronics. Proceeding from previous feasibility demonstrations, the growth and fabrication practices underlying the EoE integration process have been extensively revised and extended. The effectiveness of the resulting process has been demonstrated by fabricating the first monolithic, VLSI-complexity OEICs featuring light-emitting diodes (LEDs). As part of a research foundry project, components of this type were designed and tested by a number of groups involved in optical interconnect system development. To further realize the potential of the EoE technology, and to make its capabilities accessible to a broader user community, the focus of this work was extended beyond the development of the integration process to encompass a study of high-speed photodetectors implemented in the GaAs VLSI process, to examine the role of the EoE technology within optical interconnect applications, to formulate an analytical framework for the design of digital optical interconnects, and to implement compact, low power laser driver and optical receiver circuitry needed to implement these interconnects.

Thesis Supervisor: Clifton G. Fonstad, Jr.
Title: Professor of Electrical Engineering

Acknowledgements

I would, first of all, like to thank my thesis advisor, Professor Clifton G. Fonstad, Jr. I admire Professor Fonstad for the breadth and depth of his insight, for his lack of pretense, and for his kindness and sincerity. I am very grateful for the opportunity he has given me to learn and grow over the last six and one-half years. I would like to also thank my thesis committee members, Professors Akintunde “Tayo” Akinwande and Charles G. Sodini, for undertaking the daunting task of reading this document, for their keen observations, and for their encouragement.

This project would not have been possible without the involvement of Vitesse Semiconductor Corporation co-founder and chief technical officer James M. Mikkelson. Aside from providing access to Vitesse’s people and technology, he has, over countless telephone conversations, been a source of clarity in every phase of this work. There are many people at Vitesse to whom I am thankful, but I would like to especially acknowledge Max Helix for shepherding my handiwork through the fab in Camarillo, Robert N. Deming for helpful circuit-design consultations, and Paul Partyka for useful discussions about photodetectors.

Another key to the successful completion of this thesis has been Professor Leslie A. Kolodziejski. Her involvement, which began as my master’s thesis co-advisor and continued through the demonstration of integrated light emitting diodes, is greatly appreciated.

I owe numerous thanks to Dr. Gale S. Petrich. Charged with maintaining both the Kolodziejski molecular beam epitaxy facility and the “micro lab” fabrication facility, and as the designated troubleshooter and editor for anything needs to be repaired or proofread (including this thesis), he is, in my reckoning, the most important person in Building Thirteen.

This project began with the thesis work of Dr. Krishna V. Shenoy. He laid its foundations and gave me a start for which I am greatly indebted. Over the years, this project has also benefited from the hard work of Eric K. Braun, who unraveled the mysteries of GaAs VLSI thermal stability, Praveen T. Vaidyanathan, who toiled to develop light emitting diode integration processes and characterized the resulting devices, Steven G. Patterson, who languished in a cold, dark room growing the material that was needed, and Dr. Yakov Royter, whose help with circuit design, etching, and metallization rounded out the OPTOCHIP effort. Thank you all.

The OPTOCHIP project benefited from the skills of many collaborators. Two individuals whom I have come to know through this project, Professor Barry L. Shoop of the United States Military Academy and

Professor Charles B. Kuznia of the University of Southern California, have lent extensively of their initiative and productivity and have been a source of much-appreciated support and encouragement. Also, for their involvement in receiver array characterization, I am grateful to Dr. Tomasz Maj and the team at McGill University and to Dr. John A. Trezza of Sanders.

I have also benefited greatly from the generosity of Donald A. Hitko and Daniel R. McMahon. I'd like to thank them for their help in understanding circuit design and board layout, and for access to Professor Sodini's equipment.

I would like to also thank all the of the past and present members of Professor Fonstad's research group for seeing to all the thankless chores that have to be done before the work starts, and for providing well-needed distractions. I would particularly like to thank Dr. Isako Hoshino, Donald S. Crankshaw, and Wojciech P. Giziewicz whom I have successively conscripted into the position of system administrator.

This thesis marks the end of my career as a professional student. I've had a lot of wonderful teachers in the last twenty-two and one-half years, There are too many to list, but in addition to Professor Fonstad, I would like to mention three that have helped to define my path. George Stimson taught me physics at San Dieguito high school, shared my excitement for applying it, and told great stories. Robert Whyte, gave up his free time to help me with a term project for Mr. Stimson's class and taught me enough engineering to put four years of college into perspective. In those four years, at the University of California, San Diego, I was given the opportunity to work in the laboratory of Professor Charles W. Tu. This experience set the direction of my graduate work.

Many of the people mentioned above have crossed the line from colleague or teacher to friend. To their names I would like to add Dr. Rajni J. Aggarwal, Dr. David J. Carter, Espartaco Hidalgo, Dr. Dionisios Margetis, Dr. Edward McKigney, Courtney C. Lane, Joanna M. London, Dr. Vernon E. Shrauger, and Karen Young-Waithe. Each of you has touched my life in an important way.

Finally, for their love and support, I would like to thank my mother and father, Mehrdokht and Faradjollah, my sister and brother, Farzaneh and Farshad, and, last but not least, Kaveri. I hope to make you as proud of me as I am of each of you.

Contents

1	Introduction.	23
1.1	Electrical and Optical Interconnects	24
1.2	OE-VLSI Technologies	41
1.3	Thesis Summary.	49
2	The Epitaxy-on-Electronics Integration Process.	51
2.1	Thermal Stability	54
2.2	Historical Development	58
2.3	Dielectric Growth Window Preparation	61
2.4	Low Temperature Native GaAs Oxide Removal	67
2.5	Reduced Temperature MBE Growth	69
2.6	Polycrystalline Deposit Removal	75
2.7	Optoelectronic Device Fabrication	78
2.8	Completed OEICs	79
3	Integrated LEDs and the OPTOCHIP Project	87
3.1	LED fabrication process.	87
3.2	Bulk LEDs	94
3.3	The OPTOCHIP project.	101
3.4	Integrated LED and OEIC results.	113
3.5	Optical Crosstalk	125
3.6	Lessons From OPTOCHIP	129
4	Vitesse MSM Photodetectors.	131
4.1	MSM Photodetectors in the Standard Vitesse H-GaAs III Process.	132
4.2	Modified MSM Photodetectors in the Vitesse H-GaAs IV Process	137
4.3	A Proposed New Photodetector Structure	176
4.4	Conclusion	183

5	Digital Optical Interconnect Systems	185
5.1	Epitaxy-on-Electronics OEIC Applications	186
5.2	Characteristics of Optical Link Components	191
5.3	Digital Interconnects and Communication Channels	195
5.4	Digital Optical Interconnect Design	200
5.5	Transmitter and Receiver Circuit Requirements	231
6	A Low Power Laser Driver	233
6.1	First-generation Switched Current Mirror Laser Driver	234
6.2	Design of an Active Switched Current Mirror Laser Driver	235
6.3	Measurement Setup	242
6.4	Output Current Static Characteristics	255
6.5	Switching Characteristics	264
6.6	Eye Patterns, Delay Asymmetry, and Jitter	278
6.7	Combined Laser Driver and VCSEL Characteristics	298
6.8	Third-generation Laser Driver Designs	322
6.9	Conclusion	335
7	A Low Power Optical Receiver	337
7.1	Design of a Low-Power Continuous-Time Optical Receiver	337
7.2	Measured Receiver Characteristics	366
7.3	Conclusion	380
8	Conclusion	383
8.1	Accomplishments	383
8.2	Recommendations	386
A	Vitesse GaAs VLSI MESFET Technology	389
B	EoE Fabrication	395
C	Circuit Characterization Setups	403
D	Jitter	427
E	MESFET Drain-Lag	435

F	First-Generation Optical Receiver	447
G	Receiver-Array OEICs	455
H	A Fully-Differential Array Receiver	483
I	A Clocked Optical Receiver	497
	Bibliography	507

List of Figures

1.1	The interconnect hierarchy	25
1.2	Interconnect length distribution of the Intel Pentium II microprocessor	26
1.3	A parallel fiber optic data link (W. L. Gore & Associates)	29
1.4	An optical backplane (Ericsson “flex-foil”)	30
1.5	A planar lightwave circuit (Motorola)	31
1.6	A free-space optical interconnect demonstration (University of Southern California)	32
2.1	Overview of the Epitaxy-on-Electronics integration process	52
2.2	Metallization on Vitesse H-GaAs III and H-GaAs IV processes.	54
2.3	Vitesse interconnect sheet resistance after 5 hour thermal cycles	56
2.4	Metal 1 sheet resistance thermal stability model	57
2.5	Ohmic contact resistance after 5 hour thermal cycles	58
2.6	Metal 1-ohmic metal-implant contact resistance thermal stability model	59
2.7	The dielectric stack on Vitesse integrated circuits.	61
2.8	Initial DGW preparation technique.	63
2.9	DGWs contaminated by RIE residues	63
2.10	DGW preparation technique based on partial RIE etch of dielectric stack.	64
2.11	SEM images of “Spike” defects resulting from contaminated dry processing.	65
2.12	Successful DGW preparation process	66
2.13	SEM of DGW on an OEIC as received from Vitesse/MOSIS	68
2.14	SEM of nearly completed DGW	68
2.15	A schematic view of the use of in situ atomic hydrogen in MBE	70
2.16	Plot of bandgap vs. lattice constant of some common III-V materials	71
2.17	InGaAs/GaAs/InGaP QW-SCH laser structure.	72
2.18	Pulsed room temperature L-I characteristics of broad-area aluminum-free lasers	73
2.19	L-I characteristics of broad-area LED test structures.	74
2.20	SEM of OEIC surface and cross section showing post-growth polycrystalline deposits	76
2.21	SEM of OEIC following dry etch of polycrystalline deposits	77
2.22	First demonstration of integrated InGaP-based LEDs	79
2.23	Layout plot of the MIT-OEIC-4 test chip	80
2.24	SEM on an integrated LED.	81

2.25	FIBE cross section a completed LED	82
2.26	FIBE cross section of a completed LED; detailed views	83
2.27	EoE integration of and LED, OPFET, and electronics	84
2.28	Integrated LEDs, MSM photodetectors, and VLSI electronics	85
3.1	Cross-sectional diagram of integrated LED structure	88
3.2	SEM on an integrated LED	89
3.3	SEM of InGaP wet etching result.	91
3.4	SEM of LEDs on first OPTOCHIP OEIC showing poor interconnect metallization	92
3.5	SEM of an integrated LED in an 85 μm DGW	93
3.6	Basic LED operation	95
3.7	Optimizing LED wallplug efficiency.	96
3.8	Emission pattern from LEDs made with various contact geometries	97
3.9	Static characteristics of bulk LEDs.	98
3.10	LED contact geometry selection	100
3.11	The OPTOCHIP project	102
3.12	Photograph of a finished OPTOCHIP master die	103
3.13	LED drop in cell used in OPTOCHIP	105
3.14	LED driver drop in cell used in OPTOCHIP	105
3.15	Layout of an MSM photodetector drop-in cell used in OPTOCHIP	106
3.16	Optically-sensitive FET (OPFET) drop-in cell used in OPTOCHIP	107
3.17	GaAs photodetector optical responsivity spectrum.	108
3.18	OPTOCHIP participants	109
3.19	OPTOCHIP OEIC example: optical neural network array.	110
3.20	OPTOCHIP OEIC example: optical network interface	111
3.21	Overview of the SAPIENT optical interconnect demonstration	112
3.22	Layout of the MIT OPTOCHIP design.	114
3.23	Poor material quality at perimeter of 50 μm DGWs	115
3.24	LED angular distribution	116
3.25	Optical spectra of integrated LED emission	117
3.26	Typical static characteristics of integrated LEDs on OPTOCHIP	118
3.27	Uniformity of OPTOCHIP LEDs.	119
3.28	“Basic Optoelectronic Circuit”	120
3.29	Response of circuit shown in Figure 3.28 to an applied optical input.	121
3.30	“Number array”	122
3.31	“Smart pixel”	124

3.32	FIBE cross section of an optical shield	126
3.33	H-GaAs III cross section showing mechanism for backgate crosstalk	127
3.34	Backgate isolation rings in H-GaAs III.	128
4.1	Cross section of MSM photodetector made with standard Vitesse H-GaAs III process	132
4.2	Layout of test block from OPTOCHIP featuring various MSM photodetectors	133
4.3	Dark current of MSM photodetectors made with standard Vitesse H-GaAs III process	134
4.4	Static response of MSM photodetectors made with standard Vitesse H-GaAs III process	135
4.5	Cross section and layout of modified MSM photodetector in the Vitesse H-GaAs IV process	138
4.6	Common biasing configuration of MSM photodetectors in an integrated receiver	139
4.7	Capacitance vs. voltage of modified MSM photodetectors	140
4.8	Dark current of MSM27	142
4.9	Optical absorption coefficients of various semiconductors	143
4.10	Vertical band profiles under various bias conditions	144
4.11	Open-circuited backgate voltage of MSM27.	145
4.12	Static MSM27 photocurrent with backgate open-circuited.	146
4.13	MSM27 quantum efficiency with backgate open	147
4.14	Biasing mechanism for the source and drain Schottky diodes	149
4.15	Overview of static MSM27 photocurrent with backgate grounded	150
4.16	Static MSM27 photocurrent with backgate grounded	151
4.17	MSM27 quantum efficiency with backgate grounded	152
4.18	Comparison of devices with varying active areas	155
4.19	Evolution of a black-box model for MSM27.	157
4.20	Layout of detector test structures on MIT-OEIC-7.	159
4.21	Equivalent circuit for ground-signal-ground pads seen in Figure 4.20.	160
4.22	Measured S_{11} of MSM27.	162
4.23	Terminal admittance of MSM27	163
4.24	Equivalent circuit for MSM27 terminal impedance	164
4.25	Comparison of measured and modeled S_{11} of MSM27	166
4.26	Comparison of measured and modeled conductance of MSM27	167
4.27	Step waveforms of ECL-driven VCSEL used to produce Figure 4.28 and Figure 4.29	169
4.28	MSM27 step response at various modulation levels -- 1 of 2	170
4.29	MSM27 step response at various modulation levels -- 2 of 2	171
4.30	MSM27 impulse response	173
4.31	Linear-system model of MSM27 photocurrent at high optical power levels.	174
4.32	Modification of MSM-like structure to use backgate diode as photodetector	177

4.33	Lateral pin photodiode in the standard Vitesse H-GaAs IV process	178
4.34	75 μm diameter lateral pin photodetector	179
4.35	Band profiles in lateral pin photodetector	180
4.36	Frequency response of photogenerated carrier diffusion in lateral pin photodetector	182
5.1	EoE Tb/s smart pixel OEICs	188
5.2	EoE parallel fiber-optic data link OEICs	190
5.3	Possible integrated VCSEL structures	192
5.4	L-I characteristics of a VCSEL	193
5.5	Comparison of communication links and digital interconnects	196
5.6	Simulated eye diagrams	197
5.7	Anatomy of a digital optical interconnect	201
5.8	Example of timing and logic errors caused by signal distortion	202
5.9	Error probability vs. SNR.	204
5.10	Error probability vs. optical power amplitude and bandwidth	207
5.11	Jitter vs. bandwidth	211
5.12	Jitter vs. optical power amplitude	212
5.13	Jitter vs. transmitter rise/fall time.	213
5.14	Jitter vs. transimpedance amplifier feedback resistance	214
5.15	Jitter vs. transimpedance amplifier FET transconductance.	215
5.16	Power penalty for non-zero extinction ratio	217
5.17	Power penalty for supply noise.	220
5.18	Eye closure and jitter due to intersymbol interference	222
5.19	Power penalty for the effect of ISI on sensitivity.	224
5.20	Power penalty for the effect of ISI on jitter	225
6.1	L-I characteristics of a generic laser diode	233
6.2	Schematic of first-generation switched current mirror laser driver	234
6.3	Schematic of an active switched current mirror laser driver	236
6.4	Simulation of active switched current mirror laser driver	237
6.5	Modifications to laser driver of Figure 6.3	239
6.6	Layout of the active switched current mirror laser driver	240
6.7	Layout of die cut from MIT-OEIC-7 for laser driver characterization	243
6.8	Signal source for characterization of laser driver.	245
6.9	Data signal used to characterize laser driver	246
6.10	Signal path for electrical characterization of the laser driver.	247
6.11	Parasitic coupling of input to output in laser driver electrical test module.	248

6.12	Signal path for optical characterization of the laser driver	250
6.13	Static characteristics of a typical VCSEL	251
6.14	S_{11} characteristics of a typical VCSEL.	253
6.15	Input admittance characteristics of a typical VCSEL	254
6.16	Equivalent small signal input circuit for a typical VCSEL.	254
6.17	Simulated static characteristics of off-state current mirror	256
6.18	Simulated static characteristics of on-state current mirror	257
6.19	Dependence of laser driver static characteristics on output bias voltage.	259
6.20	Static characteristics of the off-state current mirror	260
6.21	Static characteristics of the modulation current mirror.	262
6.22	Combined output of off-state and modulation current mirrors	263
6.23	Laser driver turn-on characteristics at various modulation levels -- 20 ns time scale	265
6.24	Laser driver turn-on characteristics at various modulation levels -- 2 ns time scale	266
6.25	Laser driver turn-off characteristics at various modulation levels -- 20 ns time scale	267
6.26	Laser driver turn-off characteristics at various modulation levels -- 2 ns time scale.	268
6.27	Simulation of laser driver slow turn-on effect at low modulation level	269
6.28	Plots of rise and fall time vs. modulation bias level	273
6.29	Jitter caused by combined effect of slow laser driver settling and linear channel response	276
6.30	Comparison of transition times specifications in meeting noise margin requirements.	277
6.31	100 MHz eye pattern of laser driver output at various modulation levels	279
6.32	Close-up of transitions in 100 MHz eye patterns given in Figure 6.31	280
6.33	990 MHz eye pattern of laser driver output at various modulation levels	281
6.34	Laser output eye patterns for $V_{swg} = 0.25$ V	283
6.35	Edge timing vs. modulation level at 100 Mb/s	285
6.36	Edge timing vs. modulation level at 990 Mb/s	286
6.37	Delay asymmetry vs. modulation level.	289
6.38	Jitter vs. modulation level	290
6.39	Input signal edge timing vs. bit-rate	292
6.40	Edge timing vs. bit-rate with $V_{swg}=1.5$ V	293
6.41	Delay asymmetry vs. bit rate for $V_{swg}=1.5$ V	296
6.42	Jitter vs. bit-rate for $V_{swg}=1.5$ V.	297
6.43	Dependence of off-state current mirror on VCSEL supply voltage	299
6.44	Dependence of modulation current mirror on VCSEL supply voltage.	300
6.45	Static characteristics of off-state current mirror with VCSEL	301
6.46	Static characteristics of modulation current mirror with VCSEL	302
6.47	Combined output of off-state and modulation current mirrors with VCSEL	303

6.48	Laser rising-edge dependence on off-state level	306
6.49	Laser falling-edge dependence on off-state level.	307
6.50	Turn-on delay dependence on off-state level.	308
6.51	Laser rising-edge dependence on modulation level	310
6.52	Laser rising-edge dependence on modulation level	311
6.53	Normalized optical transition waveforms at various modulation levels	313
6.54	Transition times of optical switching waveforms measure at various modulation levels	314
6.55	Simulation of laser driver loaded by VCSEL equivalent input circuit shown in Figure 6.16	316
6.56	Simulation of laser driver loaded by a 500 Ω resistor	317
6.57	Simulation of laser driver loaded by an 85 Ω resistor	318
6.58	Simulation of laser driver loaded by an 25 Ω resistor	319
6.59	Modification of the laser driver to allow large resistive loads	320
6.60	Modified laser driver simulated with 500 Ω load	321
6.61	An improved active switched current mirror laser driver	323
6.62	Simulation of the improved active switched current mirror laser driver shown in Figure 6.61	324
6.63	Schematic of a low-power differential laser driver.	327
6.64	Circuit for generation of complementary logic signals.	328
6.65	Third stage of the differential laser driver	330
6.66	Second stage of the differential laser driver	331
6.67	First stage of the differential laser driver.	333
6.68	Simulation of the differential laser driver shown in Figure 6.63.	334
7.1	Block representation of a low-power continuous time receiver	339
7.2	Schematic of a low-power continuous time receiver.	340
7.3	Simulated static characteristics of low power receiver.	341
7.4	Simplified analysis of transimpedance amplifier used as the first stage	344
7.5	Simulated frequency response of first stage of low power receiver	345
7.6	Isolated schematic of the low power receiver's second stage voltage amplifier	347
7.7	Simulated frequency response of second stage of low power receiver.	348
7.8	Isolated schematic of the low power receiver's third stage differential amplifier	349
7.9	Simulated frequency response of third stage of the low power receiver	350
7.10	Isolated schematic of balanced-to-unbalanced conversion circuit and DCFL inverter	352
7.11	Simulated frequency response after first three stages and bal/unbal conversion.	353
7.12	Simulated transient response of the low power receiver	355
7.13	Close-up of the receiver transient simulation for a rising input transition	356
7.14	Close-up of the receiver transient simulation for a falling input transition	357

7.15	Simulated frequency response to variations in the 2 V supply	360
7.16	Simulated transient response to variations in the 2 V supply	361
7.17	Simulated frequency response to variations in the 3.3 V supply	362
7.18	Simulated transient response to variations in the 3.3 V supply	363
7.19	Layout plot of low power receiver	365
7.20	Setting VREF and IREF for receiver characterization	367
7.21	Time-average output voltage vs. input power	368
7.22	Transient simulation of receiver input with the MSM27 modeled rising edge.	371
7.23	Transient simulation of receiver input with the MSM27 modeled falling edge	372
7.24	Eye pattern of optical signal used to test the receiver	373
7.25	MSM27 source current eye patterns -- 1 of 2.	374
7.26	MSM27 source current eye patterns -- 2 of 2.	375
7.27	Receiver output eye patterns -- 1 of 2	376
7.28	Receiver output eye patterns -- 2 of 2	377
7.29	Simulated transient response with the MSM27 modeled rising edge	378
7.30	Simulated transient response with the MSM27 modeled falling edge	379
A.1	HGaAs3 Process Flow Schematic through Metal 1 [7]	391
C.1	Layout MIT-OEIC-7	404
C.2	Layout of die cut from MIT-OEIC-7 for laser driver characterization.	405
C.3	Block representation of clock and data generators	408
C.4	Clock and data generator portion of laser driver test block	409
C.5	Laser driver portion of laser driver test block	410
C.6	Close up of laser driver test block in the vicinity of the VCSEL driver	411
C.7	Schematic of line driver circuit.	412
C.8	Test board for laser driver + VCSEL characterization	415
C.9	Bonding diagram for laser driver + VCSEL characterization	416
C.10	Test board for laser driver electrical characterization	417
C.11	Bonding diagram for laser driver electrical characterization.	418
C.12	Laser module test board	420
C.13	Laser module construction and schematic	421
C.14	Layout of low power optical receiver test structure on MIT-OEIC-7	423
C.15	Block diagram of low power optical receiver test structure on MIT-OEIC-7	424
C.16	Test board for low power receiver characterization	425
C.17	Bonding diagram for low power receiver characterization.	426

D.1	Pulse shapes and data signals in communication links and digital interconnects	428
D.2	ISI in a digital interconnect.	429
D.3	Simplified waveforms used to calculate ISI induced jitter	431
E.1	Frequency response of MESFET output conductance	436
E.2	Frequency response of common-source amplifier	437
E.3	Step response of common-source amplifier	438
E.4	A simple model for drain-lag.	440
E.5	A common-source amplifier using a cascode	442
E.6	A common-source amplifier using an “active” cascode	443
E.7	Clamping controls drain lag in DCFL logic	445
F.1	MIT-OEIC-6	447
F.2	First generation optical receiver: block diagram	448
F.3	Schematic of first generation optical receiver: first-stage TIA.	450
F.4	Schematic of first generation optical receiver: differential offset stage	450
F.5	Schematic of first generation optical receiver: second gain stage	451
F.6	Schematic of first generation optical receiver: third gain stage and conversion to DCFL output	452
F.7	First generation optical link demonstration.	454
G.1	Diagram of high speed package for use with MIT-OEIC-8 and MIT-OEIC-9.	457
G.2	Layout plot of MIT-OEIC-8	459
G.3	Unit cell of the MIT-OEIC-8 receiver array	460
G.4	Bonding diagram for MIT-OEIC-8.	462
G.5	Typical characteristics from MIT-OEIC-8	466
G.6	Layout plot of MIT-OEIC-9	468
G.7	Unit cell of MIT-OEIC-9 receiver array	470
G.8	Bonding diagram for MIT-OEIC-9.	471
G.9	Static characteristics from MIT-OEIC-9	474
G.10	Plots of P_{th} and IREF vs. the REF voltage.	476
G.11	Optical link demonstration using MIT-OEIC-9	478
G.12	Parallel data transmission using MIT-OEIC-9	479
G.13	Transmission of clock signals using MIT-OEIC-9.	480
G.14	Single-channel data transmission using MIT-OEIC-9	481
G.15	Eye pattern at 1 Gb/s from optical source used to test MIT-OEIC-9 at McGill University	482
H.1	Block representation of fully-differential receiver	484

H.2	Schematic of fully-differential receiver; first stage TIA	485
H.3	Schematic of fully-differential receiver; second stage voltage amplifier	486
H.4	Schematic of fully-differential receiver; final stages	487
H.5	Simulation of fully-differential receiver: static characteristics	488
H.6	Simulation of fully-differential receiver: frequency response	489
H.7	Simulation of fully-differential receiver: long duration transient response	491
H.8	Simulation of fully-differential receiver: short duration transient response	492
H.9	Simulation of fully-differential receiver: supply frequency response	494
H.10	Simulation of fully-differential receiver: supply transient response	495
I.1	Optical bus using continuous time and clocked receivers	498
I.2	Block representation of clocked optical receiver	499
I.3	Schematic of clocked optical receiver	500
I.4	Simulation of the clocked optical receiver	502
I.5	Layout of the clocked optical receiver	503
I.6	Static characteristics of clocked optical receiver	504
I.7	Additional static characteristics of clocked optical receiver	505

List of Tables

1.1	Motivation for optical interconnects at various length scales	34
1.2	Optical interconnect applications and electronic platform technologies	47
3.1	Various ohmic contact geometries tested in conjunction with a 30 mm square mesa	9
4.1	Summary of bias conditions for MSM27 S_{11} measurements.	161
4.2	Component values for the MSM27 terminal impedance equivalent circuit	165
4.3	Summary of MSM27 step response measurement parameters	172
5.1	Comparison of rise time definitions and corresponding bandwidths for various linear systems.	208
6.1	Summary of laser driver output step characteristics	272
6.2	Summary of jitter and delay asymmetry characteristics at 100 Mb/s	287
6.3	Summary of jitter and delay asymmetry characteristics at 990Mb/s.	288
6.4	Summary of input jitter and delay asymmetry characteristics at various frequencies	294
6.5	Summary of output jitter and delay asymmetry characteristics at various frequencies	295
6.6	Operating parameters for optical switching measurements at various off-state levels.	305
6.7	Operating parameters for optical switching measurements at various modulation levels	309
6.8	Transition times of optical switching waveforms at various modulation levels	312
A.1	Vitesse HGaAs3 Fabrication Sequence [7].	390
A.2	HGaAs3 Dielectric and Metal Layer Deposition Characteristics [7].	393
C.1	Definition of pad names used in Figure C.2	406
C.2	Internal clock frequency settings	413
C.3	Taps for maximum-length PRBS generators.	414
G.1	Pad definitions for MIT-OEIC-8	461
G.2	Pinout for packaged MIT-OEIC-8	463
G.3	Pad definitions for MIT-OEIC-9	467
G.4	Pinout for packaged MIT-OEIC-9	472
G.5	Summary of parameters extracted from Figure G.9	475

1

Introduction

The last decade of the twentieth century has witnessed an exponential rise in both computer performance and telecommunication network traffic. Underlying this trend have been tremendous improvements in integrated circuit (IC) manufacturing technology. IC technology is now being driven beyond the very large scale integration (VLSI) regime in order to meet the insatiable demands of high speed computing and communication systems. However, transistor performance and integration density are no longer the primary limitations of system performance. Rather, the wires which electrically interconnect the transistors within an IC, the ICs within a circuit board, and the circuit boards within a piece of equipment and throughout a network are now recognized as barriers to higher information processing capacity.

A major technology thrust aimed at overcoming the interconnect bottleneck has focused on the use of optical interconnects. An optical interconnect consists of a signal source, such as light emitter or modulator, a means of directing the signal to its destination, and a receiver. The optical transmission of information already dominates long-distance voice and data communication because of the low loss and dispersion characteristics of optical fiber. However, optical interconnection at shorter length scales encompasses a distinct set of motivations, design considerations, and technology requirements. These concepts are introduced in Section 1.1 where some of the motivations for the use of optical interconnects are examined and a number of the architectures proposed for their implementation are reviewed.

One of the key technologies needed to realize optical interconnects is the means of fabricating optoelectronic integrated circuits (OEICs) which closely couple VLSI-complexity electronics with large numbers of optoelectronic devices such as semiconductor lasers, light emitting diodes (LEDs), optical

modulators, and photodetectors. The development of an optoelectronic VLSI (OE-VLSI) technology has been the topic of a number of research efforts and is the goal of this thesis. The technology that has been developed is known as Epitaxy-on-Electronics (EoE). Section 1.2 surveys various optoelectronic integration techniques in order to put EoE in context alongside them. This review will also serve to highlight the basic assumptions underlying the development of the EoE process.

EoE uses a completed GaAs VLSI electronics die or wafer as a “substrate” on which conventional optoelectronic devices are epitaxially grown and fabricated. Work on EoE began at MIT in 1990 [1-25] taking advantage of commercially fabricated GaAs VLSI MESFET (metal-semiconductor field-effect transistor) electronics from Vitesse Semiconductor Corporation, Camarillo CA [26]. Initial results have included the fabrication of LED-based neural-network OEICs [5] as well as the integration of resonant tunneling diodes (RTDs) [8] and surface-normal multiple-quantum-well modulators (also known as self electro-optic devices, or SEEDs) [9]. The performance of this early work was limited by a number of process shortcomings which have been addressed by this thesis. Taking a broad view of the technological requirements of optical interconnects, this work has gone beyond the development of an integration process to consider system- and circuit-level issues. A thesis outline appears in Section 1.3.

1.1 Electrical and Optical Interconnects

The motivation for the work of this thesis is the implementation of optical interconnects in high speed digital systems. In this section, interconnects will be classified by their length-scale within an interconnect hierarchy, and a number of optical interconnect architectures which have been proposed for use in various segments of this hierarchy will be reviewed. Some of the key justifications for the use of optical interconnects in these contexts will then be discussed.

The interconnect hierarchy

It is often convenient to classify interconnects of varying length scales in terms of their role in the “interconnect hierarchy.” Figure 1.1 shows the progression of interconnect length scales within computing and communication systems. On-chip interconnects, which connect transistors and gates within a die, are at the short end of the hierarchy (Figure 1.1(a)). Taking 1 cm square as a typical bound on die size, on-chip interconnects are limited to around 2 cm. Figure 1.2 shows a typical length distribution for on-chip inter-

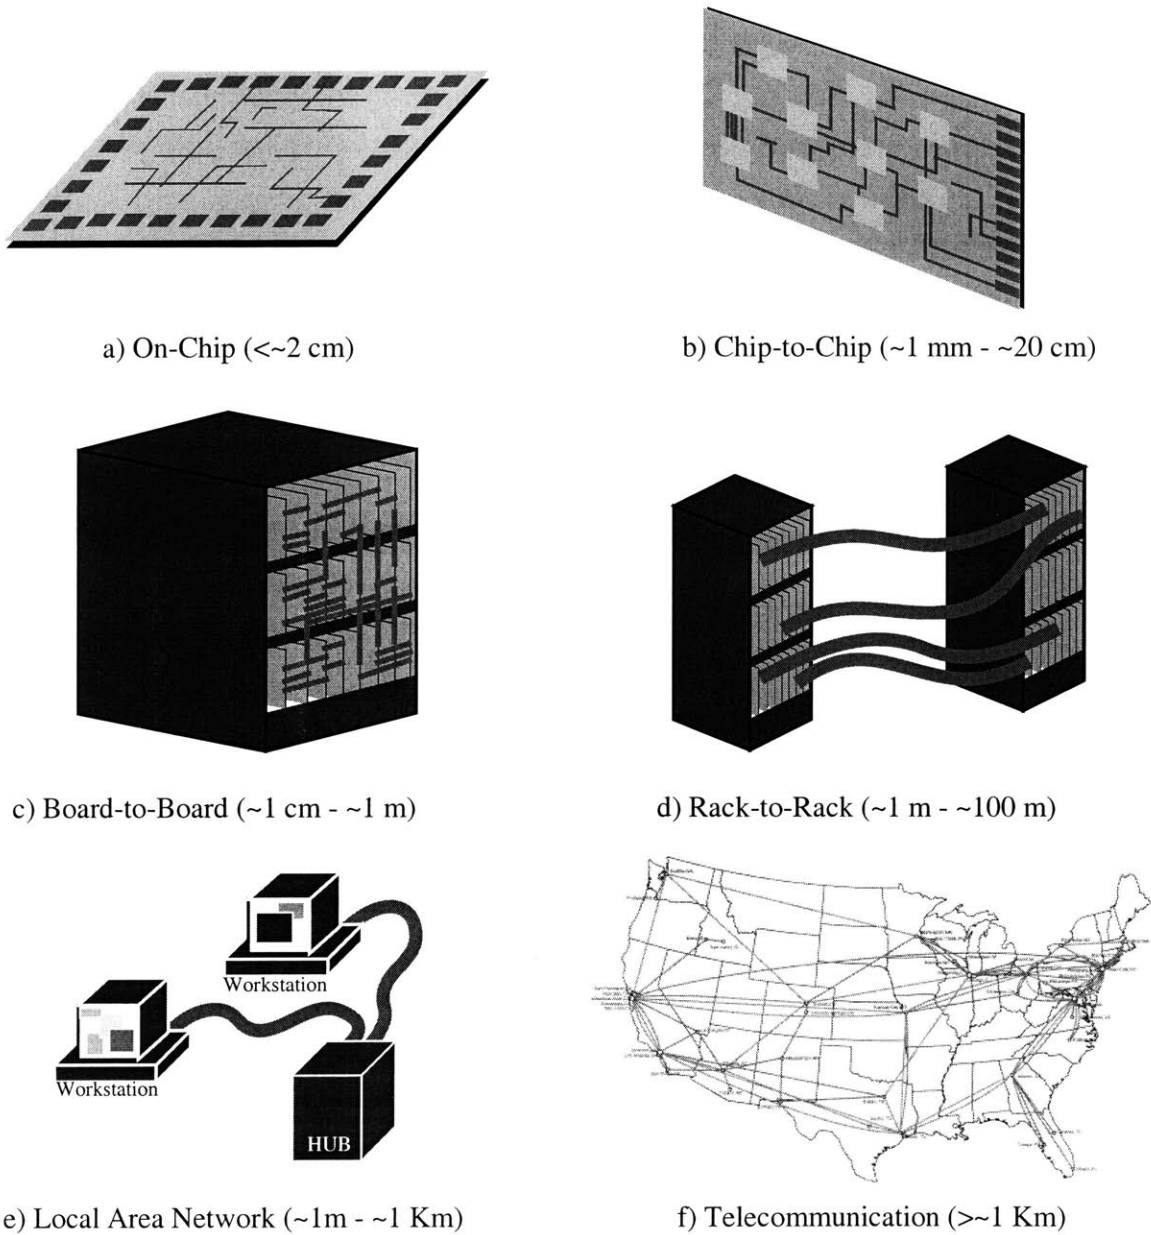


Figure 1.1: The interconnect hierarchy

Interconnect types may be classified as (a) on-chip, (b) chip-to-chip, (c) board-to-board, (d) rack-to-rack, (e) local-area-networks, and (f) long-haul communication channels [27].

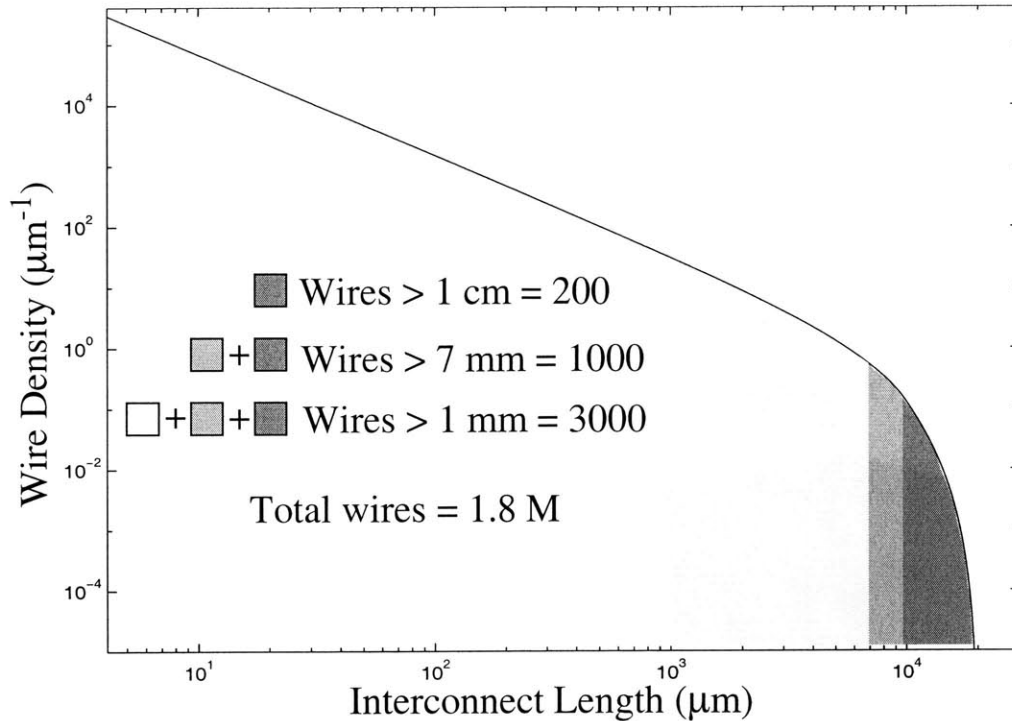


Figure 1.2: Interconnect length distribution of the Intel Pentium II microprocessor

On-chip interconnects actually consist of two sub-classes: a set of extremely numerous and short interconnects, and relatively few, long interconnects. The 1 cm square die of the Intel Pentium II microprocessor contains 5.8 million transistors and 1.8 million interconnects. As seen in the distribution, the vast majority of these wires are quite short. The distribution begins to drop rapidly as the length approaches the dimension of the die. There are 200 wires exceeding 1 cm, 1000 exceeding 7 mm, and 3000 exceeding 1 mm [28]. The basis of the interconnect density distribution plotted above is an empirical relationship, known as Rent's rule, between the number of gates or transistors, N , and the number of terminal, T , within a system: $T = kN^p$. The constant k is the average number of terminals per gate or transistor and the "Rent exponent", p , is a characteristic of the specific technology and design style in use [29,30]. The rent exponent must lie between 0 and 1 and is typically near 0.5. The number of wires in the system, W , is found to be $W = \alpha k N (1 - N^{p-1})$, where $\alpha = \text{fanout}/(\text{fanout} + 1)$ is the fraction of terminals in the system which are inputs [31]. Representing a system as an evenly spaced array of gates, and assuming that Rent's rule applies throughout, a probability density function may be found for the interconnect length:

$$f(\lambda) = \Upsilon \left(2 - 2\lambda + \frac{1}{3}\lambda^2 \right) \lambda^{-(3-2p)} \quad \text{for } \frac{1}{\sqrt{N}} \leq \lambda \leq 1$$

$$f(\lambda) = \frac{\Upsilon}{3} (2 - \lambda)^3 \lambda^{-(4-2p)} \quad \text{for } 1 \leq \lambda \leq 2$$

The variable λ is the interconnect length measured relative to the side dimension of the die, N is the total number of gates, p is the Rent exponent, and Υ is a normalization constant [28,32]. This model compares very well with actual data. For example, the interconnect distribution of every member of the Intel Pentium processor family has been modeled with a Rent exponent of 0.67 [28].

connects. It points out that these interconnects actually consist of two sub-classes: an extremely numerous set (> 1 million) of short interconnects (< 1 mm) and relatively few (\sim thousands), long interconnects (1 mm - 2 cm). Next in the hierarchy are printed wiring boards (PWBs) and multi-chip modules (MCMs). At this chip-to-chip level, there are thousands of interconnects with lengths ranging from a few millimeters to roughly 20 cm (Figure 1.1(b)). PWBs live within electronic cabinets, or racks, and hundreds of board-to-board connections are made through a “backplane” located at the rear of the cabinet (Figure 1.1(c)). Typical board-to-board interconnect lengths range from a few centimeters to around a meter. In larger systems, multiple racks may be required, and may be located up to several hundred meters apart (Figure 1.1(d)).

An important, common trait of on-chip, chip-to-chip, board-to-board, and rack-to-rack links is that they typically interface directly with digital electronics and may carry both data and timing (clock) signals. The length of rack-to-rack connections is ultimately limited by the need to maintain the integrity of these signals at a level required for error-free operation of a digital system.

Overlapping the rack-to-rack length scale, and extending up to around 1 km, are local-area networks (LANs, Figure 1.1(d)). Unlike the four previous interconnect types, LAN connections may use networking protocols to provide error-free communication in spite of an underlying, error-prone channel. The example given in the figure is of a star topology, such as used in “Giga-bit” Ethernet. Similar point-to-point links are used in ring architectures such as Token-ring and FDDI. Conventional, lower-speed Ethernet may use a multi-access bus architecture. The length scale of LAN interconnects means that they may share many of the characteristics of the previous four types, but architecturally LANs have much in common with long-haul telecommunication systems (Figure 1.1(f)).

Throughout this text, the term “interconnect” or “digital interconnect” will refer to links from the on-chip to the rack-to-rack length scales. Long-haul connections and LANs will be referred to as “communication channels”. Optical interconnects, with which this thesis is ultimately concerned, are associated with length scales ranging from ~ 1 mm to ~ 100 m, that is, from the on-chip to the rack-to-rack length scales. Some of the architectures being investigated to implement optical interconnects within this segment of the hierarchy are reviewed next.


Optical interconnect architectures

An optical interconnect consists of a signal source, such as a light emitter or modulator, a light path, and a photodetector. Various arrangements of these components have been proposed to meet the connectivity needs at each level of the interconnect hierarchy. These architectures will be described momentarily. This discussion will focus on light-emitter based implementations of the various architectures. In general, emitter based systems are more readily implemented, and receive the majority of the research community's interest [33-37]. Multiple optical interconnect architectures exist to support each level of the interconnect hierarchy. The material here is a representative, but not exhaustive, catalog of possibilities [38]. Three general categories may be used to classify the architectures: parallel optical links and optical backplanes, planar lightwave circuits, and free-space optical interconnects.


Parallel optical links and optical backplanes

Optical interconnects have been in use in telecommunication switching equipment since the mid-1980s in [39-41], and a similar history exists in data communication applications [42]. These systems, which have used discrete emitters, detectors, and optical fiber to make board-to-board and rack-to-rack connections, have evolved into parallel optical links which include linear arrays of VCSEL, photodetectors, and optical fiber ribbon. Extensive work on the implementation of highly optimized hybrid modules has been carried out by numerous groups including those at Motorola [43,44], IBM and AT&T [45], IBM and 3M [46], HP [47], NTT [48], and a European consortia [49]. These modules typically include up to 32 channels with each one operating at around 1 Gb/s at distances up to 100 m. Commercial products are nearing release from W. L. Gore & Associates [50] and from Infineon (formerly Siemens) [51,52]. The Gore module is illustrated in Figure 1.3. Work is also ongoing to examine the use of plastic optical fiber, rather than traditional glass fiber, in order to further reduce cost [53].

Closely related to the parallel optical interconnects, which use fiber-optic ribbons, are a number of optical backplane systems which allow a high density of optical connections to be managed by embedding them within either a flexible or rigid carrier [54-56]. Optical fibers, which are routed through the carrier, are terminated in parallel- or single-channel connectors at multiple points of the carrier. An example implemented at Ericsson is shown in Figure 1.4.



(a)



(b)

Number of channels	4 - 12
Data rate per channel	DC - 1.5 Gb/s
Bit-Error-Rate	10^{-15}
Jitter	300 ps peak-to-peak
Power Dissipation (Tx+Rx)	205 mW per channel
Operating Wavelength	850 nm
Fiber Type	62.5 μm Graded Index

(c)

Figure 1.3: A parallel fiber optic data link (W. L. Gore & Associates)

The Gore modules use arrays of vertical cavity surface emitting laser (VCSELs) and GaAs PIN photodetectors within compact, highly optimized transmit and receive modules, as shown in (a) [50]. An eye diagram corresponding to 1.5 Gb/s transmission of 2^7-1 bit pseudorandom data is given in (b), and preliminary specifications are shown in (c).




Figure 1.4: An optical backplane (Ericsson “flex-foil”)

The use of a large number of individual optical fibers or fiber ribbons at the back of an instrument rack is both costly and impractical. Imbedding the fibers in a flexible carrier, as in this example implemented by researchers at Ericsson [56], allows for efficient manufacturing and installation.

Planar lightwave circuits

Planar lightwave circuits implement on-chip, chip-to-chip, and chip-to-connector interconnects through optical waveguides in the plane of a printed wiring board or multi-chip module. Reports have included the use of silica waveguides on a silicon platform [57] as well as a number of polymer based [58-60] and spin-on-glass based [61] techniques which are compatible with standard printed wiring boards and multi-chip modules. In these techniques, an angled deflector may be used to couple light between the in-plane optical waveguides and the surface-normal optoelectronic components. The waveguides may intersect orthogonally without unwanted cross-coupling of the optical signal. An example of a polymer planar lightwave circuit implemented atop a conventional multi-chip module at Motorola is shown in Figure 1.5 [62].





Figure 1.5: A planar lightwave circuit (Motorola)


In this demonstration, $50\ \mu\text{m}$ wide \times $80\ \mu\text{m}$ high dielectric waveguides were fabricated from a spin-coated photopatternable epoxy. The waveguides were separated by a $60\ \mu\text{m}$ cladding layer from the alumina multi-chip module substrate (MCM-D with two copper metallization layers). In the SEM images above, a 30 mil square, 10 mil thick silicon PIN photodetector is shown flip-chip bonded astride a waveguide. The 45° angled mirror, which deflects light from the waveguide onto the detector, was formed by laser ablation [62].

Free-space optical interconnects

In a free-space optical interconnect, light signals propagate through free space and are directed by means of optical elements such as lenses, mirrors, and holographic elements [63]. These interconnects have been proposed for use at the chip-to-chip and on-chip length scales. The term “free space” in this context can include a transparent medium other than air. Some recent demonstrations are included in references [64-66]. A demonstration system using a modulator-based OEIC is shown in Figure 1.6 [66]. The three-dimensional, un-guided nature of these architectures hold the promise of very high interconnect complexity and density.



(a)



(b)

Figure 1.6: A free-space optical interconnect demonstration (University of Southern California)

In a free-space optical interconnect system, signals may be routed between chips using optical information processing and imaging techniques. The demonstration shown here used OEICs consisting of surface-normal optical modulators flip-chip bonded to CMOS electronics. The optical path between each chip includes various lenses, mirrors, beam splitters, and diffractive elements. The system is shown symbolically in (a), and a photograph of the experimental setup is shown in (b) [66].

Motivations for the use of optical interconnects

The use of optical interconnects in the contexts described above is markedly different from the use of fiber-optic connections in telecommunication systems. It is the very low loss and dispersion (the tendency of signal pulses to spread out as they propagate) of optical fibers which makes them ideal for long-haul communications. These characteristics are of limited importance at the shorter length scales of digital interconnects. A number of references have examined the motivation for the use of optical interconnects at shorter length scales [39-41,54,67-80]. A comprehensive examination of the issues will not be attempted here. Taking a pragmatic approach, Table 1.1 lists ten issues which arise at various levels of the interconnect hierarchy to limit the performance of high speed digital systems. Each issue is briefly discussed below.

Bandwidth

In the present context, bandwidth refers to the bit-rate capacity of an interconnect. This should not be confused with the information-theoretical channel capacity which may be approached using various coding schemes. The bandwidth limitation of a digital electrical interconnect stems primarily from resistive loss and is approximately given by $B \sim B_0 A l^2$ [81]. In this expression, l is the length of the connection, A is the cross sectional area of the interconnect or aggregate cross sectional area of an ensemble of interconnects, and B_0 is a constant which is roughly 10^{15} b/s for a non-equalized LC transmission line, 10^{16} b/s for an RC line, and 10^{17} - 10^{18} b/s for an equalized LC line. For a rack-to-rack run length of 10 m using 30 cm x 30 cm total cable cross sectional area, the theoretical bandwidth limit is around 1 Tb/s. On a high quality multi-chip module (5 cm wide MCM-D) the aggregate cross-sectional bandwidth would be around 750 Gb/s. On a 1 cm x 1 cm die, the aggregate bandwidth limit on long interconnects would be around 4 Tb/s. In the near future, bandwidth is not expected to pose a problem at the on-chip level. However, at the chip-to-chip, board-to-board, and rack-to-rack length scales, the bandwidth requirements of future systems are expected to exceed the capabilities of electrical interconnects.

Theoretical predictions aside, the bit-rate limit for a copper trace (~40 cm long) on an FR-4 printed circuit board is set by both skin loss and dielectric loss to around 4 Gb/s [82] (actual systems are limited by skew well below this point; see below). In high quality cabling, such as shielded twisted-pair, the bit-rate limit of a 10 m cable is around 1 Gb/s [83].

	On-Chip	Chip-to-Chip	Board-to-Board	Rack-to-Rack
Bandwidth		✓	✓	✓
Density		✓	✓	✓
Timing Skew	✓	✓	✓	✓
Crosstalk	✓	✓	✓	✓
Distributed Circuit Effects	✓	✓	✓	✓
Electromagnetic Compatibility & Interference		✓	✓	✓
Power Dissipation	✓	✓		
Simultaneous Switching Noise	✓			
Thermal Management	✓	✓	✓	✓
Partitioning	✓	✓		

Table 1.1: Motivation for optical interconnects at various length scales

Check marks indicate which issues are important in each level of the interconnect hierarchy. See text for details.

The bandwidth of an optical interconnect may be limited by dispersion, or by the optoelectronic devices used to transmit and receive the light. There is no dispersion in a free-space optical interconnect, and guided-wave interconnect at the chip-to-chip length scale are too short to be limited by dispersion. At the rack-to-rack level, if single-mode fiber is used, as in long-haul telecommunication, dispersion is again not an issue. If multi-mode fiber is used, the dispersion limited bandwidth would be around 200 MHz·km, or 2 GHz at 100 m¹. As for the optoelectronic devices, the limiting factor would be the laser. Bandwidth of a few GHz are routinely achieved in laser diodes, and compact, high efficiency vertical cavity laser diodes suitable for use in dense optical interconnects have been demonstrated with bandwidths of around 20 GHz [85].

Density

As pointed out above, the cross-sectional area of an electrical interconnect largely determines its performance. At the board-to-board and rack-to-rack length scales, the sheer bulk of cabling can be a serious practical performance limitation [39]. A related issue is the large size of the connectors needed to support this cabling [86,87]. Fiber-optic based interconnects have an important advantage in this regard. Optical fibers are very compact (typical fiber outer diameter is 125 μm, and there is usually an organic sheath bringing the total diameter to 250 μm) and flexible and may be densely packed in one- and two-dimensional arrays. Fiber optic connectors can also be quite compact, particularly in the case of parallel optical links using fiber ribbons.

Another aspect of optical interconnects leading to higher density at the chip-to-chip level stems from the ability to intersect optical paths. In a planar lightwave circuit, for instance, waveguides may intersect orthogonally without crosstalk. In free-space interconnects, optical paths can be crossed arbitrarily. Furthermore, the three-dimensional, rather than planar, connectivity of the free-space architecture may prove to be very useful². Aside from increasing chip-to-chip interconnect density, optical interconnects may provide a solution to the “escape” problem: when dense area-array flip-chip bonding is used to increase the $I/$

1. This multimode fiber bandwidth specification appears on distributor data sheets. However, 10 Gb/s data transmission using -10.5 dBm of 850 nm optical power over 125 m of standard 50 μm core multimode fiber has been demonstrated with a bit-error-rate of 10⁻¹² [84]. Evidently, with suitable design, multimode fiber may be used well beyond the manufacturer’s specification.

2. A special case of this is in interconnect intensive systems which contain some form of regularity in their connectivity. Important examples include network routing, the Viterbi decoder, and the FFT [88-91].

O density of a chip, trace congestion beneath the die makes it difficult and costly to route signal away from the IC [92].

Another density related benefit of optical interconnects appears at the on-chip level. Optical transmitters (laser driver and laser) and receivers (photodetector and receiver circuit) are typically much smaller than electrical output drivers and bondpads.

Timing skew

Timing skew refers to the difference between the arrival times of signals propagating through different paths leading to the same destination. It can occur as a result of variations in both the passive components (e.g. transmission lines) and active components (e.g. line drivers/buffers) encountered along the path, parasitic coupling to adjacent structures, and the effects of signal transition on adjacent interconnects. Timing skew is a significant problem at all levels of the interconnect hierarchy. At the on-chip level, clock distribution is a particularly menacing problem, often consuming large amounts of die area and power. At the chip-to-chip level, timing skew is the key factor limiting the off-chip clock speed¹ [82]. Board-to-board and rack-to-rack connections generally use differential signalling with twisted-pair or twin-axial cabling. The within-pair skew on high quality cable is typically 15-20 ps/m and at best around 6 ps/m. Within-pair skew limits the cable length to around 10 m at 1 Gb/s. Parallel operation is limited to around 1 m by the typical between-pair skew of 100 ps/m [83].

The delay along an optical interconnect can, in principle, be very well predicted and controlled because of the lack of parasitic interactions. This should be particularly true of free-space architectures. Waveguided systems will require good process control to minimize skew. As a point of comparison, standard 62.5 μm graded index core multimode fiber has an across-lot skew of 10 ps/m (an order of magnitude better than high-quality cabling) and a skew of 1.5 ps/m for parallel optical paths using fiber from same process lot [93]. The skew in optical links will likely to be limited by variations in the optoelectronic devices (e.g. laser threshold current variations) and the transmitter and receiver electronics. Minimizing these effects will be an important aspect of optical interconnect technology development.

1. While on-chip CPU clocks are approaching 1 GHz, off-chip data busses are stuck at ~ 100 MHz due to skew. The industry is moving toward serialization techniques to sidestep this problem, but this approach will not scale well as bus-widths and internal clock rates continue to grow.

Crosstalk

Crosstalk in electrical interconnects occurs because of parasitic capacitive, inductive, and radiative coupling. It is a significant problem at all levels of the interconnect hierarchy. Crosstalk increases as the transition time of an electrical signal decreases, and as a result, crosstalk poses a limit on the operating frequency of a digital system. In an optical interconnect, crosstalk is much less intrusive. It is virtually non-existent in a fiber coupled system. In planar-waveguide architectures, there may be crosstalk due to light scattering; this is limited by mechanical tolerances and is likely to be negligible. In free-space systems, crosstalk may result from diffraction. This is a predictable aspect of the design; for example, a free-space interconnect density limit of around 3500 connections/cm² has been predicted for one simple geometry [94]. Even when it does exist, optical crosstalk is bit-rate and pattern independent, greatly simplifying design and debug efforts.

Distributed circuit effects

An interconnect must be treated as a transmission line, rather than a lumped circuit element, when the signal propagation time, determined by the length of the interconnect and the speed of light in the dielectric material surrounding it, becomes comparable to the signal transition time. Besides determining the behavior of individual interconnects, distributed circuit effects are integral to supply and ground distribution and strongly impact crosstalk. Nearly every detail of a high speed digital design must account for distributed circuit effects [95,96]. This applies even to non-terminated on-chip interconnects exceeding a few millimeters [97].

Optical interconnects represent an extreme case of a distributed circuit, and ironically, this makes them simpler to design. The wavelength of the optical “carrier” is much shorter than the interconnect length, and the carrier frequency is much larger than the signal bandwidth. As a result, it is relatively easy to design for issues such as reflections and index matching (the optical analog of impedance matching), and there is no bit-rate or pattern dependence to these effects.

Electromagnetic compatibility (EMC) and electromagnetic interference (EMI)

Electromagnetic compatibility (EMC) is concerned with a system’s ability to function properly in a given electromagnetic environment. Electromagnetic interference (EMI) refers to emissions from a system which disrupt the normal operation of other devices [98]. Pursuant to federally mandated standards, EMC and EMI issues receive a great deal of attention in the design of commercial electronic equipment. EMI

problems grow more acute as electrical signal transition times are reduced and must be dealt with in the design of circuit boards, connectors, cabling, shielding, etc. [95]. Optical interconnects neither emit nor are sensitive to electromagnetic radiation. They provide a robust solution to many EMI/EMC problems.

EMC is also likely to prove to be a compelling reason to use optical interconnects in many lower-speed applications such as industrial control, satellites, planes/trains/automobiles, and high-end consumer audio equipment (in some existing consumer products, digital audio signals are transmitted optically between system components).

Power dissipation

A standard ECL output dissipates 20 mW (2 V supply, 1 V output swing, 50 Ω load, 50% duty cycle) in addition to the \sim 20 mW consumed by the output driver itself. An optical link consisting of a laser driver, laser, and receiver would be expected to dissipate around 10 mW (see Chapter 5, Chapter 6, Chapter 7), or one fourth that of the ECL connection. To see how this power savings scales at the system level, consider the example of an IBM AS/400 processor module consisting of 7 Bi-CMOS ICs and dissipating 220 W of power [99]. With around 2000 high-speed signals in this module, the use of optical interconnects at the chip-to-chip level would lead to a 60 W power savings, a 25% power reduction for the module. Similar arguments hold at the board-to-board and rack-to-rack levels, but with less significant savings since there are fewer interconnects.

At the on-chip level, interconnects are typically not terminated, but the need to drive large capacitive loads leads to large power consumption. Again taking CPU clock distribution as an example, the DEC Alpha microprocessor dissipates 17 W for this purpose. Optical clock distribution may require less than 7 W (using a 0.5 W off-chip source to drive 400 optical receivers spaced on a 500 μm grid across a 1 cm^2 die, with each receiver supplying the clock to an average of 5000 gates). In general, the use of on-chip optical interconnects at lengths as short as a few millimeters may provide power savings [71].

Simultaneous switching noise (SSN)

Because of the synchronous nature of common digital systems, supply current transients generated by output drivers add coherently. For example, when a 32-bit ECL bus switches from 00...0 to 11...1, the supply current jumps by 640 mA. If this happens in around 200 ps and there is around 1 nH of supply inductance, the supply voltage will fluctuate by over 3 V. Bond wires, package leads, and circuit-board vias have inductance exceeding 1 nH. Similar problems arise internal to the IC as well [100]. Clock and bus drivers

have rapidly changing current requirements. Left unchecked, the current trends suggest future supply current fluctuations of 50-100 A in each clock cycle [101]. In addition, large supply transients result when functional blocks on an IC are turned on or off as part of a power management strategy.

Flip-chip packaging helps to eliminate bond-wire inductance, but meticulous design, including the use of many power and ground connections (typically more than half the pads on a high speed IC are used for power and ground), are still required to ensure that simultaneous switching events can be tolerated. Optical interconnects can greatly reduce simultaneous switching problems by reducing current modulation requirements from 20 mA for a 50 Ω trace to 1 mA for a vertical cavity laser. Additional current transient and power savings could be gained by using optical interconnects for long on-chip connections. Optical interconnects may help in the power delivery problem by distributing the system over a larger area (see "Partitioning" below).

Thermal management

The removal of heat poses a difficult design challenge at every level of a high speed digital system. At the IC level, 150 W power dissipation levels are expected to pose unprecedented future packaging requirements [101]. A more insidious problem, though, is raised by the projection that by the year 2003, a typical 4-CPU server will exceed the thermal capacity of existing machine rooms built to meet long-standing standards; building new machine rooms will cost ~\$1000/sq. ft. [102]. Nearly every negative attribute associated with electrical interconnects scales with length leading to the tendency to compress ICs, boards, racks, and equipment clusters as much as possible. This runs counter to the need to distribute heat sources as widely as possible in order to remove heat and maintain the required ambient temperature. Since optical interconnect performance does not sharply degrade with increased length, they provide a means to spread out electronic systems in order to mitigate the thermal management dilemma.

Partitioning

The process of determining how much functionality is included on a given IC, versus how much is distributed across multiple ICs on a printed wiring board or multi-chip module, is referred to as partitioning. In general, there is an optimal choice which balances issues such as speed, size, and yield [103,104]. While issues such as power distribution, thermal management, and IC yield suggest the use of multiple, smaller die, the performance limitations of off-chip electrical interconnects have driven systems away from multi-chip partitioning. Fortunately, the exponential improvements made over the last thirty years in semicon-

ductor manufacturing technology have allowed ever-more functionality to be packed onto one IC, largely obviating the need for significant partitioning. However, by eliminating the off-chip interconnect bottleneck, the successful introduction of optical interconnects may lead to a more cost-optimal partitioning scheme.

Are optical interconnects justified?

The growth in performance of VLSI components and of voice/video/data communication systems has far outpaced improvements in electrical interconnect platforms. At length scales exceeding around 10 m, and in some cases at shorter board-to-board length scales, a number of the above arguments have already proven to be compelling. For instance, large telecommunications switching systems began using optical interconnects in the mid 1980's (AT&T Definity PBX and 5ESS central office switches; aggregate bandwidth 1-10 Gb/s). By the early 1990's optical interconnects were used at much higher densities to connect port interfaces to the core switching fabric (AT&T DACS VI-2000; aggregate bandwidth 10-100 Gb/s) [39-41]. The optical connections in these systems were based on discrete devices. The next generation telecommunication switches will have aggregate bandwidths of ~10 Tb/s and will make use of parallel optical interconnects implemented with arrays of vertical cavity lasers in highly-optimized hybrid modules (see architecture discussed above). As the cost of these modules drops, they will play an increasing role in computer systems as well. Some lower-cost/higher volume applications also currently use optical interconnects. For instance, optical interconnects play a key role in high performance RAID (redundant array of inexpensive drives) disk-drive systems. The optical implementation of the Gigabit Ethernet standard, which uses low-cost laser and fiber technology desirable for dense digital optical interconnects, has also become a commercial success. So, at the board-to-board and rack-to-rack levels of the interconnect hierarchy, it is quite safe to say that optical interconnects are justified.

At the chip-to-chip and on-chip length scales, there is currently no compelling justification for the use of optical interconnects. No insurmountable obstacle has yet been encountered in the use of electrical interconnects, and current engineering practice will continue to evolve to meet growing performance demands. Of course, this does not come without increased cost--a great deal of non-recurring engineering cost as well of some recurring cost associated with the gradual adoption of more sophisticated packaging technologies. How long this trend will continue is not certain. The growing use of optical interconnects at the board-to-board and rack-to-rack levels, and in local area networks, is expected to significantly drive

down the cost of core optoelectronic technologies. Whether and when the confluence of price and performance will favor the use of optical interconnects at the chip-to-chip and on-chip length scales is an open question.

Requirement: VLSI-complexity OEICs

Along with optoelectronic packaging, the development of optoelectronic integrated circuits (OEICs) which combine large numbers of optoelectronic devices with VLSI-complexity electronics are the key technological challenges associated with the implementation of optical interconnects. At the on-chip and chip-to-chip length scales, the realization of optical interconnects would require the close interaction of thousands of emitters and detectors with dense electronics. The board-to-board and rack-to-rack length scales, which currently use multi-component hybrid modules, OEICs would be expected to reduce cost and increase yield and reliability by simplifying packaging requirements, reducing part count and inventory overhead, reducing the number of critical alignments (since the integrated devices are automatically aligned to each other), and taking advantage of on-wafer testing [105]. The development of an optoelectronic VLSI (OE-VLSI) technology suitable for the implementation of optical interconnects in high performance computing and communication applications is the goal of this thesis.

1.2 OE-VLSI Technologies

Central to the development of an OE-VLSI technology is the choice of specific optoelectronic and electronic devices which may be harmoniously combined through a manufacturable process. This task is complicated by the incompatibilities between the III-V compound semiconductors which are commonly used for optoelectronic devices (GaAs, InP, etc.), and Si, the dominant electronic semiconductor material. The subtleties involved in the selection of the optoelectronic devices and of the electronics medium have resulted in the development of a number of different optoelectronic integration techniques based on either Si or GaAs electronics and using either a hybrid or monolithic integration technique. This section will explore a number of issues involved in choosing between GaAs and Si electronics and will review some of the most prominent optoelectronic integration technologies. The objective is not to compile a comprehensive picture of each approach, but rather, to place the work of this thesis, the Epitaxy-on-Electronics (EoE) process, in context within the spectrum of optoelectronic integration technologies. In addition, this discussion will serve to clarify the technical assumptions underlying the development of EoE, a monolithic opto-

electronic integration technology based on GaAs VLSI.

Optoelectronic integration technologies may be categorized according to the host substrate, either Si or III-V, and according to whether a monolithic or a hybrid approach is taken. Of the four possible combinations, hybrid integration on a III-V substrate has received very little attention while the others have been studied extensively. Integration on a Si substrate is desirable because of the overwhelming success of CMOS VLSI electronics. Work with III-V substrates has focused primarily on GaAs. The trade-offs between the two electronics platforms are more subtle than commonly believed, and will be addressed shortly. The choice of monolithic or hybrid integration has significant ramifications with regard to manufacturability, as will be pointed out below for each of the technologies.

Monolithic integration on Si

Various means are being pursued to develop optoelectronic devices based on Si, including Si/SiGe and erbium-doped Si [106]. Neither emitters nor modulators of adequate performance for use in optical interconnects have been realized. At this point in time, optical interconnects continue to rely on III-V optoelectronic devices. Monolithic integration of III-V devices with Si electronics thus depends on the growth of III-V materials on a Si substrate.

The lattice-constant mismatch between Si and important III-V compound semiconductors has made the growth of high quality heterostructures on Si difficult, and the large thermal expansion coefficient mismatch between Si ($2.56 \times 10^{-6}/^{\circ}\text{C}$) and III-V materials (e.g. $6.86 \times 10^{-6}/^{\circ}\text{C}$ for GaAs) has proven to be an even more intractable problem. Significant advances have been made in the growth of novel buffer layers between the Si substrate and the active III-V heterostructure, and laser diodes operating for over 2000 hours have been reported [107]. Nonetheless, much more material- and device-level research will be required to produce practical devices.

A recently introduced integration technique seeks to reconcile the material incompatibility between Si and GaAs by taking advantage of the resiliency of CMOS electronics fabricated in a thin layer of Si [108-110]. For instance, the devices in established Si-on-sapphire processes are subject to large thermal expansion stresses, with the thermal expansion coefficient of sapphire nearly equal to that of GaAs (the bi-anisotropic coefficients of thermal expansion has values of $5.0 \times 10^{-6}/^{\circ}\text{C}$ and $6.5 \times 10^{-6}/^{\circ}\text{C}$ in each of two orthogonal directions). To facilitate monolithic optoelectronic integration, the proposed technique uses

wafer-bonding to form a Si-on-insulator substrate in which the bulk substrate is GaAs rather than Si or sapphire. Efforts to implement OEIC using this strategy are ongoing.

Hybrid integration on Si

The two principle techniques in this category are epitaxial lift-off and flip-chip bonding. In epitaxial lift-off, the required optoelectronic device layers are removed from their native substrate using a selective etch and then aligned and grafted onto a completed Si die [111]. Using this technique, integration of up to an 8x8 array of LEDs has been reported [112]. The manufacturability of this technique has not been addressed.

The other common hybridization technique is flip-chip bonding. The basic technique of flip-chip bonding has been used for over thirty years as an alternative to the wire bonding of integrated circuits. In the basic technique, solder balls are applied to the IC bondpads, and the device is brought into rough alignment with a package. The solder is then reflowed by heating, and the IC and package are pulled into proper alignment by the surface tension of the molten solder. While this basic technique is very well suited to the packaging of electronic die, two significant modifications have been made in applying it to OEIC fabrication [111,113,114]. First, suitable bondpads must be prepared on the optoelectronic sample. As will be described further below, common optoelectronic devices are highly non-planar. Producing contacts to both sides of the optoelectronic device is a challenging fabrication problem, particularly given the high device densities that are desired and the bondpad co-planarity required for successful flip-chip bonding. Processes involving such a large degree of non-planarity have historically suffered severe yield limitations. Second, unlike conventional flip-chip packaging, the substrates of the optoelectronic die which are flip-chip bonded to form OEICs are typically removed, leaving individual free-standing devices. This is done in order to address the issues arising from the large thermal expansion mismatch between Si and GaAs. However, it raises important manufacturability and reliability questions that have not been addressed.

The major use of flip-chip bonding in OEIC technology has focused on multiple quantum-well modulators [37]. OEICs containing 144 optical inputs and outputs on a 2 mm x 2 mm die have been fabricated, corresponding to 3600 optoelectronic devices/cm² [111]. In spite of its relative maturity, scaling of this technology to full-wafer production has not been reported. Furthermore, although the availability of flip-chip modulator-based OEIC has spurred a great deal of optical interconnect research, the future of this field is widely believed to belong to emitter-based OEICs. In particular, vertical cavity surface emitting lasers

(VCSELs) have been compared, with respect to their importance in OEICs, to CMOS electronics and their critical role in electronic VLSI technology [33]. Extension of the flip-chip OEIC process to VCSEL integration is in progress. Initial results have shown a ~60% yield, which is currently limited by device breakage during handling, poor substrate removal, and poor electrical contact through the flip-chip bond [113,114]. Another report has addressed the concerns of planarity and substrate removal by introducing a novel VCSEL structure which minimized the level of non-planarity and by leaving the GaAs substrate intact following the flip-chip bonding process [115-117]. Due to the thermal expansion coefficient mismatch between GaAs and Si, the choice to retain the GaAs substrate has limited the size of the array which may be bonded with high yield (a 2x10 array was reported) and raises long-term reliability questions.

Monolithic integration on GaAs

A more traditional approach to optoelectronic integration has been to develop a growth and fabrication process which produces both transistors and optoelectronic devices on a common GaAs wafer. Two difficulties arise in this approach. One is the sizable investment of time and capital required to develop a VLSI electronics process. Such an investment is not easily justified without the clear existence of a profitable market. The other difficulty is the inherent height mismatch between optoelectronic devices and the field-effect transistors (FETs) used in VLSI electronics. FETs gain much of their manufacturability from their high degree of planarity which allows fine features to be photolithographically projected over them with excellent fidelity. Optoelectronic devices, however, are typically composed of numerous layers of epitaxial material, and the finished device structures may impose several microns of surface variations. (One important exception is the metal-semiconductor-metal (MSM) photodetector which exhibits superb planarity and has been incorporated into VLSI circuits.) This non-planarity disrupts basic elements of the fabrication process such as photoresist surface coverage and may also exceed the focal depth of the optical systems required to project the fine patterns needed for VLSI devices. Given these two barriers, extensive work on traditional monolithic optoelectronic integration has resulted in only small-scale demonstrations such as basic combinations of a single VCSEL, MSM photodetector, and MESFET [118], or the incorporation of a small number of in-plane lasers into an existing microwave integrated circuit process [119].

Epitaxy-on-Electronics

The barriers to conventional monolithic optoelectronic integration are circumvented in the Epitaxy-on-Electronics (EoE) process. EoE produces monolithic OEICs on a GaAs substrate. But, unlike the traditional approaches, EoE builds onto an established VLSI process to avoid the massive investment required to build complex electronics. Furthermore, EoE is inherently compatible with the height requirement of common optoelectronic devices. Both the electronic and optoelectronic fabrication processes begin with a planar surface, and thus, not only is the highly-evolved VLSI technology unaffected by the incorporation of the new devices, but the optoelectronic processing is equivalent to that of similar devices on bulk wafers.

It is apparent from the brief discussion above that monolithic-on-Si and traditional monolithic-on-GaAs technologies will not, in the immediate future, provide the complex OEICs required to implement optical interconnects. Two viable choices, however, are a hybrid-on-Si approach, such as flip-chip bonding, and Epitaxy-on-Electronics. The advantage of EoE lies in its use of monolithic integration, while flip-chip bonding succeeds in incorporating Si VLSI electronics. It is difficult to challenge the notion that monolithic integration leads to a superior level of manufacturability, reliability, density, and performance than hybrid integration. Were this not the case, the need for the continuing escalation of integration density would not be necessary. The manufacturability advantage of EoE must, then, be weighed against the advantage of using Si-CMOS rather than GaAs VLSI electronics. The trade-offs between using Si or GaAs VLSI electronics must be studied carefully.

Si and GaAs Electronics

Si-CMOS ICs account for an overwhelming majority of microelectronic components. The tremendous commercial success of Si VLSI, and the somewhat tumultuous history of digital GaAs electronics, has made CMOS the *only* choice for OEIC development in the minds of many researchers. Such a judgement is, perhaps, too hasty. CMOS offers significant power and cost advantages for applications operating below several-100 MHz and for memory-intensive applications, including microprocessors which may dedicate over forty percent of their die area to cache memory. CMOS has a dramatic advantage in static-memory intensive circuits since memory cells which are not being accessed dissipate little power. However, in non-memory-intensive architectures, such as switching and serialization/deserialization, direct-coupled FET logic (DCFL) electronics using enhancement- and depletion-mode GaAs MESFETs achieve higher speeds at lower levels of power consumption [120].

DCFL's high-speed power savings, along with the advantages of reduced crosstalk, simplified circuit design (to meet high-speed targets), and a streamlined process (rather than the twenty or more steps of a CMOS process, only thirteen mask steps are used in the four metallization layer H-GaAs III process and fifteen in H-GaAs IV process), have made GaAs VLSI an effective economical choice for high-speed logic circuits [13,120]. The most direct evidence of this fact is the dominance of GaAs VLSI in high-speed digital applications. This market consists primarily of tele- and data-communication equipment and integrated circuit test equipment. The major participants in these markets rely on GaAs VLSI electronics for the high-speed components of their products¹. This does not contradict the fact that CMOS is, rightfully, the mainstay of the microelectronics industry. While high-speed digital electronics play a pivotal and profitable role in the modern information infrastructure, their volume is a small fraction of the total IC production which includes microprocessors, memory, and a multitude of lower-speed, very-high-volume devices which control everything from smoke detectors to stereos.

The real question that must be addressed in selecting Si or GaAs electronics is the nature of the applications which require optical interconnects. As summarized in Table 1.2, the most immediate demand for optical interconnects occurs in the same markets which currently rely on GaAs VLSI electronics: tele- and data-communication systems. Tera-bit/s telecommunication switches using free-space optical interconnects and Gb/s per channel board-to-board/rack-to-rack links using parallel optical fiber ribbons are two notable examples. It is through such applications that optical interconnects will gain acceptance in the broader electronics community. For applications which do require Si electronics, such as chip-to-chip connections between microprocessors and memory, optical interconnects will not become competitive with existing electrical interconnects for a number technology generations [67]. The challenge to the optical interconnect research community is thus two-fold. First, optical interconnects for high-speed digital application must be commercialized expediently and economically. Simultaneously, development of technologies to combine III-V optoelectronics with Si-CMOS electronics must be continued in order to meet the eventual demand for optically interconnected microprocessors.

1. A partial list of tele-communication and data-communication equipment suppliers which use Vitesse GaAs VLSI components includes: Lucent, Alcatel, Ericsson, NEC, Fujitsu, Hitachi, Tellabs, Stratacom, BNR, Seagate, Newbridge, Mitsubishi, Cisco, and Cienna. In the automated IC test market, Vitesse's customers include: Schlumberger, Taradyne, Credence, Anritsu, Ando, HP, LTX, IMS, Texas Instruments, Minato, and Mistubishi [121].

Optical Interconnect Application	Electronic Technology	Time-Scale
Optically-interconnected Computer <ul style="list-style-type: none"> • Synthetic Aperture RADAR • Servers and Scientific Computers 	Si	5 yrs. (military) 10-15 yrs. (commercial)
“Smart Pixel Arrays” <ul style="list-style-type: none"> • Switching/Routing (e.g. ATM) • Error-correction (e.g. Viterbi) • Encryption/Decryption 	GaAs	3-5 yrs.
Fiber-optic Data Links <ul style="list-style-type: none"> • Single-Channel Transceivers (e.g. Gigabit Ethernet, 10-Gigabit Ethernet, FDDI, Fiber Channel) • Parallel Optical Link 	GaAs	Now

Table 1.2: Optical interconnect applications and electronic platform technologies

The choice of the electronic circuit platform used for optoelectronic integration depends on intended application. Optical interconnect applications may be categorized as shown above. Each type of application will become commercially viable on a different time scale. In an optically-interconnected computer, processing and memory building blocks may be optically linked to form high performance parallel architectures. For these computers to leverage existing CMOS microprocessor technology, the electronics process underlying the OEICs must be Si-based. Existing electrical interconnect methodologies are likely to continue to dominate these markets for several technology generations. A second application category is the so-called “smart-pixel array.” Components in this category manipulate massive data streams (>1 Tb/s) and will be needed in the telecommunication industry in the near future. These applications are currently dominated by GaAs VLSI electronics, a trend which is likely to continue irrespective of the development of optical interconnects. The third category consist of fiber-optic local-area network and rack-to-rack interconnects. These types of optical interconnects are already in commercial use, and their volume is growing rapidly. At present, transmitter and receiver functions are implemented as hybrid modules which incorporate III-V optoelectronic devices with both GaAs VLSI and Si bipolar components. These applications may benefit substantially from the use of monolithic OEICs based on GaAs.

The development of Epitaxy-on-Electronics has been based on the issues discussed above: (1) monolithic integration is the key to achieving low cost and high yield, density, and reliability in a manufacturing environment, (2) III-V optoelectronic devices are needed, (3) the fabrication process must allow for the large thickness of common optoelectronic devices, and (4) the use of GaAs VLSI electronics is a reasonable economic and technical choice for the high-speed digital applications which will require optical interconnects for the foreseeable future.

1.3 Thesis Summary

The topic of this thesis is the development of a monolithic optoelectronic VLSI technology. The work has encompassed aspects of device-, process-, circuit-, and system-level engineering. This wholistic approach is appropriate at the early stage of development because of the close interdependency of each distinct facet. The significant progress which has been made in each area has resulted in this document's large size. The technology will benefit in the future from more specialized research.

Chapter 2 describes the Epitaxy-on-Electronics (EoE) integration process. This work builds on that of Shenoy [1-5], which provided a proof-of-concept for the EoE technique. Although the framework of the process remains intact, each element of it has been substantively altered. The revisions and additions were necessary for the implementation of user-ready OEICs.

The demonstration of the first VLSI-complexity, monolithic, emitter-based OEIC is presented in Chapter 3. The EoE process was recently used to fabricate LED-based OEICs as part of a research foundry project known as OPTOCHIP. The goal of the research foundry was to catalyze progress in optical interconnects by making sophisticated OEICs available to systems researchers. Simultaneously, close interaction with the designers has been used to guide further technology development. Chapter 3 will describe the OPTOCHIP project and detail the design, fabrication, and characterization of the OPTOCHIP LEDs and a number of optoelectronic circuits into which they have been incorporated.

The experience of the OPTOCHIP project provided much of the direction for the remainder of the thesis. The key objectives which were identified include the introduction of integrated vertical cavity surface emitting lasers (VCSELs) in place of the LEDs, improvements in the integrated photodetectors which are implemented as part of the VLSI electronics process, and the design of high speed circuits to interface between the digital electronics and the optical emitters and detectors. This thesis addresses the photodetector and interface circuit objectives.

Chapter 4 deals with the implementation of metal-semiconductor-metal (MSM) photodetectors in the standard, or slightly modified, GaAs VLSI MESFET process from Vitesse Semiconductor Corporation [26]. Two versions of these devices were characterized and their shortcomings identified. Potential directions for future photodetector efforts are examined.

System-level aspects of the OEIC technology are addressed in Chapter 5. A prototypical optical interconnect application is used to formulate guidelines for the interface circuits which, in conjunction with the

optoelectronic devices, determine the behavior of the digital optical interconnects. The structure of the digital optical interconnect is examined, particularly in comparison with established optical communication systems. The performance of the digital optical interconnect, as manifest in its error probability and jitter characteristics, is formulated analytically in order to understand the role of the various physical parameters.

The guidelines of Chapter 5 are applied in Chapter 6 to the design of a laser driver and in Chapter 7 to an optical receiver. In each case, the circuits are expected to support 1 Gb/s operation in a VLSI environment while consuming minimal power. The laser driver is thoroughly characterized, the limitations of the current design identified, and improved designs examined. In the case of the receiver, the characterization was limited by the performance of the integrated photodetector examined in Chapter 4. Additional receiver designs appear in Appendix F, Appendix H, and Appendix I. The first of these proceeded and influenced the design of Chapter 7. The second was optimized for use in a stand-alone receiver array OEIC. The receivers of Chapter 7 and Appendix H are used in each of two such array chips which are documented in Appendix G. The final design, Appendix I, is a clocked receiver which provides a valuable power consumption and area reduction in systems where its use is appropriate.

Chapter 8 summarizes the accomplishments of the thesis and highlights some directions for future work.

2

The Epitaxy-on-Electronics Integration Process

The EoE process is outlined in Figure 2.1. It begins with a commercially-fabricated, fully-metallized GaAs VLSI integrated circuit featuring enhancement- and depletion-mode MESFETs and four levels of aluminum-based interconnects. Metal-semiconductor-metal (MSM) and optically sensitive MESFET (OPFET) photodetectors can also be implemented in the standard electronics process. The unrestricted placement of the LEDs or other heterostructure optoelectronic devices occurs, along with the electronics, as part of the routine layout of the IC. In the regions designated for the optoelectronic devices, the SiO₂-based interconnect dielectric stack is removed to reveal the underlying GaAs substrate. These regions, known as dielectric growth windows (DGWs), expose the substrate material for subsequent epitaxial growth using molecular beam epitaxy (MBE). Figure 2.1(a) shows the IC just prior to MBE growth. The n-type source/drain implant is used as the bottom n-contact to the optoelectronic device. Using solid- or gas-source MBE, epitaxial material is grown in the DGWs while polycrystalline material is deposited on the over-glass, as shown in Figure 2.1(b). Standard processing methods are then used to remove the polycrystalline deposits, fabricate the optoelectronic devices, and interconnect them with the electronics through top-side electrical contacts. Figure 2.1(c) illustrates a completed OEIC containing, in this case, LEDs.

EoE is a monolithic integration technique which uses established fabrication processes to facilitate scaling to full-wafer production with high yield and reliability. By building on fully-metallized GaAs VLSI electronics from a commercial manufacturer, in this case Vitesse Semiconductor Corporation [26], the massive up-front investment in developing a VLSI process is avoided. However, the thermal stability of the completed electronics now constrains the integration process. As summarized in Section 2.1, previous

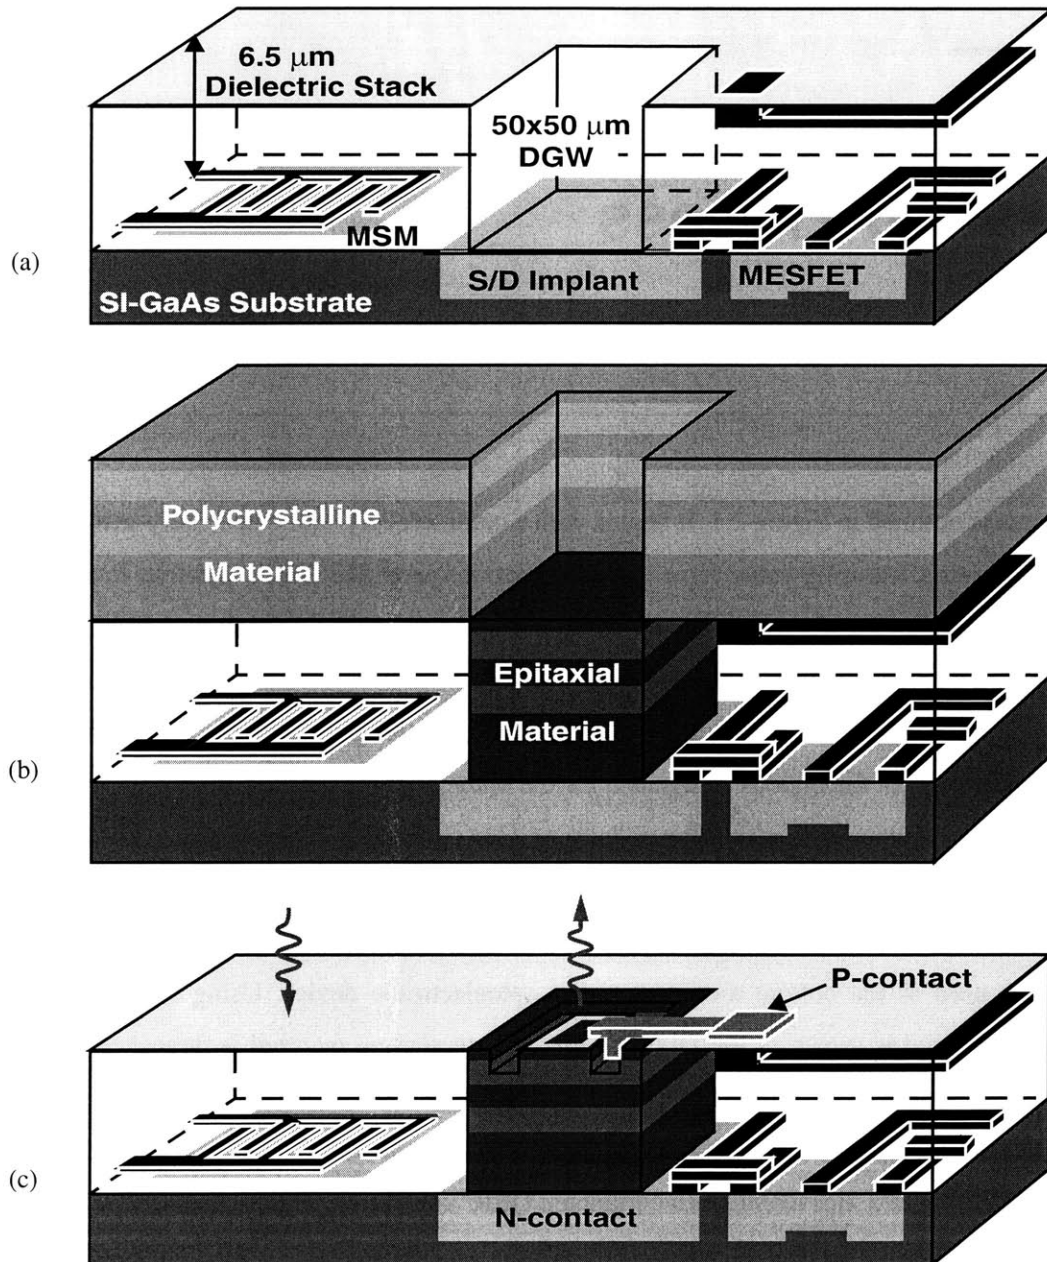


Figure 2.1: Overview of the Epitaxy-on-Electronics integration process

The EoE process monolithically integrates optoelectronic devices with GaAs VLSI electronics. (a) Dielectric growth windows (DGWs) are cut into fully-metallized die or wafers to expose regions of the GaAs substrate. (b) Conventional MBE is used to epitaxially grow optoelectronic heterostructures in the DGWs while polycrystalline material is left on the overglass. Standard techniques are then used to remove the polycrystalline deposits and fabricate the integrated optoelectronic devices. (c) An integrated LED.

studies of Vitesse self-aligned GaAs MESFET circuits, which use refractory gate and ohmic metallization and aluminum-based interconnect metallization, have demonstrated that five-hour temperature cycles of up to 475°C do not degrade the electronics [3,6,7]. Similar thermal stability characteristics have also been observed for GaAs VLSI MESFET ICs produced by Motorola [7]. At higher temperatures, the interconnect resistance increases as a result of a metallurgical reaction between the AlCu_x core and the WN_x claddings that make up the interconnect metallization. In extreme cases, such as temperature excursions to ~580°C, the reacted interconnect metals lose their malleability, preventing wire-bond adhesion.

Most optoelectronic devices require epitaxial growth of materials such as $\text{In}_x\text{Ga}_{1-x}\text{As}_y\text{P}_{1-y}$ and $\text{Al}_x\text{Ga}_{1-x}\text{As}$ using molecular beam epitaxy or metalorganic chemical vapor deposition (MOCVD). The typical substrate temperature during MOCVD growth is in the range of 600°C-850°C while MBE growth uses substrate temperatures of 450°C-700°C. The coincidence of the MBE substrate temperature range with the GaAs VLSI circuit thermal stability limit is an important aspect of EoE integration. Much of the EoE technology development work up to this point has focused on realizing MBE growth procedures which stay within the sub-475°C EoE temperature constraint. Full thermal compatibility of the EoE process with the GaAs VLSI technology is needed to preserve the high performance of the electronics. As importantly, it allows the device models and design tools which already exist for the well-established electronics process to be used in optoelectronic circuit design.

Two advances have led to full compatibility with the EoE temperature-cycle limitation: first, the use of atomic hydrogen produced by a high temperature cracker for the pre-growth removal of the native GaAs oxide, without conventional thermal desorption at ~580°C, and second, the use of the aluminum-free InGaAsP material system to allow sub-475°C growth of high quality optoelectronic device material. In addition to the thermal cycle limitation, the unique aspects of EoE integration include the formation of dielectric growth windows (DGWs) and the removal of polycrystalline deposits following growth. A detailed description of the current EoE process can be found in Section 2.2. This section also describes the early evolution of the EoE process in order to clarify the reasoning that has gone into its development and to point out the significance of the key process innovations.

2.1 Thermal Stability

Before detailing the EoE process in Section 2.2, the current section will describe the thermal stability limitations of the VLSI electronics which constrain the material growth process used for EoE integration. Figure 2.2 points out the various elements which make up a Vitesse GaAs MESFET VLSI circuit, a more detailed description of which can be found in Appendix A. As with silicon CMOS, the achievement of VLSI level integration in GaAs MESFET technology depends on the self-alignment of source and drain implants to the gate. This means that the gate material must be able to withstand the high temperature anneal (800°C for one hour in the case of the Vitesse process) that is required to activate the implant. In CMOS, this is accomplished by using polycrystalline silicon for the gate. In commercial self-aligned MESFET processes, high temperature stability of the Schottky gate is attained by using a refractory metal




Figure 2.2: Metallization on Vitesse H-GaAs III and H-GaAs IV processes

Vitesse's self-aligned MESFET process uses a refractory WN_x gate metal, Ni-Ge ohmic contacts, and $Al_{0.99}Cu_{0.01}$ interconnect metallization sandwiched between WN_x barrier layers. The gate metal withstands a high temperature (800°C) activation anneal and the ohmic contacts are sintered at 550°C. The upper metallization layers are exposed to temperatures up to 380°C during the PECVD deposition of the SiO_2 dielectric stack. See Appendix A for details.

based contact, such as $W_{1-x}N_x$. The fact that the gate can withstand an extended exposure to such high temperatures inspired investigation into the thermal stability of completed ICs [3]. These studies established the possibility of MBE growth in foundry opened epitaxy windows on fully processed chips.

Thermal Stability Results

Investigation of the thermal stability of Vitesse GaAs MESFET VLSI ICs by Braun, *et. al.*, has led to an understanding of the degradation mechanisms and illuminated the boundaries of the time/temperature “growth envelope” [6,7]. Thermal cycles were performed on electrical test structures located on Vitesse H-GaAs III process control monitors. As summarized in Appendix A, the Vitesse process includes Schottky contacts for MEFSET gates, ohmic contacts to the source/drain regions, and “upper level” (metal 1 through metal 3 on the chips tested) metal interconnects. The Schottky gate contact itself was found to be stable for all thermal cycles [7], as originally expected. However, the upper level metal sheet resistances and ohmic contact resistances were found to degrade as the result of thermally activated metallurgical reactions [6].

Five hour thermal cycles, corresponding to a typical growth time for integrated structures, were carried out at various temperatures. Figure 2.3 shows the gate, ohmic, and interconnect metal sheet resistances as a function of anneal temperature. The gate and ohmic metal layers displayed almost no sheet resistance change even after 5 hours at 600°C. However, between 400°C and 600°C, the aluminum-based metal 1 through 3 layers increased significantly in sheet resistance. This resistance increase is believed to be the result of a metallurgical reaction between the $AlCu_x$ core and the WN_x cladding. Increasing the nitrogen content in the WN_x layers produced interconnects that did not exhibit an increase in sheet resistance when annealed below 550°C for five hours [6]. However, increasing the nitrogen content also lead to an unacceptable increase in interlayer contact resistance, precluding its inclusion in the standard Vitesse process.

The metal 1 sheet resistance increased linearly in time, the underlying reaction being thermally activated. The sheet resistance increase saturated at approximately 15 times the non-annealed value. Using an Arrhenius based model, shown in Figure 2.4, a set of theoretical curves was fitted to the measurements of metal 1 sheet resistance for arbitrary times in the temperature range of 400°C to 600°C [6].

The ohmic contact resistance increase was also thermally activated and progressed linearly in time, but the onset of the resistance increase was delayed in time. This delay was found to depend strongly on the distance between the metal 1 via edge and the ohmic contact metal edge. Figure 2.5 shows the average




Figure 2.3: Vitesse interconnect sheet resistance after 5 hour thermal cycles

The ohmic and gate metallization are found to be stable at up to 600°C while the interconnect metallization sheet resistance increases when exposed to temperatures in excess of 475 °C [6].

measured MESFET source resistance and small and large transmission line model (TLM) resistances per unit contact length as a function of anneal temperature. (In the figure, the “critical temperature” is defined as the temperature where the resistance per unit contact length exceeded 2000 ohm- μm .) A reaction between the Al-Cu-W-N interconnect and Ni-Ge contact is the cause of the ohmic contact resistance increase. The time delay is associated with the diffusion of the aluminum compound to the ohmic contact metal edge through the WN_x barrier. The barrier consists of the metal 1 bottom cladding layer, the ohmic metal top layer, and the via top spacing layer, which explains the correlation between the ohmic contact to metal 1 via spacing and the delay of the onset of the resistance increase [6].

As with the interconnect metal sheet resistance, the strong temperature dependence of the ohmic contact resistance can be modeled as a thermally activated process. Figure 2.6 gives a set of curves suitable for approximating this resistance for arbitrary thermal cycle temperatures and durations [6].

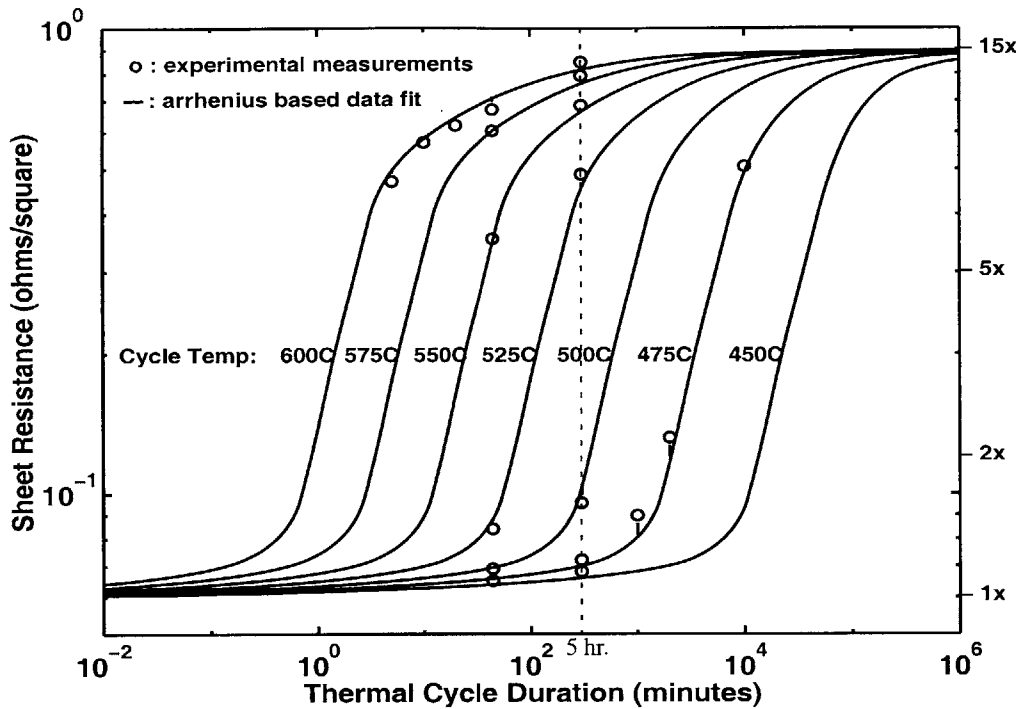


Figure 2.4: Metal 1 sheet resistance thermal stability model

The Arrhenius model based on the measured points indicates an activation energy of 3.5 eV for the degradation of the metal 1 sheet resistance[6].

Finally, in addition to systematic measurements made on individual MESFETs and interconnect test structures, the effects of the thermal cycles on actual circuit performance were tested on 23-stage direct-coupled field effect transistor logic (DCFL) ring oscillators. Oscillation frequency was unchanged for anneals up to 450°C. Between 450°C and 540°C, the ring oscillator period increased by 10%. No oscillations were observed for anneals at or above 550°C. Using the measured ohmic contact resistances, HSPICE simulations of the ring oscillator correctly predicted that the circuit would not oscillate after a 540°C anneal [6].

Thermal Budget

The importance of the thermal stability results to the present work is in establishing safe time/temperature exposure limits for Vitesse ICs used in EoE MBE growth. Figure 2.4 and Figure 2.6 indicate that a thermal cycle of 5 hours at 475°C will not compromise electronic circuit performance. A substrate temperature of 470°C has been chosen for the growth of the aluminum-free InGaP material to be discussed in the

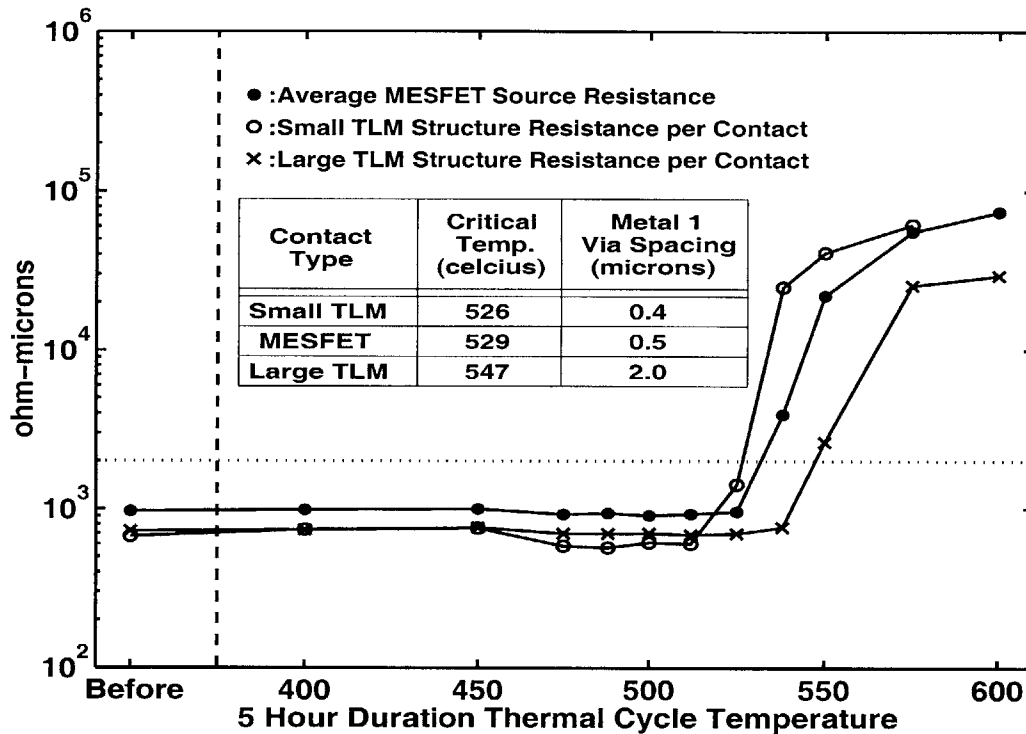


Figure 2.5: Ohmic contact resistance after 5 hour thermal cycles

Ohmic contacts on MESFET sources and on large and small TLM structures were studied. The inset gives the 2000 Ω - μm critical temperature and metal 1 via spacing for each ohmic contact structure. [6]

next section. It is also important to note that even a brief exposure to temperatures near 600°C causes significant degradation of the interconnects, including a loss of malleability which inhibits bond-wire adhesion. This fact has led to the development of a low-temperature process, to be described below, for the pre-growth removal of the GaAs native oxide.

2.2 Historical Development

The first Epitaxy-on-Electronics demonstration, by Shenoy, *et. al.*, used a process control monitor from the Vitesse GaAs foundry [4]. A portion of the monitor sample containing FETs was masked with wax, and hydrofluoric acid was used to completely remove the dielectric stack in the unmasked region. GaAs/AlGaAs LED material was grown on the exposed semi-insulating GaAs substrate using solid source molecular beam epitaxy (MBE). The growth was carried out at 530 °C, preceded by a conventional temper-

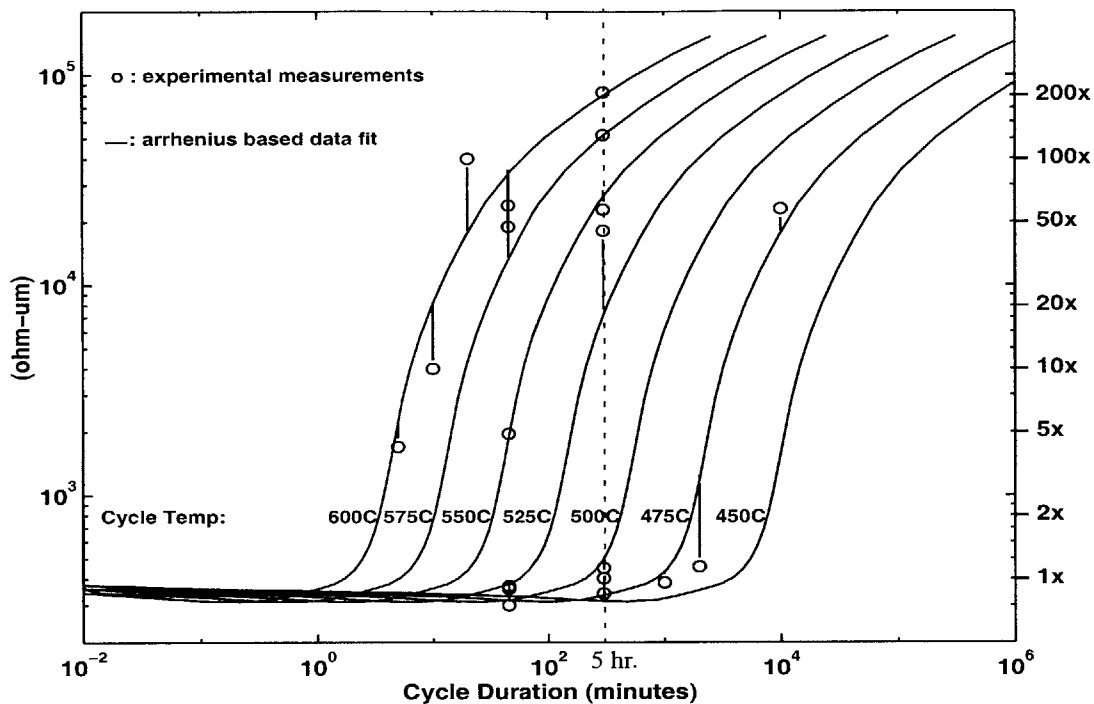


Figure 2.6: Metal 1-ohmic metal-implant contact resistance thermal stability model

For the composite structure formed of metal 1 and ohmic metal atop a source/drain implant, the Arrhenius based model fit to the measured contact resistances indicates an activation energy of 2.5 eV to 2.8 eV for the degradation process [7].

ature excursion to 580°C to desorb the native GaAs oxide. The 530°C growth temperature was selected on the basis of the thermal stability data available at the time [3]. The polycrystalline material deposited on the “circuit side” of the sample was wet-etched with the “epi-side” masked with photoresist. LEDs, with top side contacts to both the p- and n-type layers, were then fabricated from the epitaxial material.

This demonstration established the general practices of surface preparation, reduced-temperature MBE growth, post-growth polycrystalline deposit removal, and optoelectronic device fabrication that make up the EoE process. Each of these steps required major revision. Due to the reduced growth temperature, the LEDs on both the chip and on a simultaneously-grown bulk substrate had a low efficiency of ~0.03%. However, this 530°C growth temperature, as well as the 580°C oxide desorption temperature, were still high enough to degrade the electronics. Furthermore, the optoelectronic and electronic devices were not in close proximity.

The next phase of EoE development produced an OEIC with closely coupled LEDs and electronics [5]. Rather than crudely exposing a large region of the GaAs substrate on the chip, device-sized regions, referred to as dielectric growth windows, were etched using $\text{CHF}_3/\text{CF}_4/\text{He}$ reactive ion etching (RIE) at the Vitesse foundry. The DGW patterns were defined on a standard mask layer as part of the circuit layout. Furthermore, the regions of the substrate beneath the DGWs were designated as source/drain n^+ implants and were electrically contacted using standard Vitesse ohmic contacts. Thus, the bottom n-type contacts to the LEDs were made automatically, and only a top p-type contacts to the devices remained to be fabricated. This innovation of growth within an implanted DGW is the key to overcoming the non-planarity problems which have long plagued traditional monolithic optoelectronic integration processes.

Still using 580°C thermal desorption of the GaAs native oxide and growing 0.03% efficient GaAs/AlGaAs LEDs at 530°C , EoE produced functional winner-take-all optical neural network circuits [5]. It still remained, however, to bring the growth cycle into compliance with the electronic thermal stability limit which was subsequently established at 475°C for a five-hour growth [6].

Also, in the author's first attempt at EoE integration of Al-free LEDs, a problem was identified with the DGW preparation technique [10]. The foundry RIE used to produce the DGWs left fluorinated hydrocarbon residues on the substrate which slowly reacted with the GaAs.¹

The work in [1-7] established a proof-of-concept for the EoE process, including an understanding of the interconnect thermal stability constraints and of the basic EoE process. The task of the current thesis has been to develop the basic process into a practical optoelectronic integration technology. This task has required the modification of every aspect of the process and the addition of some new practices, including 1) a new DGW formation procedure, 2) the introduction of a low-temperature GaAs native oxide removal technique, 3) the use of Al-free materials, and 4) revised fabrication processes to accompany the new materials and to improve the process robustness.

In the remainder of this chapter, each step of the EoE process will be detailed. The specifics of the fabrication procedure for integrated LEDs will be deferred to Chapter 3. The starting point for the integration process is fully-metallized die received from the MOSIS service [122] and fabricated by Vitesse using their standard H-GaAs III or H-GaAs IV process. The foundry's passivation etch, which normally exposes the

1. In addition to the generation of fluorinated hydrocarbon residues, the RIE subjected the GaAs substrate to energetic ions, potentially generating surface damage which would seed defects in the epitaxial layers. Although no direct evidence of this had been observed, the possibility was a source of concern.




Figure 2.7: The dielectric stack on Vitesse integrated circuits.

Regions intended for growth are designated as source/drain implants bordered by ohmic contacts. A 30 nm Si_3N_4 layer covers the GaAs substrate in the source/drain regions while two such layers, separated by a thin SiO_2 layer, cover the remaining, non-active, surface of the substrate. PECVD SiO_2 and spin-on glass make up the remainder of the 6.5 μm thick dielectric stack.

bondpads, is excluded in order to protect the bondpads during subsequent processing. The die used in this report were approximately 1 cm square. Although the cost of full wafers or larger die is prohibitively high for a university research program, all of the EoE process steps are fully compatible with full-wafer processing, and in many instances using full wafers would simplify handling and would be expected to increase yield.

2.3 Dielectric Growth Window Preparation

The interconnect dielectric stack is pictured in Figure 2.7 and is described more thoroughly in Appendix A. It is roughly 6.5 μm thick, the bulk of which consists of layers of PECVD (plasma-enhanced chemical-vapor deposition) SiO_2 interleaved with layers of SiO_2 -like spin-on glass (SOG). The SOG layers are used to improve surface planarity following each step of metallization. The cured SOG may be wet- or dry-etched using the same chemistries used for the deposited SiO_2 , although the etch rate is typically higher for the SOG. *In situ* mass spectroscopy within the MBE chamber indicates that the SOG does not outgas during growth [123]. At the base of the dielectric stack is a thin layer (~30 nm) of PECVD Si_3N_4 over the

MESFET active regions. Over the non-implanted areas there are two such layers, but these regions are not used for DGWs. The Si_3N_4 is deposited directly over the GaAs substrate in order to prevent the loss of As during the 800°C implant activation anneal. During this process, the Si_3N_4 is believed to incorporate some Ga and As atoms and is densified by the high temperature process. As a result, the etch rate of the Si_3N_4 layer may be substantially higher than a typical, non-annealed film (the etch rate of which, using buffered hydrofluoric acid, is roughly 10 nm/min, or about one tenth that of SiO_2 .) Significant variations in this etch rate have been observed among the various Vitesse ICs used throughout this project.

Because of the large thickness and isotropic nature of the dielectric layers, and due to the etch-rate difference between the PECVD SiO_2 and the SOG, a directional reactive-ion etch (RIE) is required to form the DGWs. The approach taken in the original EoE demonstration was to expose the GaAs substrate using the $\text{CHF}_3/\text{CF}_4/\text{He}$ RIE used at the foundry. This is depicted in Figure 2.8, where it is pointed out that a residue remains on the GaAs surface following the etch. Figure 2.9(a) shows DGWs contaminated by this residue. SIMS (secondary-ion mass spectroscopy) analysis of the residue showed the presence of carbon and fluorine [124]. It has also been observed at Vitesse that residues of this form appearing on a dielectric surface are readily removed by O_2 plasma, but residues that come in contact with GaAs are not easily removed [125]. It is thus believed that they react chemically with the GaAs. In Shenoy's initial work, a combination of CF_4/O_2 RIE and BOE (buffered hydrofluoric acid) wet etching was successful in removing this residue prior to growth. However, this approach did not work after the same chips had "matured" for a time. With prolonged exposure to an energetic O_2 plasma, the residue was removed, but as shown in Figure 2.9(d), the underlying GaAs was highly textured and not suitable for use as a starting surface for epitaxial growth.

Therefore, while RIE must be used to define the majority of the DGW, Vitesse's dry etch must not be allowed to come in contact with the substrate¹. Rather, after RIE etching to within a few tenths of a micron of the substrate, the remaining dielectric should be removed with a gentle, selective BOE wet etch. This approach, shown in Figure 2.10, was taken on the test chip entitled MIT-OEIC-3. Unfortunately, precise control of the RIE has not been pursued at the Vitesse foundry since it is not required in the VLSI process. As a result, several microns of dielectric actually remained at the base of the DGWs, and the thickness of

1. This does not rule out the use of other dry etch processes which do not chemically or physically degrade the GaAs surface.




Figure 2.8: Initial DGW preparation technique

In this approach, used by Shenoy, *et. al.*, the foundry RIE was used to expose the GaAs substrate within the DGW regions.

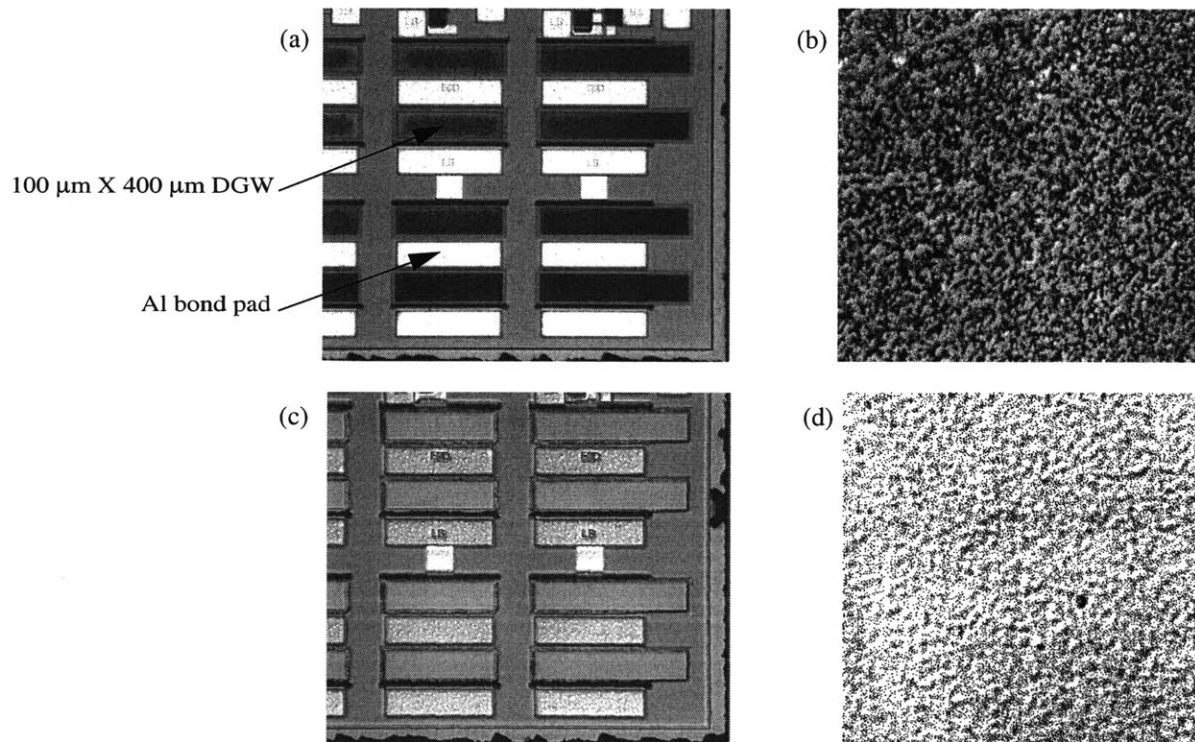


Figure 2.9: DGWs contaminated by RIE residues

Nomarski micrographs at 50X (a) and at 1000X (b) of DGWs containing RIE-generated fluorinated hydrocarbon residues. Corresponding images at 50X (c) and at 1000 X (d) after cleaning with O_2 plasma and BOE. All residue is removed, but GaAs substrate has a rough morphology




Figure 2.10: DGW preparation technique based on partial RIE etch of dielectric stack

To prevent RIE damage of the GaAs substrate, a revised DGW preparation technique attempted to use the foundry RIE to remove the dielectric stack to within $\sim 0.5 \mu\text{m}$ of the substrate. BOE would then be used to selectively etch the remaining dielectric. Unfortunately, limited control of the RIE etch rate and dielectric thickness resulted in a much thicker portion of the dielectric stack retained above the substrate. Completion of the DGW preparation process at MIT was difficult.

this layer was highly non-uniform over each die. The processing required to complete the DGWs was torturous, though not insurmountable for a properly motivated individual [8].

The task was further complicated by contamination in the available dry etching equipment which led to the unique growth defects shown in Figure 2.11. Defects of this type appeared in high densities on GaAs substrates which were directly exposed to contaminated plasmas, including O_2 plasma used for routine post-lithography ashing. On ICs where only the dielectric above the GaAs substrate was exposed to contaminated plasma (i.e. prior to actually exposing the GaAs), such defects may appear at a reduced density. Consequently, the use of both RIE and O_2 plasmas was discontinued in the DGW preparation process with the exception of the uncontaminated processes carried out at Vitesse. To reiterate, the circumstances leading to the defects in Figure 2.11(b) were linked to the specific pieces of equipment available for use at the time and do not, in general, constrain the use of dry processing in EoE. The decision to eliminate dry processes in DGW preparation is significant in the present context only because it impacts how the process




Figure 2.11: SEM images of “Spike” defects resulting from contaminated dry processing

Both the O₂ plasma asher and the RIE used in much of the EoE development work were found to be contaminated. (a) shows a GaAs substrate after exposure to an O₂ plasma ash. MBE growth on OEICs and bulk substrates treated in this way resulted in the unique defects shown in (b). The contamination coincided with the presence of sulfur on the GaAs surface suggesting that SF₆, which had been used in the same equipment, was as the source of the contamination.

was actually carried out, and thus the results achieved. In particular, it was found to have significant consequences in the LED demonstration shown in Chapter 3.

A successful DGW preparation procedure

To overcome the lack of RIE etch depth control, the solution outlined in Figure 2.12 has been employed starting with the test chip MIT-OEIC4. Each DGW is specified as a region of source/drain implant, which will form the n-contact to the optoelectronic device, beneath a region of Metal-1. At the foundry, the dielectric stack is etched down to Metal-1 using a CHF₃-based RIE. The aluminum-based Metal-1 acts as an etch stop in this process, eliminating the need for precise control. Figure 2.13 is an SEM of a DGW at this stage. Upon receipt of the die from the manufacturer, HCl and H₂O₂ are alternately applied to etch the AlCu_x and WN_x of the Metal-1 layer. With this layer removed, only a thin (~0.5 μm) dielectric remains over the substrate. While protecting the overglass and DGW sidewalls with a photoresist mask, the dielectric in the DGW regions is BOE wet-etched to expose the underlying GaAs. Because of the low etch rate of the densified Si₃N₄ film at the bottom of the dielectric stack, very long etch times may be

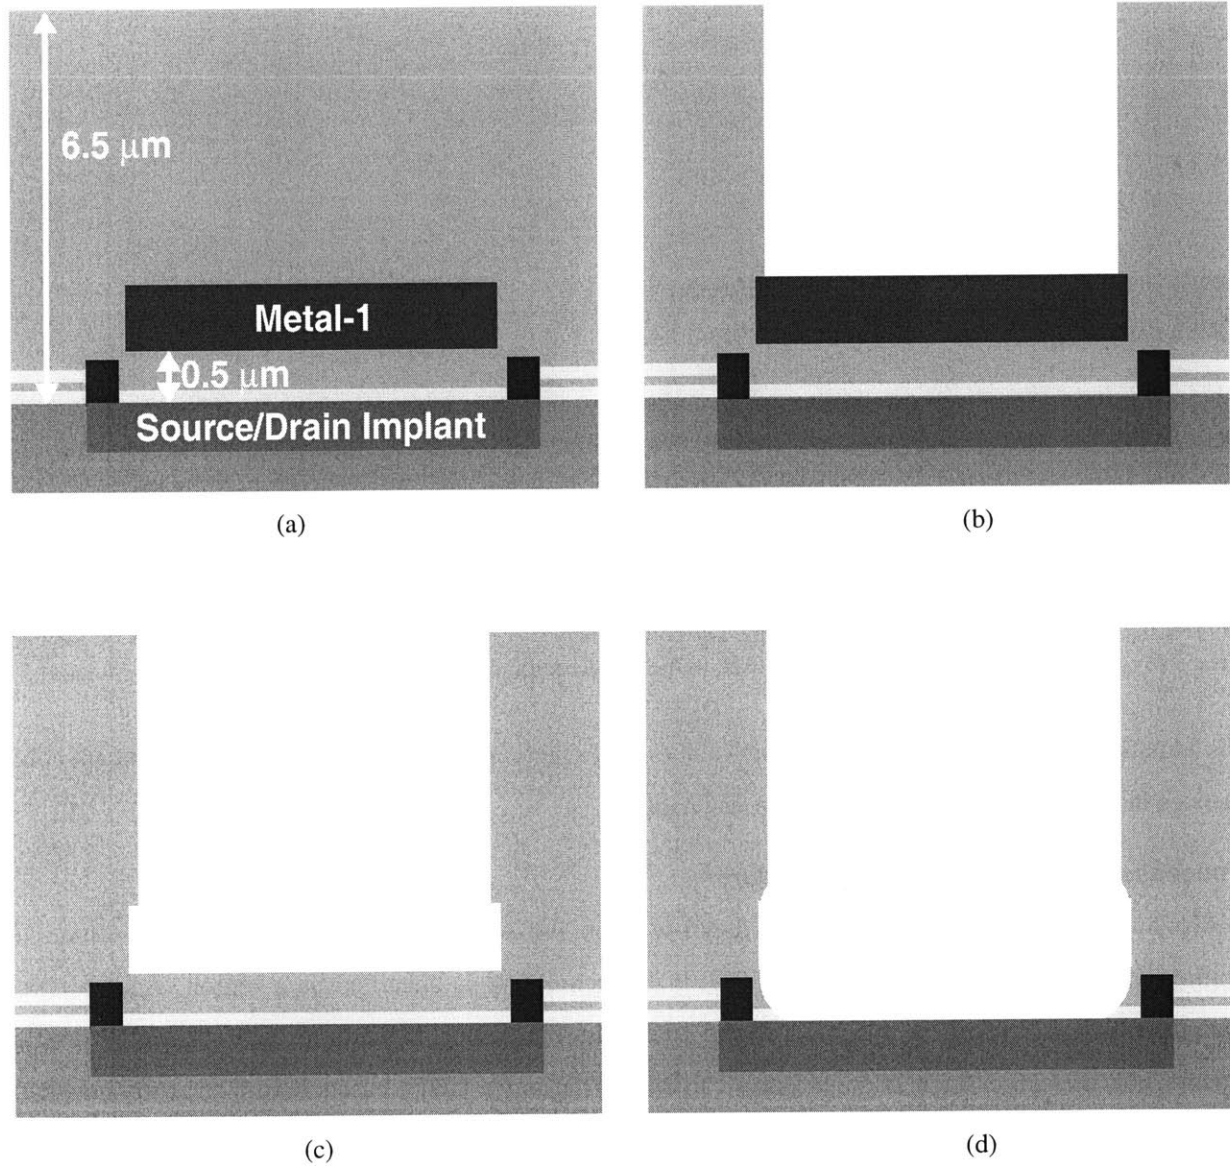


Figure 2.12: Successful DGW preparation process

(a) The initial structure consists of a source/drain implant, ohmic contacts, and metal-1 etch-stop. (b) The dielectric stack is etched using the foundry RIE. The metal-1 layer acts as an etch stop. (c) The metal-1 material is removed by selectively wet etching. The remaining dielectric is then selectively wet etched resulting in the final structure shown in (d).

required to completely expose the GaAs. Photoresist films applied over the 6.5 μm deep DGW holes typically held up to BOE for around ten minutes (SC-1827 positive photoresist [126] and NR8-3000 negative photoresist [127] were both used). In many cases, samples were patterned and BOE etched four times, for ten minutes each, in order to remove the dielectric remaining below the metal-1 etch stop. In other cases, however, the Si_3N_4 was more readily removed. Improved process control at the IC foundry should eliminate this variability, preferably in favor of the more readily removed Si_3N_4 film. Nonetheless, the procedure described here was repeatably used to produce damage-free GaAs surfaces within DGWs. A step-by-step process description is given in Appendix B. An SEM view of an almost-fully-cleaned DGW is shown in Figure 2.14 (a fully-cleaned DGW was not imaged in order to avoid pre-growth contamination).

One of the difficulties in applying the above procedure is in patterning photoresist within the 6.5 μm deep DGWs. Photoresist tends to be difficult to develop out of the corners of the rectangular DGWs (circular DGWs may be a possible solution, but have not been investigated). Largely to blame for the difficulty is the use of contact lithography which does not sharply expose the resist at the base of the DGWs. By comparison, the DGW sidewalls formed at Vitesse are well-defined and vertical.¹ Incomplete removal of photoresist from DGWs has led to the degradation of LEDs produced in the OTOCHIP project (Chapter 3), but proper DGW formation has been demonstrated in other instances.

An additional issue that had to be addressed within the DGW preparation step was the appearance of cracks in the overglass of ICs following MBE growth. The cracks did not seem to damage the electronics, but were a natural cause for concern. To prevent their formation, a 0.6 μm layer of SiO_2 was sputtered over the ICs prior to carrying out the above procedure. The added strength of this layer was sufficient to prevent the formation of cracks during the MBE temperature cycle.

2.4 Low Temperature Native GaAs Oxide Removal

Following DGW preparation, the OEIC die and an epi-ready GaAs sample, which serves as a RHEED (reflection high-energy electron diffraction) and pyrometer target, are indium soldered to a silicon wafer and loaded into the MBE system. Prior to growth, the native GaAs oxide must be removed in order to expose the underlying crystalline surface. As noted above, conventional thermal desorption of the GaAs

1. To achieve this result, the stepper focal depth must be specifically adjusted for this process step.






Figure 2.13: SEM of DGW on an OEIC as received from Vitesse/MOSIS
The 50 μm square recess was formed by RIE at the Vitesse foundry.



(a)



(b)

Figure 2.14: SEM of nearly completed DGW

Only a thin dielectric layer covers the GaAs substrate within the 50 μm square DGW in (a). The effect of the undercutting of the photoresist by the BOE etch may be seen in the close-up view given in (b).

oxide at $\sim 580^\circ\text{C}$ damages the circuit interconnects and interferes with wire-bonding. Motivated by recent reports of the use of atomic hydrogen in surface cleaning, the author's S.M. thesis included an investigation of the use of atomic hydrogen for the pre-growth removal of the GaAs native oxide [See Chapter 4 of Ref. 10 and citations therein]. To take advantage of this technique, the gas-source MBE (GSMBE) reactor was modified to include a high temperature hydrogen cracking cell [128]. As outlined in Figure 2.15, atomic hydrogen is produced within the MBE system by passing high purity hydrogen gas through the cracker with a filament temperature of $\sim 2100^\circ\text{C}$. The chamber pressure is maintained at $\sim 2 \times 10^{-5}$ Torr, which is comparable to the background pressure during GSMBE growth. The native oxide is removed at the growth temperature of 470°C by exposing the substrate to atomic hydrogen for 10 to 15 minutes [11]. Removal of the oxide layer is confirmed by the observation of a non-diffuse RHEED pattern as shown in Figure 2.15. Oxide removal has also been carried out at a substrate temperature of 300°C , and lower temperature are likely to be possible. Carrying out the process at the 470°C growth temperature has been found to be most expedient.

2.5 Reduced Temperature MBE Growth

Following removal of the GaAs native oxide, conventional MBE growth is used to produce the desired optoelectronic device heterostructure. Full EoE thermal compatibility is attained, while producing material of high optical quality, by a judicious material choice. The previous EoE work used traditional AlGaAs-based emitters. The AlGaAs material system owes its popularity to the fact that nearly exact lattice matching is maintained, relative to a GaAs substrate, over the entire Al composition range. This is shown by the vertical dashed line in the plot of bandgap-vs.-lattice constant in Figure 2.16. Thus, various bandgap energies may be produced without the complications related to lattice-mismatched growth. However, owing to the high oxygen affinity of aluminum, AlGaAs is optimally grown near 700°C . Below 700°C , oxygen is incorporated into AlGaAs with unity sticking coefficient and is found to accumulate at AlGaAs/GaAs interfaces [129,130]. Once incorporated in AlGaAs, oxygen acts as an efficient nonradiative recombination center [131,132]. If located in the active layers of a light emitter or at the interfaces to the active layer, non-radiative recombination centers dramatically reduce the efficiency of the device. (In heavily doped, majority-carrier-dominated layers, nonradiative recombination does not play a significant role.) In early EoE work, Shenoy, *et. al.*, found the broad-area pulsed room-temperature threshold current of InGaAs/GaAs/

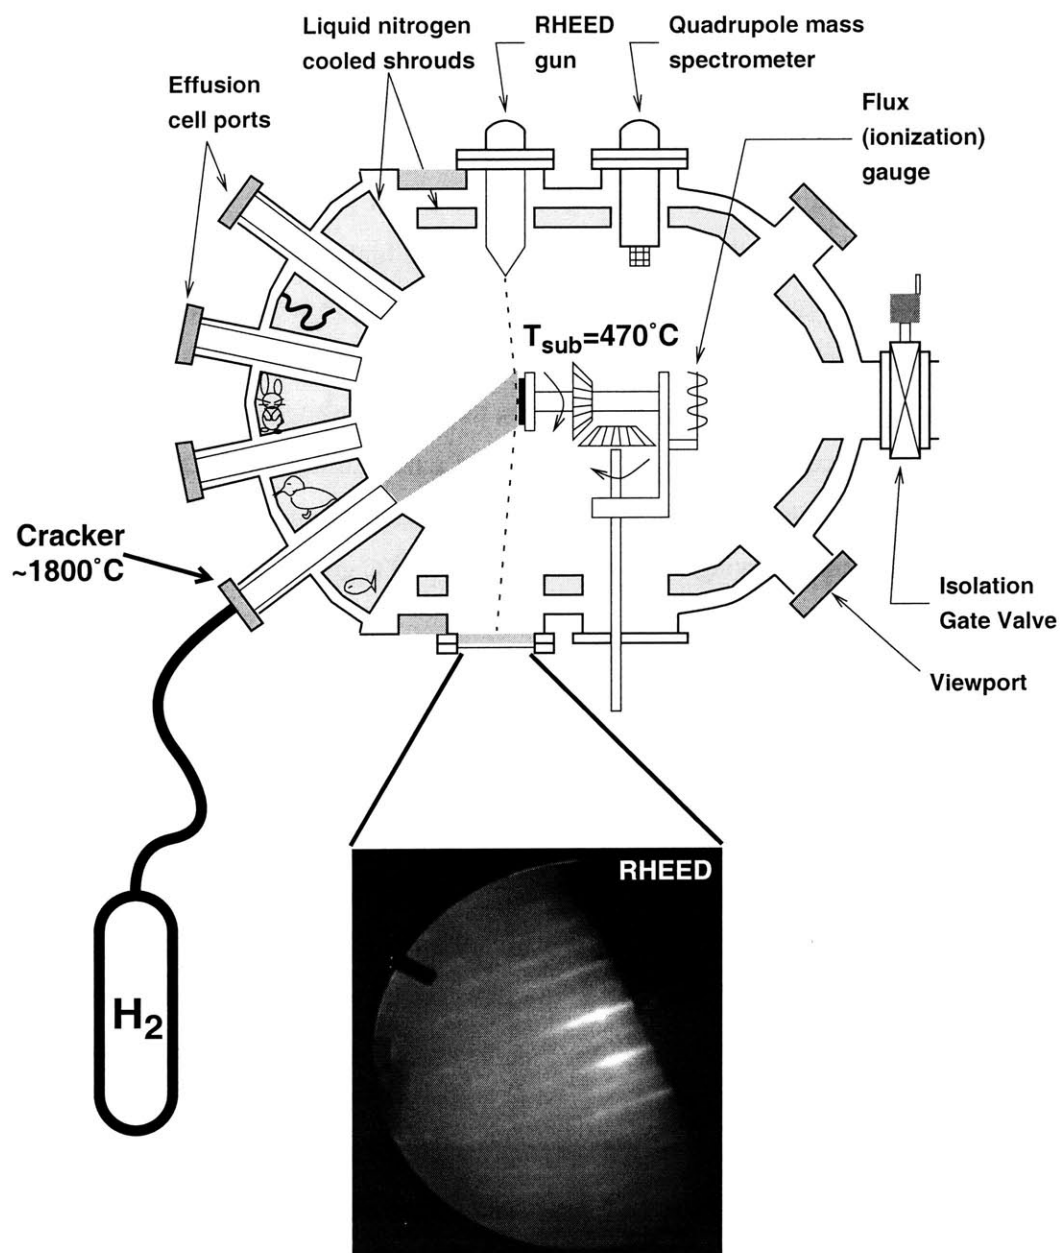


Figure 2.15: A schematic view of the use of *in situ* atomic hydrogen in MBE

High purity hydrogen gas (H_2) is cracked by a high temperature tungsten filament to produce a beam of atomic hydrogen (H) which is incident on the substrate. Removal of the GaAs oxide may be monitored using RHEED (reflection high-energy electron diffraction).

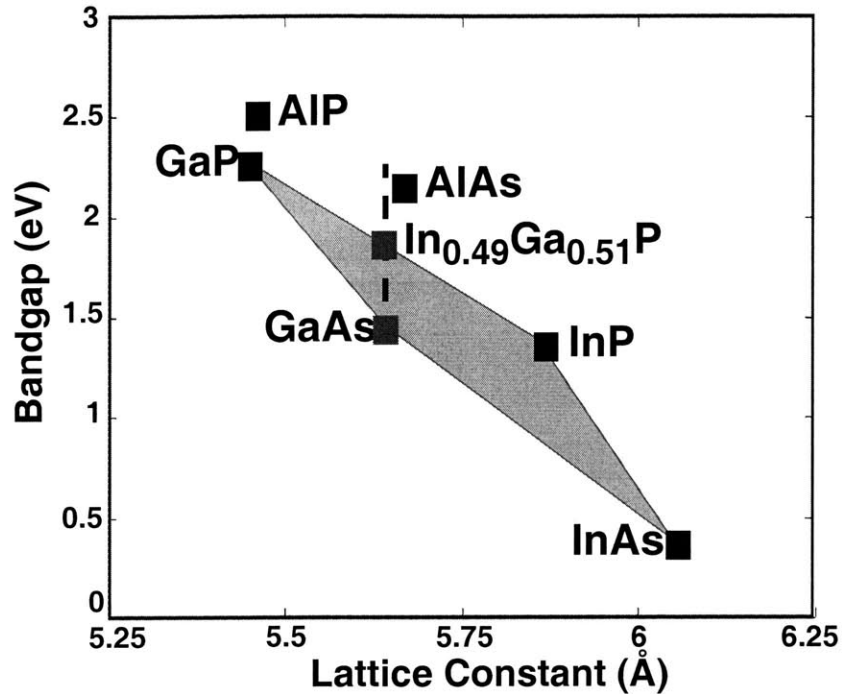


Figure 2.16: Plot of bandgap vs. lattice constant of some common III-V materials

This plot represents a simplified design space for the implementation of various optoelectronic devices. An important advantage of the ternary $\text{Al}_x\text{Ga}_{1-x}\text{As}$ material system, indicated by the dashed line, is the “automatic” lattice matching to GaAs which exists over its full composition range. However, the high oxygen affinity of Al hampers the growth of this material system at reduced temperatures. On the other hand, the Al-free quaternary $\text{In}_x\text{Ga}_{1-x}\text{As}_y\text{P}_{1-y}$ material system (shaded region) may be used to produce high-quality material at EoE-compatible temperatures using conventional growth techniques. In particular, a ternary member of this material system, $\text{In}_{0.49}\text{Ga}_{0.51}\text{P}$, is lattice matched to GaAs and may be used as a replacement for the wide bandgap AlGaAs layers in many optoelectronic devices.

AlGaAs quantum well laser diodes to increase from 740 A/cm^2 to 1800 A/cm^2 when the growth temperature was lowered from 600°C to 530°C [1].

Aluminum-free materials such as InGaAsP are normally grown at much lower temperatures than AlGaAs, and the recombination velocity at an $\text{In}_{0.5}\text{Ga}_{0.5}\text{P}/\text{GaAs}$ interface, 1.5 cm/s , is dramatically lower than the 210 cm/s figure measured for an $\text{Al}_{0.4}\text{Ga}_{0.6}\text{As}/\text{GaAs}$ interface grown at an *optimal* temperature [133]. $\text{In}_{0.49}\text{Ga}_{0.51}\text{P}$ (hereafter referred to as InGaP) is lattice-matched to GaAs and has been used as a replacement for the wide-gap AlGaAs cladding layers of in-plane laser diodes used in high power applica-

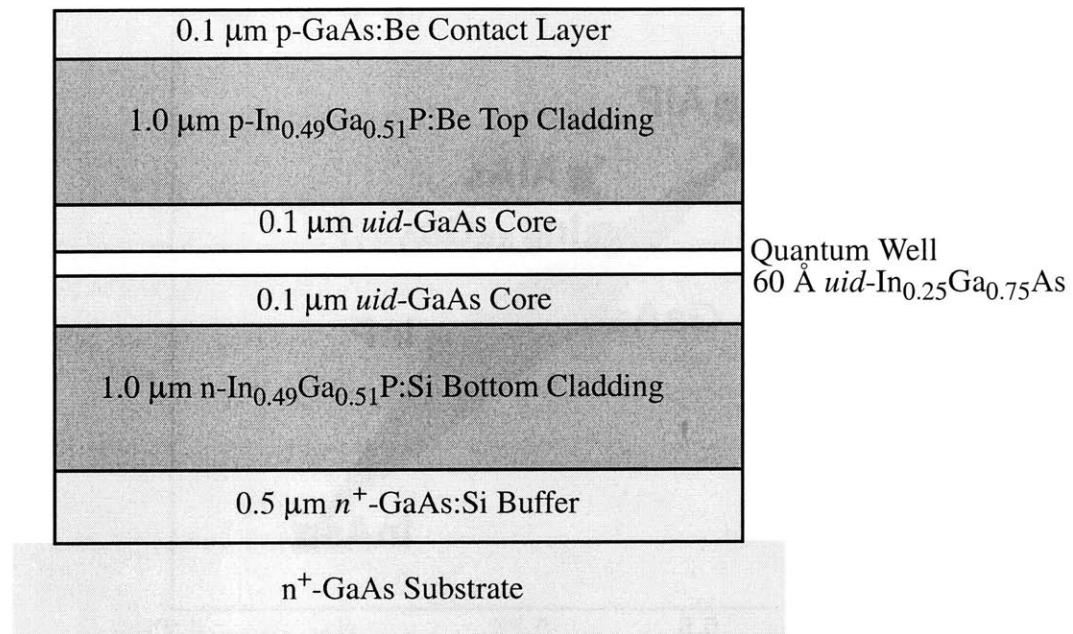


Figure 2.17: InGaAs/GaAs/InGaP QW-SCH laser structure

Strained InGaAs/GaAs/InGaP single quantum well separate confinement heterostructure (QW-SCH) used to demonstrate high-quality EoE compatible material growth.

tions where the use of Al-containing materials can increase the occurrence of device failure due to catastrophic optical damage [134]. These InGaP-based lasers currently find wide commercial usage [135].

In order to demonstrate the potential for producing high performance optical emitters at EoE-compatible temperatures, an in-plane laser diode similar in structure to that used in [134] (Figure 2.17) was grown entirely at 470°C using GSMBE (further details in [10]). The pulsed, room-temperature light-versus-current characteristics of a number of broad-area devices fabricated from this material are shown in Figure 2.18. The low threshold current of 200 A/cm² confirms that high quality optoelectronic device material may be produced at EoE-compatible temperatures using this material system. The corresponding value for the optimized device reported in [134] is 72 A/cm². The difference may be attributable to an error in the InGaP composition (53.3% In as opposed to the targeted 49%; sample number R170), the inclusion of Ga spitting defects which were produced by the GSMBE at the time, and the specifics of the broad-area device processing [10].




Figure 2.18: Pulsed room temperature L-I characteristics of broad-area aluminum-free lasers

Pulsed room temperature light output power vs. input current density for strained InGaAs/GaAs/InGaP QW-SCH broad area lasers. Ohmic contact stripes were made with non-annealed Cr/Au or Ti/Au metallization. Cavities of various lengths were tested. A low threshold current density of 200 A/cm^2 was attained with a $792 \mu\text{m}$ long Cr/Au contacted device. A $284 \mu\text{m}$ Ti/Au contacted device output over 80 mW without failure.

In addition to the above laser structure, GaAs/InGaP LED heterostructures which are the basis of the devices used in Chapter 3, were grown at 470°C and compared with a similar GaAs/AlGaAs LED heterostructure grown at 530°C [10]. Light-versus-current characteristics of these structures are shown in Figure 2.19. The GaAs/InGaP electroluminescent efficiency of 3 mW/A (2% quantum efficiency) measured at 50 A/cm^2 was 10 times greater than that of the GaAs/AlGaAs sample. In fact, it is near the theoretical quantum efficiency limit set by total internal reflection of light at the GaAs-to-air interface. This result further confirms that high quality optoelectronic devices which are fully compatible with the 475°C EoE thermal limit may be produced by employing the aluminum-free InGaAsP/GaAs material system.




Figure 2.19: L-I characteristics of broad-area LED test structures


The test structure is a 500 μm diameter metal contact. The light output power vs. current curves exhibit a “threshold” corresponding to the transition from non-radiative- to radiative-dominated recombination mechanisms. The InGaP-based samples, grown at 470°C, vary in material quality (primarily corresponding to variations in composition and thus lattice mismatch) and may be compared to the AlGaAs-based sample grown at 530°C. The inset shows the layer structure of samples R196, R210, R245, and R248. R176 and R177 instead begin the structure with a 0.5 μm GaAs buffer layer. Sample 6042 uses a 1.2 μm p⁻-GaAs core and 0.3 μm (top) and 0.2 μm (bottom) Al_{0.3}Ga_{0.7}As claddings. Please see [10] for further details.

2.6 Polycrystalline Deposit Removal


While material is epitaxially grown in the DGW regions, polycrystalline material is deposited over the remaining die surface. An SEM micrograph of an OEIC following growth is shown in Figure 2.20. The procedure used by Shenoy at this point to remove the polycrystalline deposits consisted of masking the epitaxial material with photoresist and applying a standard GaAs wet etchant such as $\text{H}_3\text{PO}_4:\text{H}_2\text{O}_2:\text{H}_2\text{O}$ (1:1:5). The material change from AlGaAs to InGaAsP necessitated the use of a different etching chemistry. The use of a wet process also raised the concern that overetching could allow the etchant to penetrate the photoresist mask and attack the GaAs at the base of the DGW, thereby cutting off electrical contact to the device which is made through the shallow (~ 100 nm) source/drain implant. Thus, this process step was altered to use a dry etch in order to enhance its overall robustness. The GaAs/InGaAsP polycrystalline deposits have been removed using Cl_2 ion-beam-assisted etching with an ECR source (ECR-IBAE) [136]. Other dry-etch processes compatible with the GaAs/InGaAsP chemistry may also be used.

A detailed recipe for polycrystalline deposit removal is given in Appendix B and is described briefly here. Following growth, roughly $0.4 \mu\text{m}$ of SiO_2 is sputtered over the die. The DGW regions, plus a $5 \mu\text{m}$ border, are covered by photoresist and the SiO_2 is etched using BOE. With SiO_2 masking the epitaxial material, Cl_2 ECR-IBAE at 195°C is next used to etch the deposits. The manufacturer's overglass, which has been retained over the aluminum bondpads, provides an etch stop. An SEM of the result of this step is shown in Figure 2.21. The overglass is now covered with photoresist and the SiO_2 over the DGW regions and their perimeter is removed with BOE. The polycrystalline material exhibits poor adhesion to the overglass, and as a result, the ring of polycrystalline material left under the perimeter SiO_2 is also removed in this step.


An alternative to the use of the latter BOE step in removing the polycrystalline rings is to simply scrub the chip surface with a foam swab while the sample is immersed in propanol. The thin rings are sheared off and the surface of the epitaxial material remains covered with a protective SiO_2 film. No damage to either the IC surface or to the epitaxial material has been observed following this procedure. This technique can be applied to full wafers using commercial equipment used to clean wafer surfaces following chemical-mechanical polishing.



(a)



(b)



(c)

Figure 2.20: SEM of OEIC surface and cross section showing post-growth polycrystalline deposits

(a) Shows a region of an OEIC surface covered with polycrystalline deposits which includes DGWs and electronics. (b) Cross section, made by FIBE [137], of a DGW which is filled with epitaxial material and surrounded by polycrystalline deposits, and (c) is a close-up of the left side of (b).




Figure 2.21: SEM of OEIC following dry etch of polycrystalline deposits

(a) DGW following the dry etch of the polycrystalline deposits. (b) Close-up of the lower right corner of (a). A ridge of polycrystalline material remains around the DGW due to the 5 μm outset SiO_2 mask.

2.7 Optoelectronic Device Fabrication

The EoE process concludes with optoelectronic device fabrication. By adjusting the total height of the optoelectronic device structure to match the thickness of the interconnect dielectric stack, a planar die surface is attained following removal of the polycrystalline deposits. This level of planarity, which is limited by the height variation resulting from the on-chip interconnects, allows a conventional device process developed on a bulk substrate to be used for integrated device fabrication.

A specific device fabrication flow is presented in Chapter 3 for the case of integrated LEDs. One key aspect of the device fabrication procedure which is specific to EoE is the final metallization step needed to electrically connect to top-side contact of the device with the on-chip electronics. In a commercial environment, the EoE process could actually be inserted prior to completion of the back-end process, and the connection to the optoelectronic device would be made upon resumption of the standard back-end process. In the present context, however, the top-side connection to the optoelectronic device is equivalent to a fifth IC metallization layer and must address the same surface non-planarities that are encountered commercially. A similar metallization process is thus employed: aluminum is conformally sputtered and the interconnects are defined by etching. The feature sizes required in this step are large enough to allow a patterning of the Al using PAN etchant (77% H_3PO_4 , 20% acetic acid, 3% HNO_3). Details of the metallization process are given along with that of the LED in Chapter 3.

In using an Al interconnect to the optoelectronic device, care must be taken to ensure compatibility between the interconnect and the ohmic contact metallizations. In particular, a reaction between Al and Au may result in the condition known as “purple plague” [138]. This condition was encountered on the first demonstration of an InGaP-based LED integrated with non-degraded electronics. As shown in Figure 2.22, the LED top contact had to be made by direct probing while the bottom contact was made through on-chip interconnects. To eradicate the purple plague, a Ni barrier layer was deposited over the Ti/Au p-type ohmic contacts being used. The Al-free ohmic contact Ti/Pt was also investigated, but electron-beam evaporation of Pt using the available equipment raised the sample temperature to the point where the photoresist being used to pattern the metal (via a lift-off process) was damaged.

2.8 Completed OEICs

Chapter 3 will present results of integrated GaAs/InGaP LEDs. The first demonstration of these devices, detailed in Figure 2.22, was carried out on the MIT-OEIC-4 test chip fabricated in the Vitesse H-GaAs III process. A layout of this test chip is shown in Figure 2.23. Figure 2.24 shows a scanning electron micrograph of the top surface of an integrated LED in a 50 μm DGW. This device appeared on the OPTO-




Figure 2.22: First demonstration of integrated InGaP-based LEDs

The first demonstration of InGaP-based LEDs, grown at EoE-compatible temperatures, was made on the MIT-OEIC-4 test chip. In this figure, the emission from an integrated LED is observed with a microscope equipped with a CCD camera. The probe tip on the right connects to the n-type LED terminal through the bondpad and on-chip wiring which connect to the source/drain implant at the base of the DGW. The p-type contact to the LED, however, is made by directly probing the top of the device. This was necessary because of a metallurgical reaction, known as “purple plague,” between the Au-based ohmic contacts and the Al interconnect metallization.




Figure 2.23: Layout plot of the MIT-OEIC-4 test chip

Fabricated in the Vitesse H-GaAs III process, the 4.7 mm x 4.7 mm MIT-OEIC-4 test chip was used for the first demonstration of an integrated GaAs/InGaP LED.

CHIP described in Chapter 3. The Ti/Au/Ni p-contact to the $30\ \mu\text{m} \times 30\ \mu\text{m}$ current confinement mesa is interconnected by Al to nearby electronics. The DGW is surrounded by an optical shield, designed to block stray LED emission from coupling into nearby electronics, which is formed by stacking interconnect metal and via patterns up to the metal-3 layer. In Figure 2.25, a similar LED has been cross-sectioned using focused ion-beam etching [137]. Figure 2.26 gives a closer view of this cross section. The figures show




Figure 2.24: SEM on an integrated LED

SEM micrograph of an integrated GaAs/InGaP PiN LED. Epitaxial material was grown inside a 50 μm square DGW surrounded by an optical shield. A 30 μm mesa confines current injected by a Ti/Au/Ni p-type ohmic contact. Al interconnects the p-contact to on-chip electronics. The n-contact is made through a source/drain implant below the DGW.

the sidewall profile of the epitaxial material, the ohmic contact to the n^+ source/drain implant beneath the LED, and the top-side LED contact.

Examples taken from completed OEICs are shown in Figure 2.27 and Figure 2.28. The LEDs, photodetectors, and MESFET electronics are seen to fit together seamlessly.




Figure 2.25: FIB cross section a completed LED

Cut along the symmetry axis of the device, through the contacts [137]. A more detailed view is given in Figure 2.26.

Thermal Stability Revisited

To verify that the procedures developed to reduce the MBE growth and GaAs native oxide desorption temperatures have been successful in eliminating degradation of the electronic performance, 23-stage ring oscillators on five chips were characterized before and after EoE processing (different sets of chips are used in each case since bondpad openings are not yet present on chips used for integration). The mean and standard deviation of the extracted inverter delay is 67 ± 2 ps both before and after growth and fabrication. The consistency of the pre- and post-integration inverter delay is a strong indication that the electronics performance is not affected by the EoE process.




Figure 2.26: FIB cross section of a completed LED; detailed views

Close ups of left (a) and right (b) sides of DGW cross sectioned in Figure 2.25 [137]. (a) includes the optical shield structure while (b) shows the ohmic contact to the DGW and the metallization leading to it. The top-side contact structure is also seen in (b).

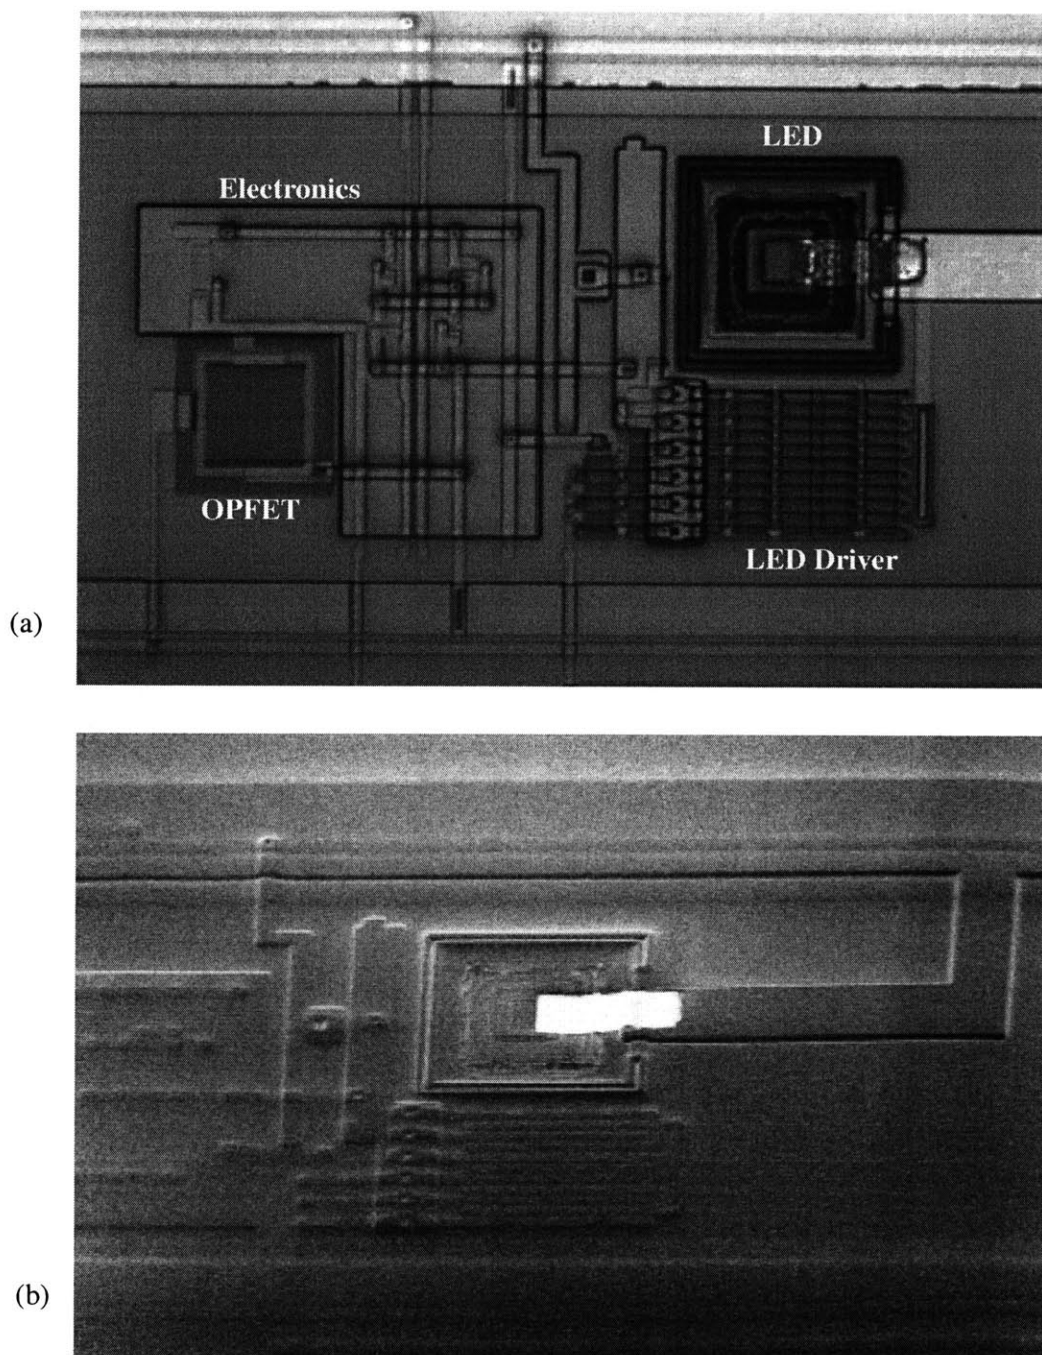


Figure 2.27: EoE integration of and LED, OPFET, and electronics

(a) Photograph and (b) SEM of an optoelectronic integrated circuit fabricated as part of the OPTOCHIP project. It includes an LED, an OPFET photodetector, an LED driver, and some processing electronics.



Figure 2.28: Integrated LEDs, MSM photodetectors, and VLSI electronics

Photograph of a portion of a completed OEIC. Six integrated LEDs in 50 μm DGWs are visible along with ten MSM photodetectors and VLSI MESFET electronics. This image was taken from an OPTOCHIP OEIC.

Summary

This section has reviewed the development and details of the Epitaxy-on-Electronics process. The starting point for EoE integration is fully-metallized die fabricated by Vitesse using their standard H-GaAs III or H-GaAs IV process. Die have been used thus far due to cost and equipment limitations. However, the EoE process should be scalable to full-wafer production in a manufacturing environment since it relies only on standard growth and fabrication techniques. At its core, Epitaxy-on-Electronics consists of: 1) commercial fabrication of GaAs VLSI electronics, 2) formation of DGWs to expose regions of the GaAs substrate, 3) conventional growth of optoelectronic device heterostructures using molecular-beam epitaxy, 4) removal of polycrystalline growth deposits leading to a planar surface, and 5) optoelectronic device fabrication using well-established processing techniques. Thus, EoE produces monolithic optoelectronic-VLSI circuits by merely augmenting conventional electronic and optoelectronic device processes with the two EoE-specific steps of DGW formation and polycrystalline deposit removal. By making use of commercial electronics, EoE allows the fabrication of complex OEICs without a massive up-front investment in VLSI process development. By growing epitaxial heterostructures within implanted DGWs, EoE overcomes the non-planarity limitation of traditional optoelectronic integration approaches. And finally, by using well-established growth and fabrication techniques, EoE lends itself to high-volume, high-yield, and high-reliability manufacturing.

3

Integrated LEDs and the OPTOCHIP Project

As shown in Chapter 2, efficient lasers and LEDs may be grown at EoE compatible temperatures by making use of $\text{In}_{0.49}\text{Ga}_{0.51}\text{P}$, hereafter referred to as InGaP, rather than $\text{Al}_x\text{Ga}_{1-x}\text{As}$ as the wide bandgap layer in these structures. While integration of lasers has been a long term goal of the EoE effort, LEDs were chosen as an initial demonstration vehicle in order to avoid the additional complications associated with the growth and fabrication of lasers. Section 3.1 will describe the process used to fabricate both bulk and integrated LEDs. LED design considerations and bulk LED results will then be highlighted in Section 3.2. Integrated LEDs were used as part of a research foundry project known as OPTOCHIP. This effort will be described in Section 3.3 and the characteristics of LEDs and OEICs completed as a part of it will be given in Section 3.4. Section 3.5 will then address a number of optical crosstalk issues which played a role in the OPTOCHIP OEICs. Finally, Section 3.6 will make a number of key observations based on the OPTOCHIP project and the LED and OEIC results.

3.1 LED fabrication process

Figure 3.1 shows the LED structure used to demonstrate the EoE integration process and employed in the OPTOCHIP project. Figure 3.2 (repeated from Figure 2.24) is a scanning electron micrograph of an integrated LED in a $50\ \mu\text{m}$ DGW. The Ti/Au/Ni p-contact to the $30\ \mu\text{m} \times 30\ \mu\text{m}$ current confinement mesa is interconnected by Al to nearby electronics. The DGW is surrounded by an optical shield, designed to block stray LED emission from coupling into nearby electronics, which is formed by stacking interconnect metal

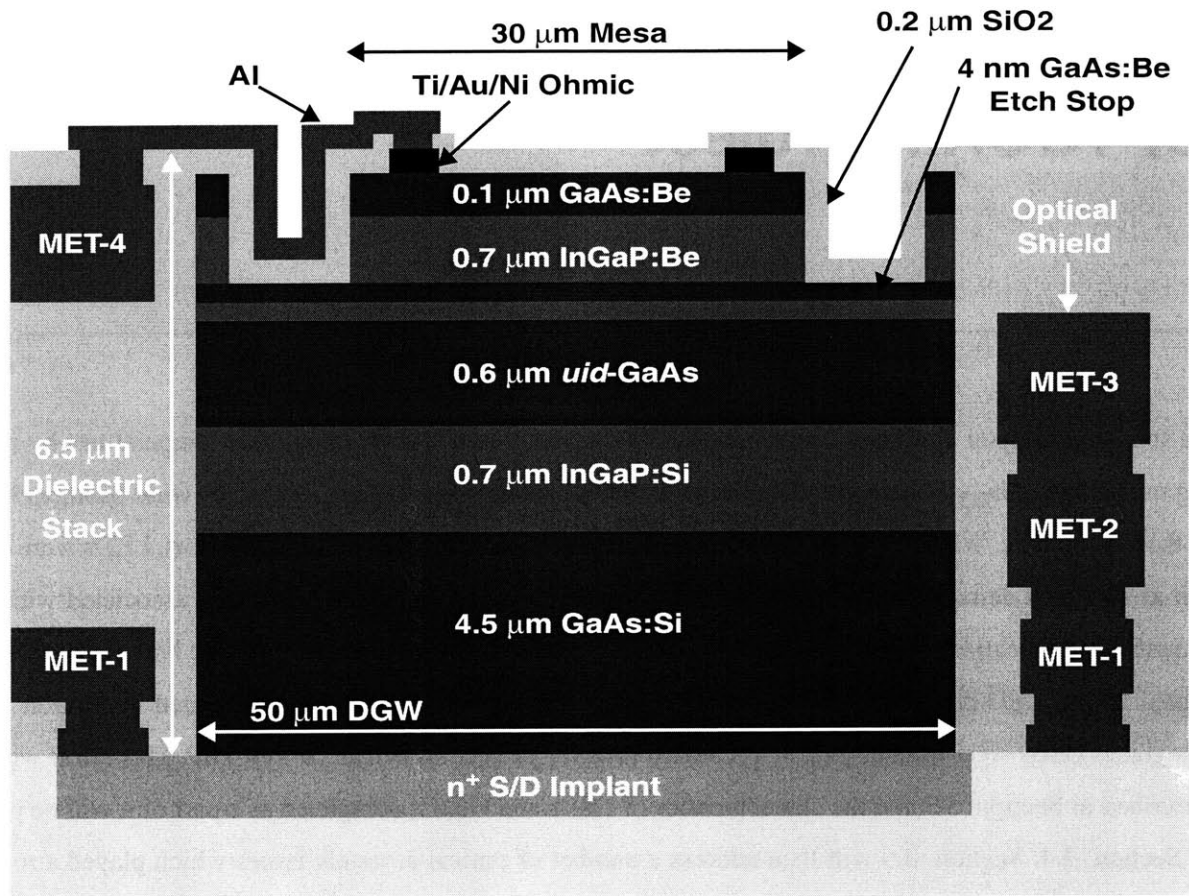


Figure 3.1: Cross-sectional diagram of integrated LED structure

The PiN LED heterostructure is grown within a 50 μm DGW on top of an n^+ source/drain implant which provides the n-type contact to the device. The active layer of the LED is 0.6 μm of unintentionally doped GaAs and is surrounded by wider-bandgap n- and p-type InGaP layers. A 4 nm GaAs etch stop layer is placed 0.1 μm above the GaAs core (within the p-type InGaP layer). A 0.1 μm layer of p-type GaAs is placed above the p-InGaP layer to aid in ohmic contact formation and, in conjunction with the etch stop, to control the mesa etch which is used for lateral current confinement. The p-type ohmic contact consists of an adhesion layer of Ti (30 nm) followed by a conduction layer of Au (0.15 μm) and a Ni barrier layer (50 nm). The Ni barrier prevents the Au from reacting with the Al metallization used to interconnect the LED to the on-chip electronics. Also shown is an optical shield formed by stacking on-chip metallization layers up to Metal-3. The shield surrounds the DGW and prevents stray emission from the LED from coupling into nearby electronics.




Figure 3.2: SEM on an integrated LED

SEM micrograph of an integrated GaAs/InGaP PiN LED. Epitaxial material was grown inside a 50 μm square DGW surrounded by an optical shield. A 30 μm mesa confines current injected by a Ti/Au/Ni p-type ohmic contact. Al interconnects the p-contact to on-chip electronics. The n-contact is made through a source/drain implant below the DGW.

and via patterns up to the metal-3 layer. The optical shield will be further discussed in Section 3.5.

The light-emitting core of the LED is 0.6 μm of unintentionally-doped GaAs between 0.7 μm of p-type InGaP on top and 0.7 μm of n-type InGaP on the bottom. A 4 nm GaAs etch stop layer is placed 0.1 μm above the core (inside the p-InGaP layer). A 0.1 μm p-GaAs contact layer on top of the device facilitates ohmic contact formation and mesa etching (see below). A 4.5 μm thick n-GaAs buffer layer at the bottom of the heterostructure is grown to achieve planarity with respect to the surface of the dielectric stack. A foundry-contacted source/drain implant in the DGW area provides the bottom contact to the LED. The doping of the p- and n-regions is nominally $2 \times 10^{18} \text{ cm}^{-3}$. This heterostructure was adapted from previ-

ous EoE work using AlGaAs-based LEDs [1] and has not been optimized for the Al-free materials. The LED heterostructure was grown by GSMBE at 470 °C with the exception of the majority of the thick GaAs buffer layer which was grown at 400°C in order to reduce the time spent at 470°C to less than the allowed five hours (see Section 2.1). No degradation of the conductivity of this layer or of the quality of the subsequent layers was observed as a result of this practice.

LEDs were fabricated following the GSMBE growth and the removal of the polycrystalline deposits. The fabrication process is nearly identical to that of similar devices on bulk wafers. The primary difference is in the use of thick photoresist (~3 μm) which is needed in order to insure step-coverage over chip-surface non-planarities. Contact lithography is effective at patterning the 2-3 μm minimum feature sizes, however, nonuniform resist thickness over the limited die area leads to nonuniform mask contact and, in turn, variation in device feature sizes and characteristics. As pointed out in Chapter 2, step coverage issues also lead to the use of aluminum interconnect metallization. The fabrication process is outlined below, and further details are given in Appendix B.

Ohmic contact metallization

The LED fabrication process begins with ohmic contact formation using lift-off. The electron-beam evaporated contacts consist of 30 nm of Ti followed by 0.15 μm of Au and 50 nm of Ni. The Ni layer is used to avoid the undesirable reaction (“purple plague”) of the gold-based ohmic contact with the Al interconnect metallization [138]. Following lift-off, the sample is rapid thermal annealed at 420°C for 20 seconds.

Mesa etch

A 5 μm wide “moat” is next etched for current confinement around a 30 μm square active mesa. The GaAs contact layer is first etched for 1 minute using 1 H₃PO₄ : 1 H₂O₂ : 25 H₂O (which does not etch InGaP). The p-InGaP layer is subsequently etched for 35 minutes down to the GaAs etch-stop using 1 H₃PO₄ : 1 HCl : 1 H₂O (which does not etch GaAs). This latter etch exhibits electrochemical properties which require the presence of both the etch-stop as well as the conductive GaAs contact layer in order to achieve controllable, well-defined etching. Furthermore, the etch is sensitive to material quality. This tendency has not been systematically studied but appears to track the compositional accuracy of the InGaP. Tightly lattice matched InGaP films lead to good etch characteristics such as shown in Figure 3.3. Lower




Figure 3.3: SEM of InGaP wet etching result

A bulk sample of LED material (heterostructure shown in Figure 3.1) was etched using 1 H_3PO_4 : 1 H_2O_2 : 25 H_2O to etch the 0.1 μm GaAs contact layer and 1 H_3PO_4 : 1 HCl : 1 H_2O to etch the 0.6 μm InGaP layer down to the 4 nm GaAs etch stop. In this high-quality sample, the surface at the base of the mesa is smooth and the mesa sidewalls are well-defined and vertical.

InGaP quality leads to a rough etched surface (see, for instance, Figure 3.5). Both the inherent InGaP quality, and the degraded etch characteristics may impact device performance.

Passivation and via etch

Next, a 0.2 μm layer of SiO_2 is sputtered over the device/OEIC surface. A via is etched through this dielectric over a designated region of the ohmic contact using buffered-HF. Openings are also made, using CHF_3/O_2 reactive ion etching, to allow access to contact pads (for electrical connection to the LED top contact) and to bondpads.

Interconnect metallization

The LED fabrication process concludes with the interconnect metallization. A one minute *in-situ*, argon-sputtering back-etch is used to remove the aluminum surface oxide (on the OEIC contacts and bond-




Figure 3.4: SEM of LEDs on first OPTOCHIP OEIC showing poor interconnect metallization

Nonplanarity of the OEIC surface leads to nonuniform photoresist coverage and exposure. On the initial fabrication run of OPTOCHIP OEICs, this resulted in a large number of open circuits following the wet etching of the Al interconnect metallization. On subsequent runs, the interconnect geometry was modified to eliminate this problem.

pads) prior to sputtering $0.75\ \mu\text{m}$ of aluminum. This layer is patterned using PAN etchant ($77\% \text{H}_3\text{PO}_4$, 20% acetic acid, $3\% \text{HNO}_3$) to interconnect the LED top contacts to the OEIC electronics. As already pointed out, sputtering is used in this step in order to insure film coverage over the chip-surface non-planarities. Surface height variations are a result of the underlying metallization layers. There is also an abrupt $1\text{-}1.5\ \mu\text{m}$ step at the boundary of the overglass cut exposing the metal-4 contact pad used for the LED p-type connection. Al sputtering is effective in covering these features.

A further issue, however, arises in patterning the Al film using a wet etch. Wet etching was selected here as a practical matter in order to minimize dependence on the dry-etching equipment that was available. Due to variations in resist coverage and exposure resulting from surface non-planarities, the size and geometry of the interconnect patterns must be selected in order to avoid inadvertent open circuits. Figure 3.4 shows LEDs on the first (preliminary) OPTOCHIP OEIC to be fabricated. The Al metallization pattern used in this case resulted in a high occurrence of open circuits due to poor etching. In response, the metal-




Figure 3.5: SEM of an integrated LED in an 85 μm DGW

This LED was fabricated inside an 85 μm DGW. Superior results were achieved, as compared with the 50 μm DGWs, because of the greater ease of DGW preparation and the greater local planarity during LED fabrication. The epitaxial material is seen to have good surface morphology, and the mesa and metallization patterns are well defined. Unlike the etch test shown in Figure 3.3, however, the bottom of the “moat” which isolates the mesa is seen to be rough. This is a result of material quality limitations which were common to both the integrated LEDs and the bulk devices which were grown simultaneously.

lization pattern was changed to the more robust one seen in Figure 3.2. This came at the expense of reduced LED efficiency since the new pattern covered a larger area of the LED active mesa. Better planarization in the IC process--for instance, through the use of chemical-mechanical polishing--would certainly improve this situation, as would carrying out the EoE process prior to completion of the IC metallization.

A closer look at a completed, integrated LED is shown in Figure 3.5. This device is built in an 85 μm DGW.

3.2 Bulk LEDs

To be useful in optical interconnect demonstrations, the integrated LEDs should combine high output power and power conversion efficiency (“wallplug” efficiency) with a low operating voltage. While the demonstrative nature of the LED development did not justify the iterative process needed to optimize the the material heterostructure, a range of design options in the geometrical layout of the device were examined. Relevant design considerations include current density, current spreading, contact resistance, and extraction efficiency.

In the double-heterostructure PiN LEDs being used, electrons and holes are injected from the n- and p-type InGaP layers, respectively, into the lower-bandgap GaAs active region. The basic operation of this device is reviewed in Figure 3.6. Within the active region, these carriers may recombine by either a radiative or a nonradiative mechanism. For a carrier concentration n , non-radiative recombination proceeds as An while radiative (“bimolecular”) recombination goes as Bn^2 . For low carrier concentrations, non-radiative recombination dominates. Above a certain concentration, determined by the relative sizes of A and B , radiative recombination becomes dominant. In the device L-I (light-output vs. current) characteristics, this behavior is manifest as a threshold below which light output is negligible and above which light output is proportional to current (above “threshold” and in steady state, Bn^2 becomes proportional to current). The crossover can also be seen as a “knee” in the I-V (current vs. voltage) characteristics since each recombination mechanism give rise to a different ideality factor. Alternately, the LED characteristics may be viewed as arising from a parallel combination of a pair of diodes corresponding to the two recombination mechanisms.

It is desirable to maximize the current density of the LED in order to operate well above the non-radiative to radiative cross-over. This suggests minimizing the mesa area. However, power efficiency also depends on ohmic contact resistance. To minimize the resistance of the ohmic contact, the metallization area should be maximized. In the simple mesa confined devices being used, the ohmic metallization blocks the light emission and reduces external efficiency. These issues are reviewed in Figure 3.7. As in previous demonstrations, the LED mesa was chosen to be $30\ \mu\text{m} \times 30\ \mu\text{m}$. This mesa size is consistent with the typical spot size of the optical systems in which the LED would be used. Furthermore, as a practical matter, feature sizes of the ohmic metallization had to be kept to $\sim 2\text{-}3\ \mu\text{m}$ and additional space had to be allowed on

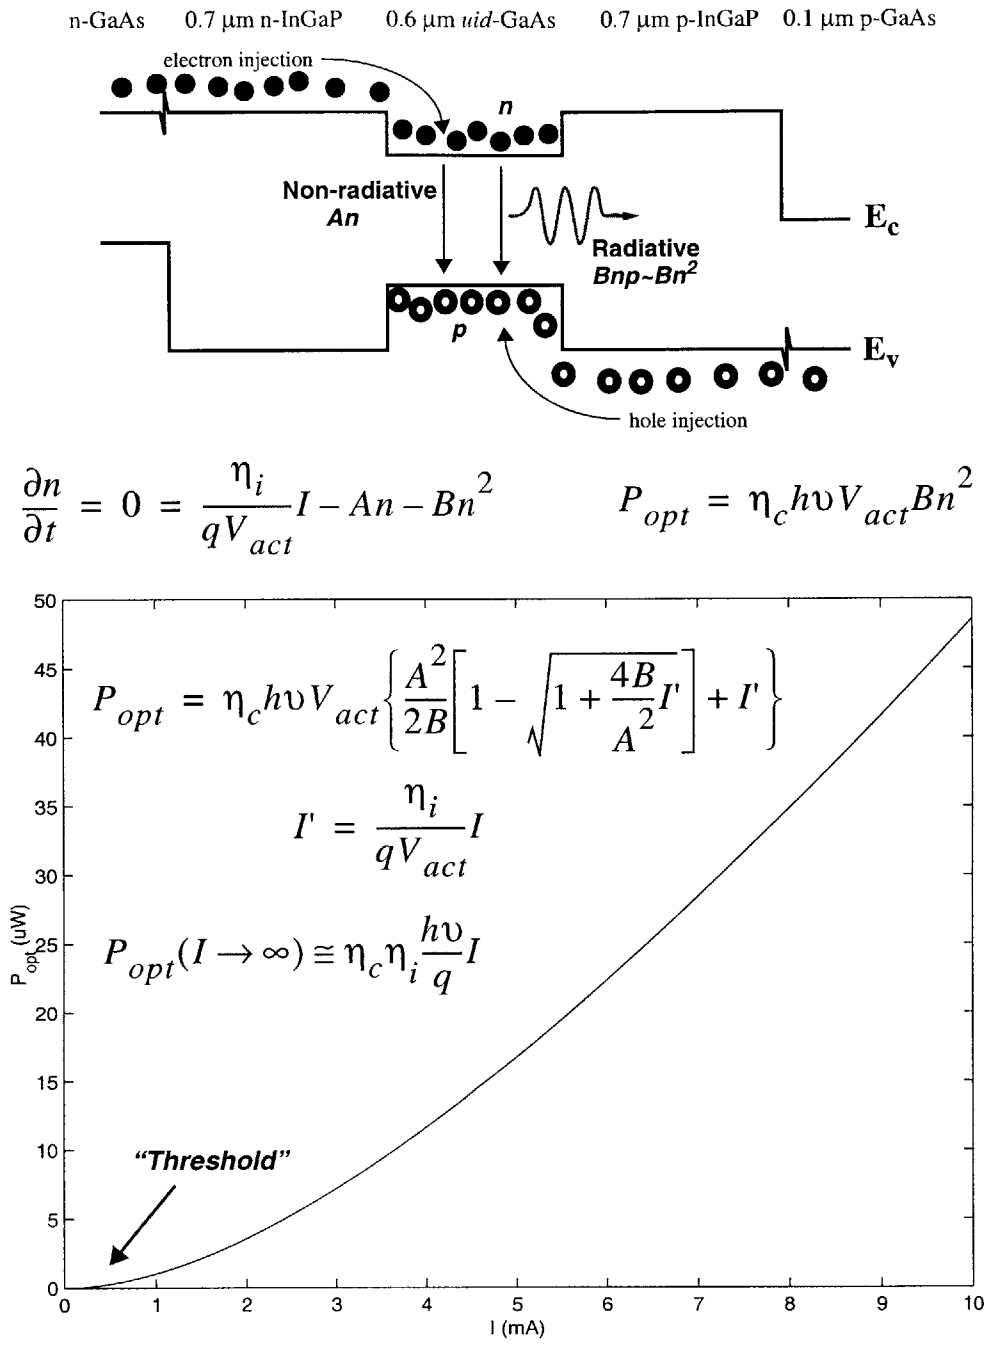


Figure 3.6: Basic LED operation

A rate-equation formalism is used to balance the terminal current with radiative and non-radiative recombination. The theoretical steady-state output power is plotted for a 30 μm square LED with A=2x10⁷ s⁻¹ and B=2x10⁻¹¹ cm³s⁻¹, as determined in [14]. The internal efficiency and collection efficiency were set to η_i=0.9 and η_c=0.01, respectively.

$$\begin{aligned} \text{External Quantum Efficiency:} \quad \eta_{ext} &= \frac{L}{I} \propto \frac{A_{act} - A_c}{A_c} \\ \text{Wallplug Efficiency:} \quad \eta_{wall} &= \frac{L}{IV} = \frac{\eta_{ext} I}{I(V_J + IR_c)} \\ \text{Contact Resistance:} \quad R_c &= \frac{r_c}{A_c} \end{aligned}$$

$$A_c \uparrow \quad \eta_{ext} \downarrow \quad R_c \downarrow \quad \eta_{wall} \uparrow \downarrow ?$$

Figure 3.7: Optimizing LED wallplug efficiency

The external quantum efficiency, η_{ext} , wallplug (power) efficiency, η_{wall} , and contact resistance, R_c , are related to the active device area, A_{act} , the ohmic contact area, A_c , the junction voltage, V_J , and the specific contact resistance, r_c . Increasing the contact area lowers the external quantum efficiency by blocking more light, but it also lowers the contact resistance, reducing the terminal voltage. For a given operating current and specific contact resistivity, which is determined by the ohmic contact metallization, there is a certain contact area which optimizes the wallplug efficiency.

the mesa for a via structure leading to the interconnect metallization. The mesa area was thus made as small as possible while allowing adequate room for an ohmic pattern which leaves most of the mesa unblocked.

The shape of the ohmic is also significant in that it effects current spreading. It is desirable to have injected current spread away from the ohmic so that the generated light is not blocked. A number of different ohmic patterns of varying size and shape were tested in conjunction with the 30 μm X 30 μm mesa. These are shown in Table 3.1 along with the ohmic contact area for each. To assess the current spreading characteristics of the various designs, LEDs were viewed with a CCD camera while in operation. Several representative emission patterns are shown in Figure 3.8. It is seen that the annular contact provides good current spreading while blocking less mesa area than the multi-finger designs. The cross pattern does not promote good current spreading.

The I-V and L-I characteristics of LED with the various contact are shown in Figure 3.9. Notice the threshold in the L-I curves and the corresponding kink in the I-V curves which originate from the non-radiative to radiative crossover mentioned above. It must be noted that the light output reported in Figure 3.9 is









Contact Structure	Contact Area (μm^2)	Pattern
A	329	
B	197	
C	225	
D	326	
E	413	
F	169	
G	253	
H	297	

Table 3.1: Various ohmic contact geometries tested in conjunction with a $30\ \mu\text{m}$ square mesa

Basic annular designs of varying contact area (A, C, F) were augmented with multi-finger designs (D, E, G, H) and a cross pattern (B) in order to study the effect of contact area and geometry on contact resistance, current spreading, and quantum efficiency. Further details may be found in [14].




Figure 3.8: Emission pattern from LEDs made with various contact geometries

LEDs were fabricated on bulk material with $30\ \mu\text{m}$ square mesas and the contact geometries indicated. The LEDs were biased at $0.5\ \text{mA}$ and their emission was viewed through a microscope with a Si CCD camera. The annular pattern (F) is seen to produce good current spreading. It is evident that the cross pattern (B) does not support good current spreading, and that the multi-finger design (G) does not give any advantage over the basic annulus.

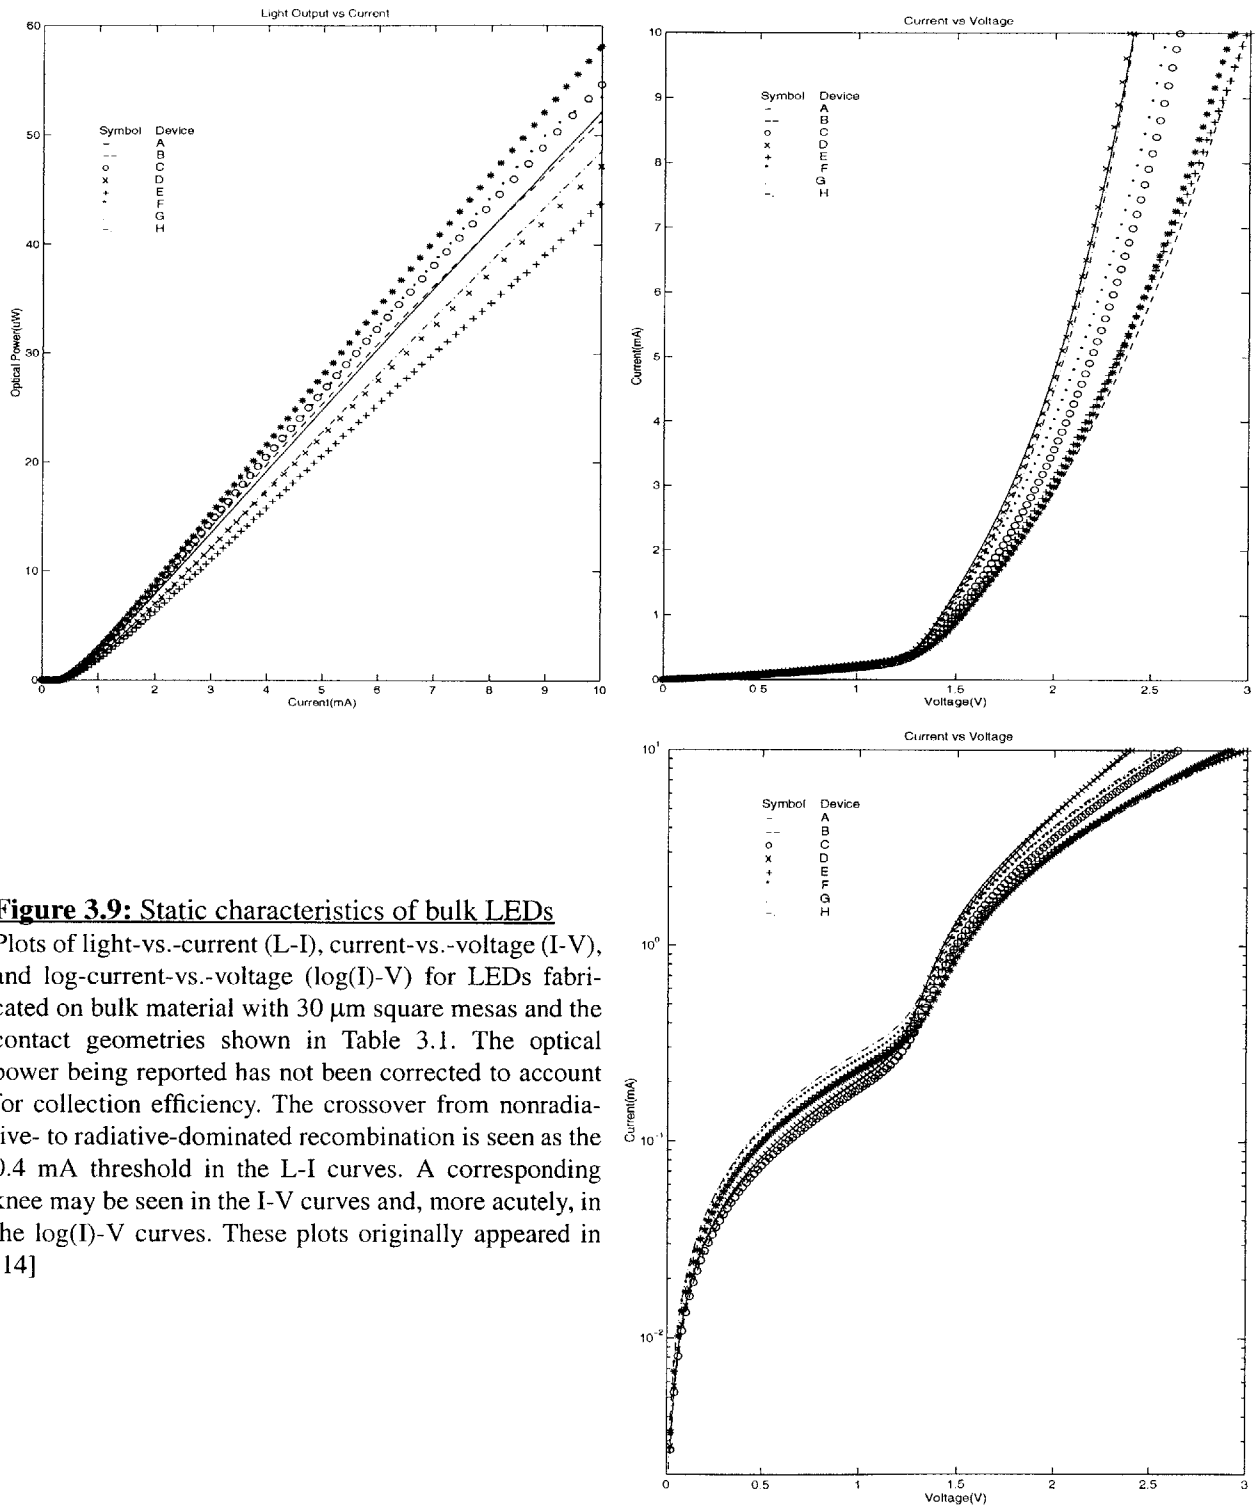


Figure 3.9: Static characteristics of bulk LEDs
 Plots of light-vs.-current (L-I), current-vs.-voltage (I-V), and log-current-vs.-voltage (log(I)-V) for LEDs fabricated on bulk material with $30\ \mu\text{m}$ square mesas and the contact geometries shown in Table 3.1. The optical power being reported has not been corrected to account for collection efficiency. The crossover from nonradiative- to radiative-dominated recombination is seen as the 0.4 mA threshold in the L-I curves. A corresponding knee may be seen in the I-V curves and, more acutely, in the log(I)-V curves. These plots originally appeared in [14]

not calibrated to reflect the collection efficiency of the measurement setup. A more careful assessment has shown some devices to achieve an external quantum efficiency in excess of $12 \mu\text{W}/\text{mA}$ (0.85%).

Figure 3.10 allows the relative optimality of the various contact designs to be judged (note: collection efficiency is not calibrated). Notice that while contact F, which is an annular contact with the smallest contact area, has the largest quantum efficiency (top-right plot), it is device A, which is also annular but with a significantly larger area, and thus a lower contact resistance, that gives the highest power efficiency and the lowest operating voltage. An ohmic contact similar to A was thus chosen for the integrated LEDs used in the OPTOCHIP project.

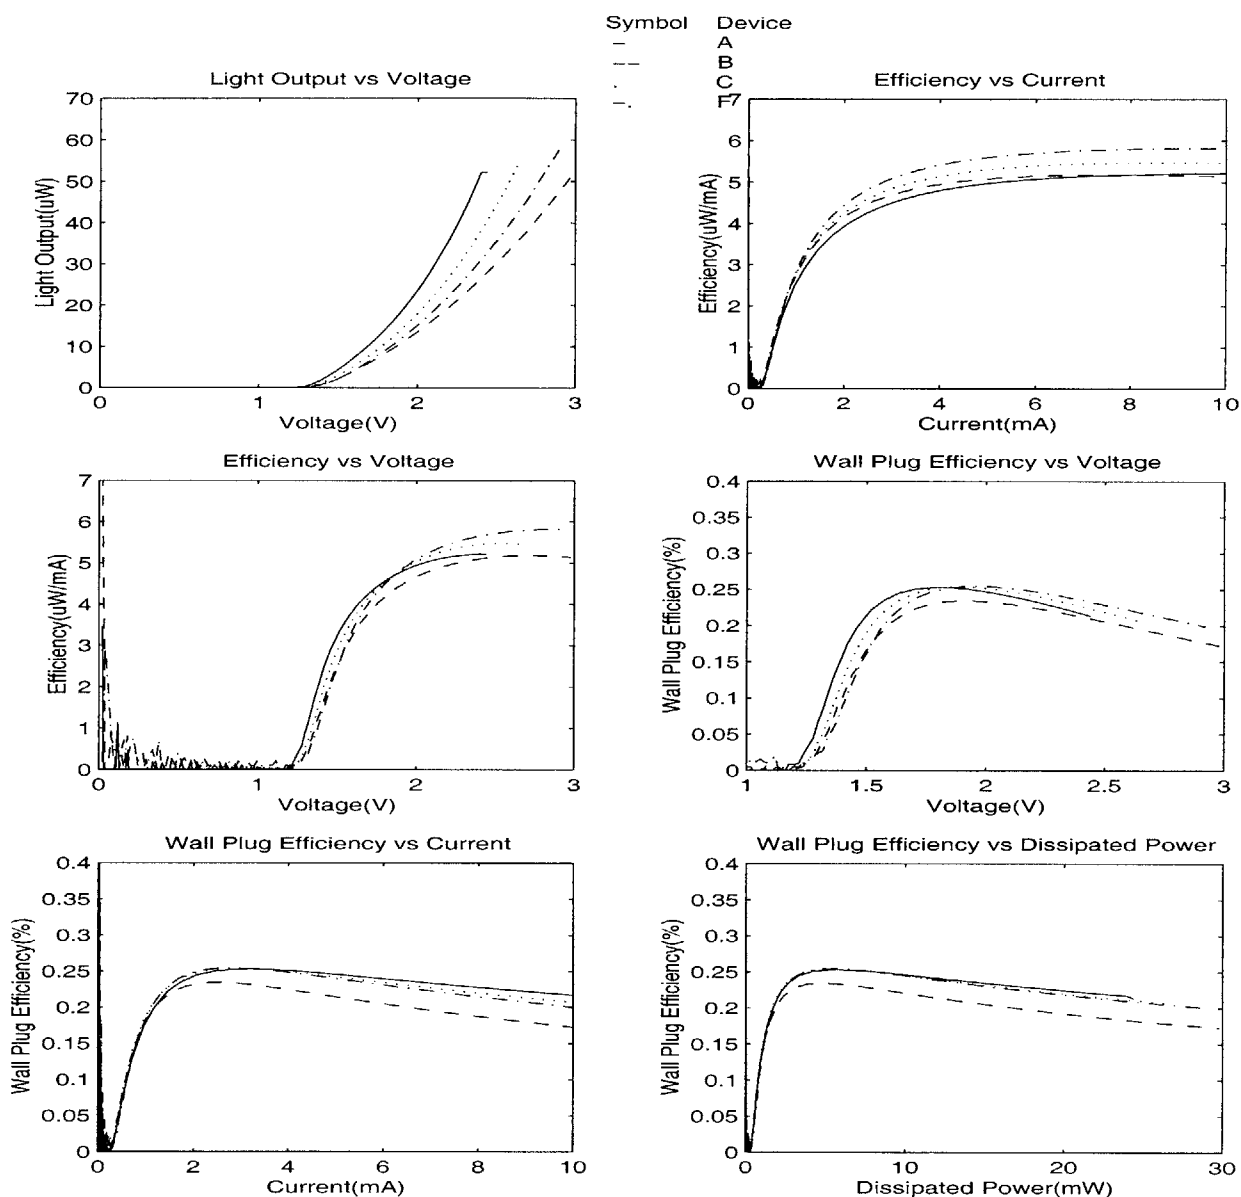


Figure 3.10: LED contact geometry selection

Figure 3.8 showed that the annular contact geometries (devices A, C, and F in the above plots and in Table 3.1) have good current spreading characteristics. The above plots compare these devices, along with the cross pattern B, in order to optimize the contact area. The light output and efficiency above are not corrected to account for the light collection efficiency of the measurement setup. “Efficiency” is defined above as the output optical power divided by the terminal current, while the “wall-plug efficiency” is defined as the output optical power divided by the total LED power dissipation. It is seen that while device F, which has the smallest contact area, is optimal with regard to efficiency at a given current level, device A, with the largest contact area, is actually preferable in all other scenarios. Specifically, it produces more light at a lower voltage and achieves a higher wall-plug efficiency under most circumstances. This is due to its lower contact resistance. These plots originally appeared in [14]

3.3 The OPTOCHIP project

One of the barriers to the advancement of optical interconnects has been the unavailability of complex OEICs to optical interconnect system researchers. Although no commercial source exists for OE-VLSI components, large-scale collaborations between OEIC technologists and users can catalyze optical interconnect development. A notable example of such an effort has been the AT&T/DARPA hybrid-SEED smart-pixel workshop which made OEICs based on flip-chip bonded multiple-quantum-well modulators available to the participants [139]. A similar research foundry, the OPTOCHIP project, has delivered emitter-based OEICs fabricated with the EoE process. The goals of OPTOCHIP have been two-fold. First, to advance the development of optical interconnects by providing much-needed OEICs to system researchers, and second, to refine the EoE technology on the basis of valuable feedback from the designers.

In the OPTOCHIP project, the EoE process was carried out to fabricate OEICs which combined monolithically integrated GaAs/InGaP LEDs with the electronics and photodetectors available on the standard Vitesse H-GaAs III process. As outlined in Figure 3.11, nine research groups involved in optical interconnect system development took part in the project, with two groups combining their designs in a shared area. The project was kicked-off at an informal meeting of the participating groups in February 1996. Each of the participants completed the design of a 2 mm x 2 mm die. Final layouts of the eight contributed designs were collected at MIT in April, combined with a ninth 2 mm x 2 mm design, and tiled into a 3x3 array to form a "master" die. With scribe line patterns added around each submission, the master die measured 7.3 mm x 7.3 mm. This layout was submitted to the MOSIS service for inclusion on the May 1996 H-GaAs III fabrication run [122]. The electronic master die were returned by MOSIS in August. Following equipment related delays, growth and fabrication of integrated LEDs began in January 1997. A total of five of the master die were processed for the OPTOCHIP project. A finished master die is pictured in Figure 3.12. The completed masters were sawed to separate the individual 2 mm x 2 mm designs. The last of these completed OEICs was shipped to the OPTOCHIP participants in May 1997. The various delays encountered in the one year period from the date the design was submitted to MOSIS until chips were sent to the participants added over six months to the cycle. Under more optimal circumstances the sequence could be executed, in a research environment, in as little as three months.

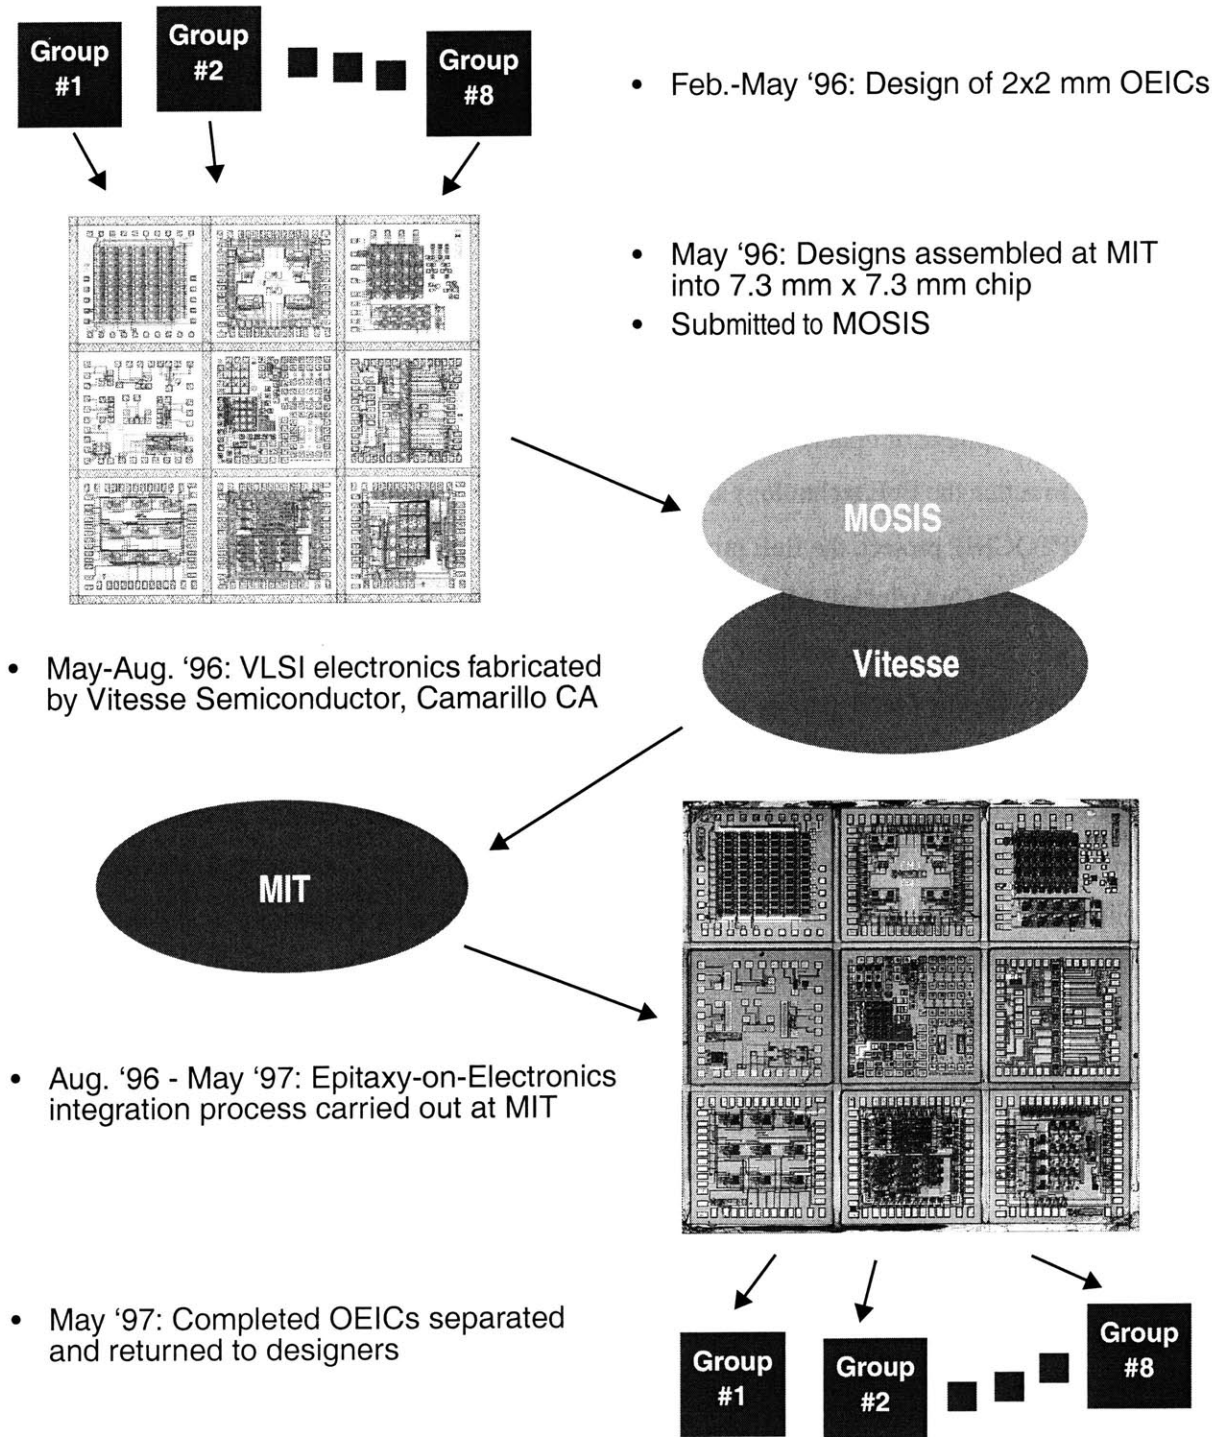


Figure 3.11: The OPTOCHIP project




Figure 3.12: Photograph of a finished OPTOCHIP master die

The master die is composed of nine 2 mm x 2 mm designs. With the addition of scribe lines, the total master die measures 7.3 mm x 7.3 mm. The nine designs are identified in Figure 3.18.

Design tools

The design of EoE-OEICs for the OPTOCHIP project followed a standard flow common in electronic design and used standard VLSI CAD tools. The layout of the electronic circuitry, implemented using enhancement- and depletion-mode MESFETs and four metallization layers, followed the H-GaAs III design rules established by Vitesse. Layouts were completed in the designers' preferred tools, including Cadence, Magic, and LEdit, and were simulated in HSPICE with device models provided by Vitesse. A number of groups used software from Cascade Design Automation Corporation (now Duet) to synthesize circuit layouts on the basis of a high-level hardware description language. The Cascade suite also included input/output pad drivers and a pad frame for the Vitesse process. Others employed full-custom circuit layout and a custom pad frame.

Optoelectronic standard cells

In addition to the design rules and HSPICE models which support the electronic design, the OPTOCHIP participants were provided with standardized drop-in cells which provided optical functionality. These included metal-semiconductor-metal (MSM) and optically sensitive FET (OPFET) photodetectors implemented as part of the H-GaAs III process, LEDs, and an LED driver circuit.

The LED cell included all layers needed at the Vitesse foundry and in the EoE process, and the designers were required to make electrical connections to the LED using metal-1 for the n-contact and metal-4 for the p-contact. The standard OPTOCHIP LEDs used a 50 μm x 50 μm dielectric growth well (DGW) surrounded by an optical shield. The layout of the LED drop-in cell is shown in Figure 3.13 along with a photograph of a completed OPTOCHIP LED. The EoE design rule regarding placement of the LED cells stipulated only that the surrounding layout geometries, with the exception of the electrical connections to the n- and p- contacts, be placed beyond the optical shield by the minimum space set in the Vitesse design rules. Although not a requirement, most LEDs in a typical design were connected to a common metal-4 p-contact designated as the LED supply rail and were modulated by switching the current at the n-contact.

The LED driver circuit, shown in Figure 3.14, is similar in design to the first-generation laser driver described in Chapter 5.

The MSM photodetector drop-in cell is shown in Figure 3.15. These devices will be dealt with in Chapter 4. Briefly, they are implemented by forming interdigitated gate-metal fingers over a MESFET

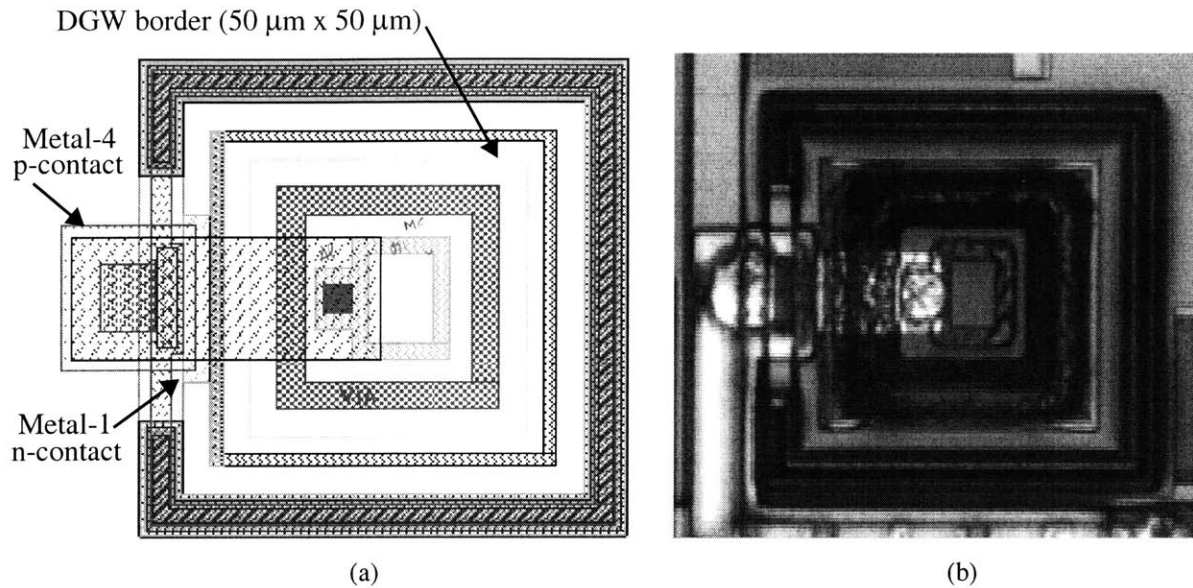


Figure 3.13: LED drop in cell used in OPTOCHIP

(a) The layout provided to the designers, which includes all Vitesse layers needed to implement the DGW and optical shield, as well as all EoE layers needed to fabricate an integrated LED. Electrical contact to the LED is made via metal-1 and metal-4 at the points indicated. (b) A photomicrograph of an OPTOCHIP LED.




Figure 3.14: LED driver drop in cell used in OPTOCHIP

Schematic (a) and layout (b) of an LED driver circuit provided to the OPTOCHIP designers [140]. The operation of this circuit is similar to the basic switched-current-mirror laser driver described in Chapter 5.




Figure 3.15: Layout of an MSM photodetector drop-in cell used in OPTOCHIP

Two versions of a drop in MSM photodetector cell were implemented using $20\ \mu\text{m}$ square active regions and $1.2\ \mu\text{m}$ and $1.0\ \mu\text{m}$ interdigitated gate metal fingers spaced $1.6\ \mu\text{m}$ and $1.8\ \mu\text{m}$ apart, respectively. An $8\ \mu\text{m}$ wide ring of source/drain n^+ implant surrounded the cell to provide backgate isolation.

active region. MSM detectors of this design will operate at several hundred Mb/s and have a responsivity of approximately $0.2\ \text{A/W}$. Typically, a transimpedance amplifier is required to use an MSM in an optical receiver.

An OPFET photodetector is implemented using a floating gate enhancement-mode MESFET. The OPFET drop-in cell is shown in Figure 3.16(a). This detector may be combined with a diode connected depletion-mode MESFET load to form an optical-input direct-coupled FET logic (DCFL) inverter, as shown in Figure 3.16(b). Optical input logic thresholds of $0.1\ \mu\text{W}$ are readily achieved, making OPFET receivers well-matched to LED transmitters. However, OPFET bandwidths are typically 10-100 KHz. Many of the OPTOCHIP designs have made use of OPFET receivers and are thus limited in speed by the detector rather than the LED.

Wavelength

As will be quantified below, the OPTOCHIP LEDs emit at $873\ \text{nm}$ corresponding to band-to-band emission from the GaAs core. A wavelength of $\sim 850\ \text{nm}$ is more common in short-haul optical interconnects (it is well detected by both GaAs and Si detectors), but emission at $850\ \text{nm}$ requires the growth of a quaternary InGaAsP material. This does not pose an EoE-specific problem, but does require greater control in




Figure 3.16: Optically-sensitive FET (OPFET) drop-in cell used in OPTOCHIP

As shown in the layout in (a), two OPFET photodetectors were implemented with $40\ \mu\text{m}$ square active regions and $1.2\ \mu\text{m}$ and $1.0\ \mu\text{m}$ floating gates. The cell was surrounded by an $8\ \mu\text{m}$ wide ring of source/drain n^+ implant to provide backgate isolation. The schematic in (b) shows how the OPFET may be used as part of an “optical input inverter”. In this configuration, the OPFET takes the place of the drive EFET of a DCFL (direct coupled FET logic) inverter. With a DFET W/L ratio of $6\ \mu\text{m} / 18\ \mu\text{m}$ and a $1.2\ \mu\text{m}$ OPFET gate, the optical inverter achieves a logic threshold of $0.1\ \mu\text{W}$ using $850\ \text{nm}$ incident light.

growth conditions as compared to the ternary InGaP or the binary GaAs. Since the $873\ \text{nm}$ emission from the basic GaAs core may be detected with the GaAs-based photodetectors available on the Vitesse process, though with somewhat reduced sensitivity, the more difficult quaternary growth was not implemented for this demonstration project. A plot of the spectral characteristics of an H-GaAs III MSM photodetector is shown in Figure 3.17.

The OPTOCHIP designs

The designs completed in the OPTOCHIP project include applications in free-space optical interconnects and smart pixels, image processing, neural networks, and sensing. Figure 3.18 shows the OPTOCHIP groups, including descriptive project titles and the principal investigators. The position of each group in this figure corresponds to their location on the master die.

The designs varied substantially in terms of their use of digital and analog electronics, OPFET and MSM photodetectors, LED density, and circuit complexity. Each master die contained a total of 195 inte-




Figure 3.17: GaAs photodetector optical responsivity spectrum

Measurements were made on a Vitesse H-GaAs III MSM photodetector and are normalized to unity at 880 nm incident wavelength [141]. The same curve applies to the OPFET photodetectors. The GaAs active region of the integrated LEDs emit light at 873 nm, and may thus be used in conjunction with the Vitesse photodetectors.

grated LEDs with the number of LEDs on each submission varying according the designers' goals. The highest LED density implemented on OPTOCHIP, $6800/\text{cm}^2$, appears on a demonstration circuit to be described in the next section. The highest density smart pixel array implemented was a 6×11 , $4444 \text{ pixel}/\text{cm}^2$, analog optical neural network element designed at Cal Tech and shown in Figure 3.19 [145]. Each pixel in this example consists of one LED, two OPFET detectors, and a 20-transistor processing circuit.

An additional, noteworthy OPTOCHIP example is the network interface chip ("SAPIENT") implemented at USC and occupying the lower half of the OEIC in Figure 3.20 [144]. This digital circuit was

<p><i>High order correlation connectionist network</i></p> <p>D. Psaltis, California Institute of Technology</p>	<p><i>Optical neural network</i></p> <p>M. W. Haney, George Mason University</p>	<p><i>High order correlation connectionist network</i></p> <p>C. W. Wilmsen, Colorado State University</p>
<p><i>Dynamic smart pixels</i></p> <p>F. A. P. Tooley, McGill University</p>	<p><i>EoE Test Cells</i></p> <p>MIT</p>	<p><i>Sensor applications of integrated emitter-detector arrays</i></p> <p>W. R. Babbit, University of Washington</p> <p>[142,143]</p>
<p><i>Optoelectronic error diffusion neural network</i></p> <p>B. L. Shoop, United States Military Academy</p> <p>[20-23]</p>	<p><i>OPTOCHIP Implementation of a Smart Pixel Array Network Interface Chip (SAPIENT)</i></p> <p>A. A. Sawchuk, University of Southern California [144]</p> <p>&</p> <p><i>Optical multi-channel interconnection network interface</i></p> <p>T. M. Pinkston, University of Southern California</p>	<p><i>OPTOCHIP controller interface for a distributed ensemble of remote sensors</i></p> <p>L. Cheng, Texas Christian University</p>

Figure 3.18: OPTOCHIP participants

The participants position in this chart corresponds to the location of their design in the master die shown in Figure 3.12.

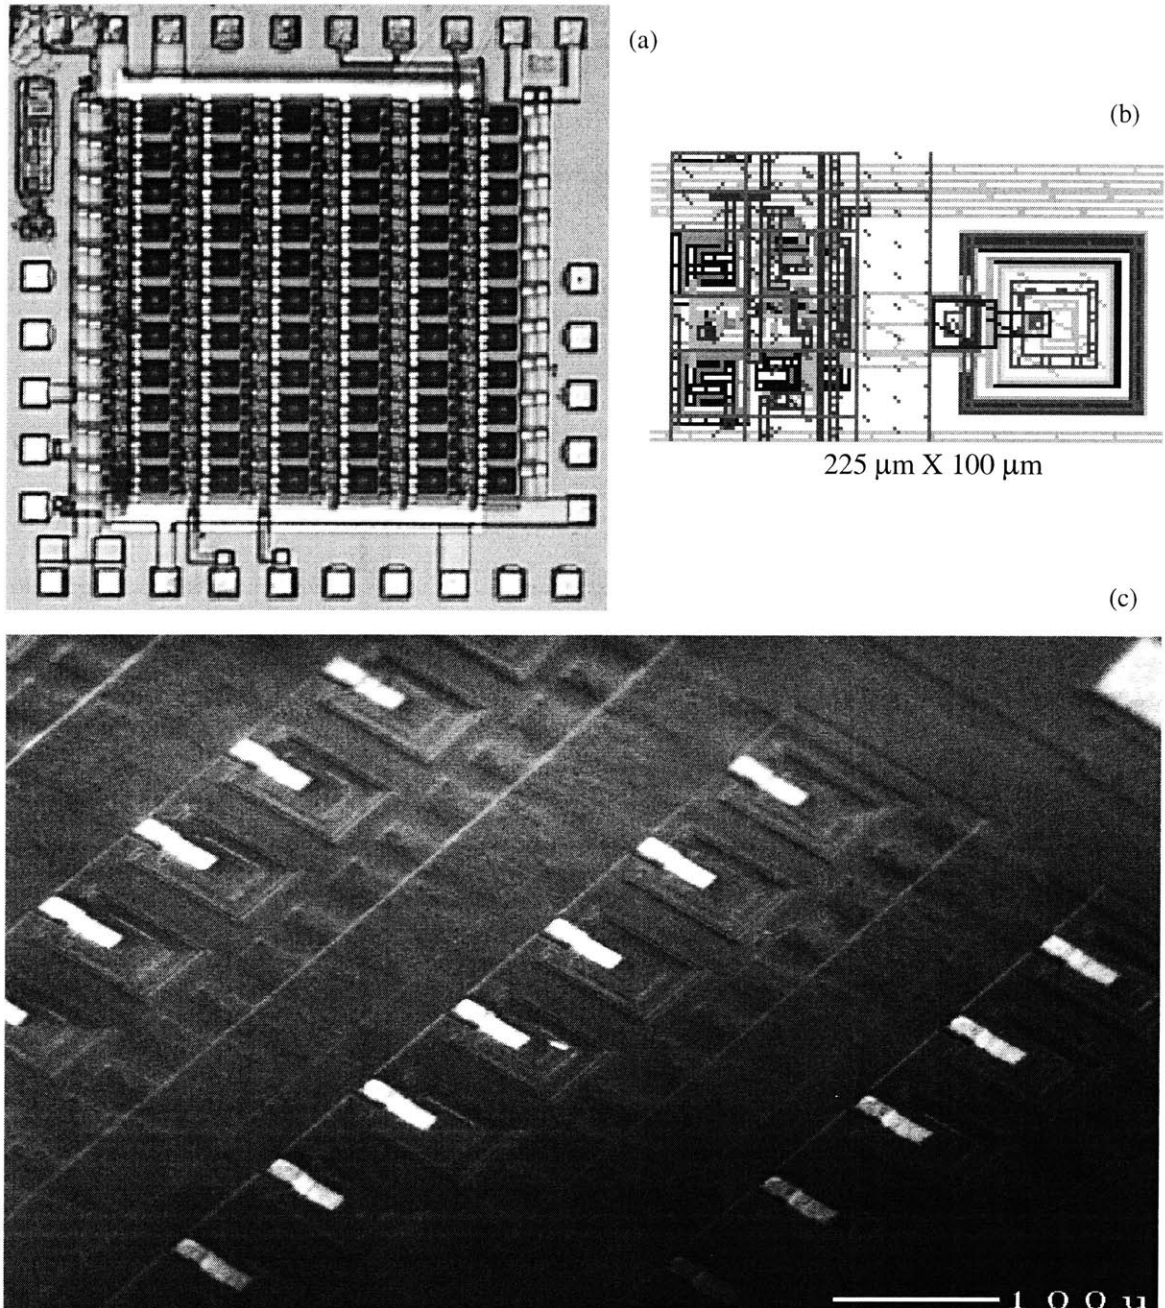


Figure 3.19: OPTOCHIP OEIC example: optical neural network array

A 2 mm x 2 mm OPTOCHIP design of an optical neural network smart pixel array [145]. Each pixel consists of one LED, two OPFET photodetectors, and a 20-transistor processing circuit. The 6x11 array achieves a density of 4444 pixels/cm². (a) is a photograph of a finished die, (b) is a layout of the unit cell used for form the array, and (c) is an SEM view of a portion of the array.

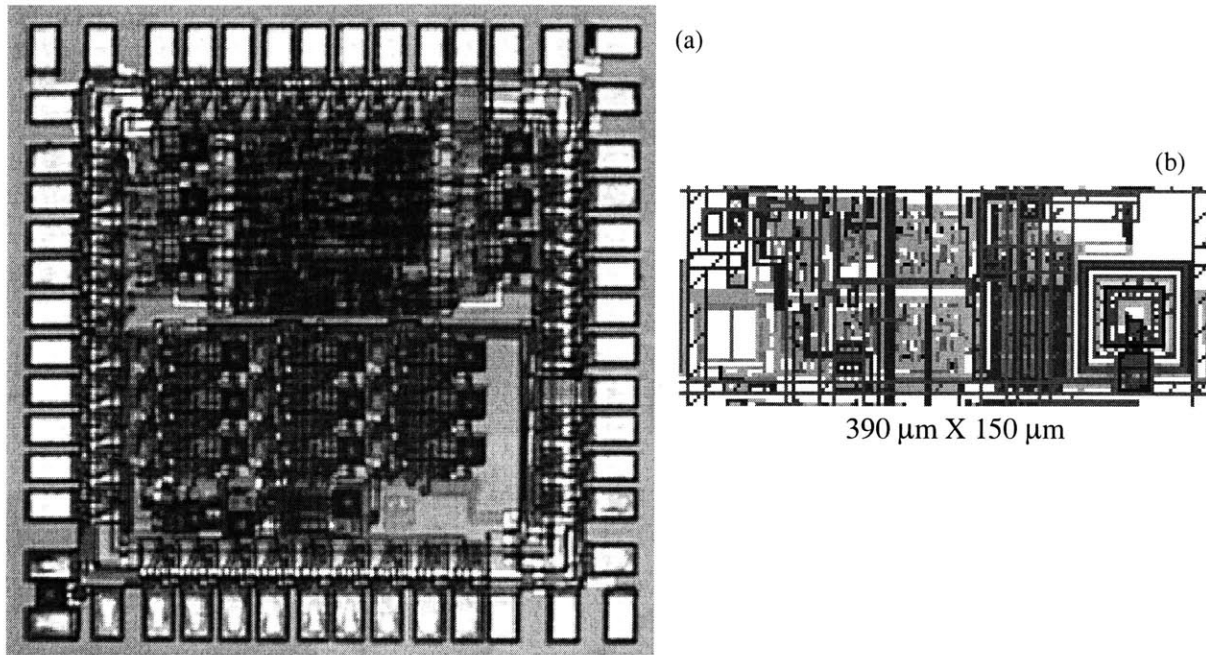


Figure 3.20: OPTOCHIP OEIC example: optical network interface

The OPTOCHIP submission from USC includes two designs, seen in top and bottom halves of the die photo in (a), and utilized computer-synthesized layout based on a hardware description language. The lower design, known as “SAPIENT” (SmArt PxEI Network inTeface), has been deployed in an optical interconnect demonstration as shown in Figure 3.21. The layout of the unit cell which makes up the SAPIENT smart-pixel array is shown in (b) [144].

designed in a hardware description language from which the layout was synthesized. Two of these OEICs were deployed in an optical system, shown in Figure 3.21, and test data was transmitted between them. As will be pointed out in the next two sections, this demonstration was limited by issues of LED yield and optical crosstalk. Notwithstanding these limitations, it is a good example of the rich capabilities of the EoE technology and represents a milestone in the technology’s development.

Summary

The OPTOCHIP project has addressed the goal of advancing optical interconnect development by providing emitter-based OEICs of unprecedented density and complexity to the user community. In this initial research foundry offering, much of the participants’ efforts were focused on qualifying the EoE technology for future full-scale optical interconnect demonstrations. Free space optical interconnection between a pair of OPTOCHIP circuits has been demonstrated by one of the OPTOCHIP groups. The current LED-based

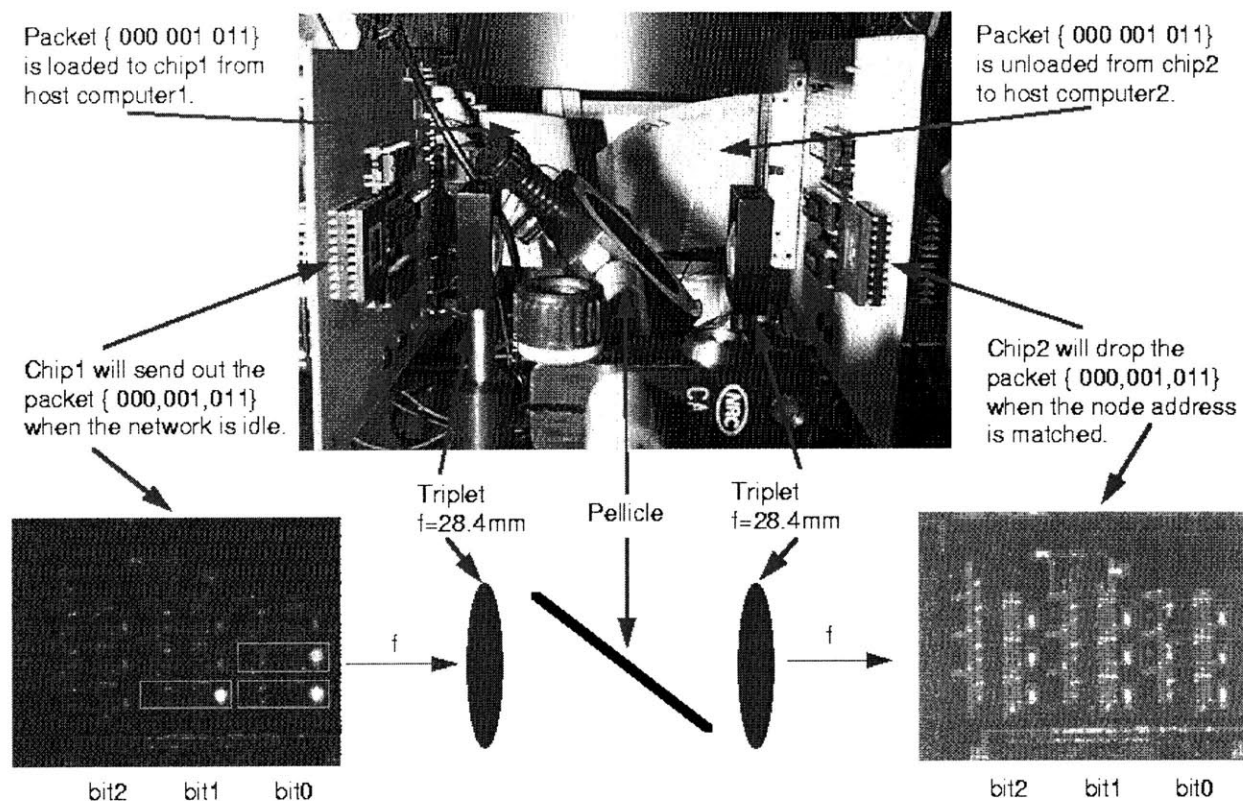


Figure 3.21: Overview of the SAPIENT optical interconnect demonstration

The OPTOCHIP implementation of “SAPIENT” (SmArt P*IX*El Network inTerface) was deployed in an optical interconnect demonstration by its designers at USC. This figure, taken from reference [144], shows the physical implementation of the link in which two of the OEICs shown in Figure 3.20 are used to transmit packets of binary data through the optical system.

OEICs are adequate for a basic implementation of many optical interconnect architectures. The performance of these demonstrations is limited, however, by the characteristics of the LEDs and photodetectors to be presented in Section 3.4 and by optical crosstalk effects explained in Section 3.5. Notwithstanding these limitations, the experience of the OPTOCHIP participants has validated the design flexibility of the EoE technology. Their feedback and the device and circuit results will be examined in Section 3.6.

3.4 Integrated LED and OEIC results

The majority of the results presented in this section are based on test cells found on the MIT portion of the OPTOCHIP master die. A layout of the MIT design is shown in Figure 3.22. Additional results are based on the West Point OPTOCHIP design (lower left corner of Figure 3.12 and Figure 3.18), and on observations of the OPTOCHIP participants.

A scanning electron micrograph of an integrated LED in a 50 μm DGW is shown in Figure 3.2. The optical shield around the DGW perimeter is formed by stacking interconnect metals up to Metal-3. The purpose of this shield, as discussed in Section 3.5, is to prevent coupling of LED emission to nearby electronics. The present section is concerned only with the effect of the shields on LED performance. Two shield designs (referred to as #1 and #2) differing in the sizing and placement of metal and via patterns, and consequently differing in surface non-planarity above the shield, were implemented around 50 μm DGWs. Unshielded 85 μm DGWs were also used.

In many of the 50 μm DGWs, the material within a ~ 7 μm border of the DGW was seen to be rough, as in Figure 3.23. This effect was absent in the 85 μm DGWs. Examination of five completed OPTOCHIP OEICs suggests that the extent and frequency of this border depends both on the DGW size and on the surface non-planarity of the surrounding shield and electronics. It is believed that the roughness originated during the DGW formation stage of these chips as a result of poor photolithography which led to incomplete removal of the photoresist along the DGW perimeter. In earlier work, growth was obtained in DGWs as small as 10 μm square [2], and 50 μm DGWs have been used both before and after the OPTOCHIP project, and with both AlGaAs- and InGaP-based devices, without similar complications. More work is required to determine the minimum practical DGW size and the maximum density with which the LEDs and electronics may be placed.

Wavelength and directionality

The angular distribution of the output light from the integrated LEDs has been found to be Lambertian (intensity proportional to the cosine of the normal angle), as plotted in Figure 3.24 [146]. The angular distribution is important in determining the coupling efficiency of the LED emission into an optical system. Knowledge of the angular distribution is also used to calibrating the output power measurements shown below.

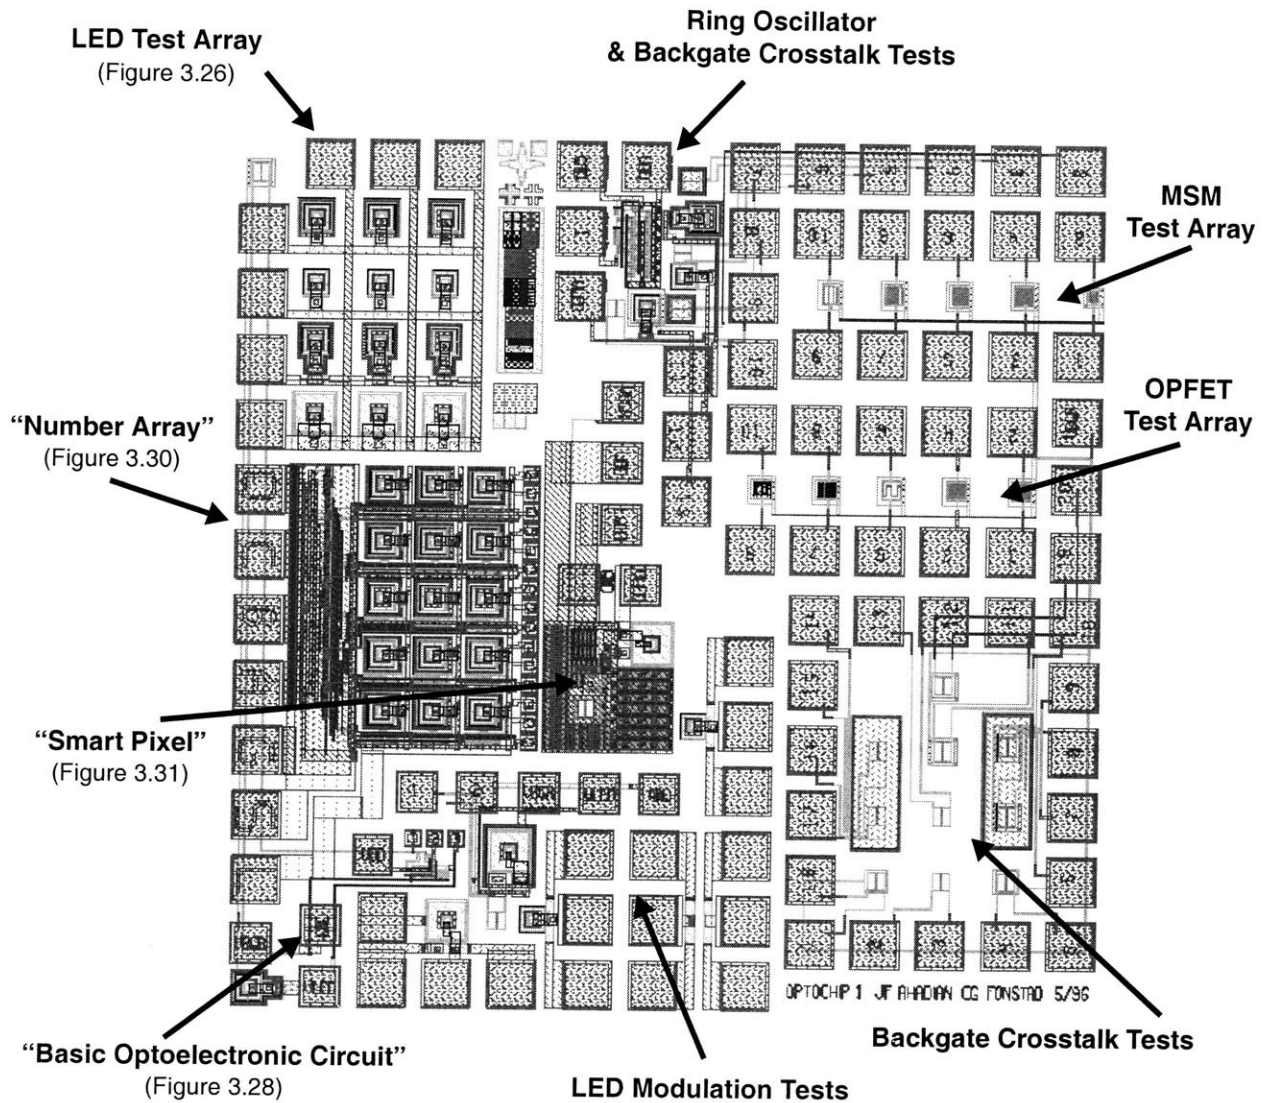


Figure 3.22: Layout of the MIT OPTOCHIP design

The 2 mm x 2 mm MIT design, at the center of the OPTOCHIP master die in Figure 3.12, consists of a number of test cells. These include the three optoelectronic circuit demonstrations shown in this section, test structures used to measure the static and dynamic properties of the LED, the ring oscillator used in measurements mentioned in Section 2.8., and additional cells used to characterize the MSM and OPFET photodetectors and the backgate crosstalk effects taken up in Section 3.5.

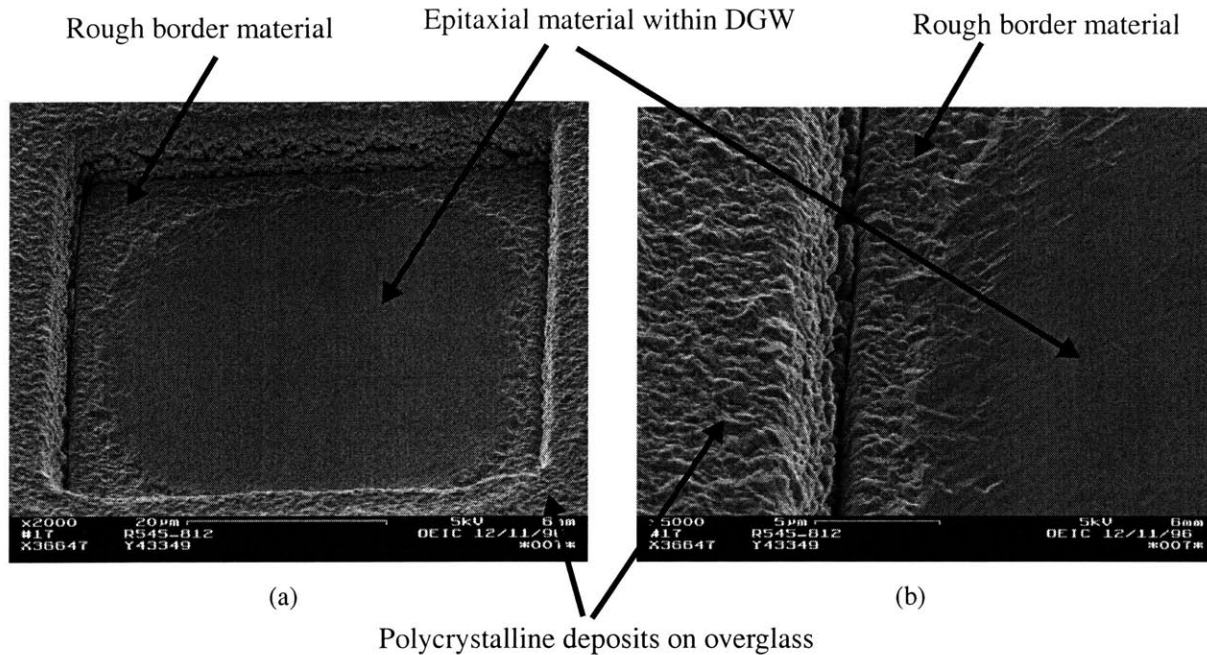


Figure 3.23: Poor material quality at perimeter of 50 μm DGWs

SEM of 50 μm DGW following MBE growth (a) and close up of left edge of same DGW (b). The polycrystalline deposits on the surface of the OEIC have not yet been removed. Contrast the smooth morphology of the epitaxial material at the center of the DGW with the rough surface along the border of the DGW region. The material is believed to be degraded in these areas due to incomplete removal of the photoresist during the DGW preparation process.

The emission spectrum of an integrated LED operating at three power levels is shown in Figure 3.25 [23]. As determined from a measurement of ten LEDs taken from two OPTOCHIP master die, the LED emission is centered at a wavelength of 873 ± 1 nm and has a spectral full-width-at-half-maximum (FWHM) of 31 ± 2 nm, or roughly 50 meV [20]. As expected, this is consistent with the thermally broadened band-to-band emission from the unintentionally-doped GaAs core, and it agrees with electroluminescence spectra taken on bulk LED material [10]. The spectral peak shifts by only 3 nm over a wide operating range, and the uniformity of the spectral characteristics is seen to be quite good.

Light-Current-Voltage characteristics

The current/voltage/light-output characteristics of typical, working LEDs in each of the three types of DGWs are shown in Figure 3.26. In each of the three sets of data, the LED current begins to gradually rise




Figure 3.24: LED angular distribution

Plot of the angular distribution of the light emission from an LED fabricated in an 85 μm DGW. The solid line is a plot of the cosine of the normal angle and is known as a “Lambertian” distribution [146].

starting at 1-2 V. At an LED current of 1.5 mA, the rate of current rise increases dramatically. Light emission begins at this “knee” current. These characteristics are similar to those of the bulk LED discussed in Section 3.2. The primary difference is in the value of the knee current--1.5 mA in this case rather than 0.4 mA in the case of the bulk LEDs. The voltage corresponding to the knee varies from 1.5 V to 3 V in the data given. The overall device performance is characterized by the value of the knee current and the differential external quantum efficiency above the knee.

The 1.5 mA knee current corresponds to a current density of 170 A/cm^2 within the 30 μm x 30 μm LED mesas. Electroluminescence studies on bulk LED material grown alongside the OEICs exhibited a comparable knee current density. This tends to indicate that the knee current difference between these LEDs and those shown in Section 3.2 originates in the material growth rather than in the integration process.




Figure 3.25: Optical spectra of integrated LED emission

The optical spectrum of an OPTOCHIP LED is shown at three bias points [23]. Similar measurements on ten LED on two copies of the West Point OPTOCHIP give the mean emission wavelength as 873 ± 1 nm and the FWHM as 31 ± 2 nm, or roughly 50 meV. This is consistent with thermally broadened band-to-band emission from the GaAs core of the LED.

The bulk material grown alongside the OEICs was broken during LED processing, precluding a direct comparison of fully-processed bulk and integrated LED efficiencies above the knee. Instead, the differential external quantum efficiency of the integrated LEDs may be compared with a theoretical limit of 1%. This limit is set by the total internal reflection of 98% of the spontaneously emitted light within the device, and by metal coverage of 50% of the active area. At a current of 10 mA (8.5 mA above the knee) the differential external quantum efficiency is 1.0% for the 85 μm DGW case and 0.3% and 0.6% for 50 μm DGWs with shield styles #1 and #2, respectively.

The lower efficiency of the 50 μm DGW LEDs is presumably linked to their rough border. While better process control will be needed to prevent a repeated occurrence of this problem in the 50 μm DGWs,




Figure 3.26: Typical static characteristics of integrated LEDs on OPTOCHIP

Voltage vs. current and light output vs. current characteristics of EoE-integrated GaAs/InGaP LEDs. Results are given for LEDs in unshielded 85 μm DGWs (solid lines) and for LED in 50 μm DGWs with two different optical shield styles (dashed and dotted lines).

the 85 μm DGW results are already significant. The high differential efficiency of the 85 μm DGW LEDs indicates that optoelectronic devices integrated with VLSI electronics using the EoE process can retain the performance achieved by similar discrete devices.

Uniformity

The variability of the rough border in the 50 μm DGWs translated into poor uniformity among the LEDs fabricated in these DGWs. Greater uniformity is found among the 85 μm DGW LEDs. Out of twelve 85 μm DGW LEDs measured on four OEICs, only one device failed due to a damaged interconnect. At 10 mA, the remaining eleven LEDs had a power efficiency mean and standard deviation of $0.3 \pm 0.1\%$. The 30% variability in this figure is a manifestation of the poor process control afforded by the handling of bare die in a research facility and can be expected to improve in a manufacturing environment using full wafers.

The uniformity of the 50 μm DGW OPTOCHIP LEDs is significantly lower than the 85 μm DGW result. In this case, the large variability in the border roughness in the 50 μm DGWs translated into significant variability in the OPTOCHIP LEDs. Figure 3.27 shows L-I plots of 18 LEDs on two copies of the




Figure 3.27: Uniformity of OPTOCHIP LEDs

Light-vs.-current measurements were made on eighteen LED taken from two copies of the West Point OPTOCHIP [23]. Around one third of these devices have characteristics similar to those of Figure 3.26. Another third of the devices were half as efficient, and the remaining third functioned very poorly. These uniformity characteristics are a result of the rough border material shown in Figure 3.23.

OPTOCHIP OEIC designed at West Point [23]. It is apparent that approximately one third of these LEDs exhibited the typical characteristics shown in Figure 3.26. Another third were half as efficient, and the remaining third performed very poorly. This level of variability has hindered a number of the OPTOCHIP designs which depend on emitter uniformity for their essential functionality.

Modulation bandwidth

The small-signal bandwidths of integrated LEDs in both 50 μm and 85 μm DGWs was measured by Vaidyanathan, and found to vary from 30 MHz to 40 MHz with increasing bias [14]. The low bandwidth of




Figure 3.28: “Basic Optoelectronic Circuit”

The FET dimensions are specified as W/L [$\mu\text{m}/\mu\text{m}$]. An optical-input direct-coupled FET logic (DCFL) inverter drives a chain of three cascaded DCFL inverters and a pull-down EFET. This EFET modulates the integrated LED current and, therefore, the light output. When an optical input is applied, the LED is turned on.

these unoptimized LEDs is due to the thickness of the emission region and due to the lack of p-type doping in this region. Improved design could increase the LED bandwidth to over 100 MHz [14].

“Basic optoelectronic circuit”

The schematic of an optoelectronic circuit consisting of an OPFET photodetector, GaAs electronics, and an integrated LED is shown in Figure 3.28. As seen in Figure 3.22, the OPFET was placed at the upper left corner of the MIT OPTOCHIP design while the LED was located at the lower left corner. The OPFET, a 40/1 (W/L [$\mu\text{m}/\mu\text{m}$]) floating-gate EFET, is loaded by a diode-connected DFET to form an optical-input direct-coupled FET logic (DCFL) inverter. The output of this stage, labeled “Optical Inverter” in the figure, passes through three successively larger DCFL inverters before driving a 50/1 EFET that modulates the LED. The drain of this EFET is labeled “EFET Pull-down” in the figure. The LED in this circuit uses a 50 μm DGW with a style #2 optical shield. The characteristics of this LED are similar to those shown in Figure 3.26, with the electrical turn-on at 2 V and the knee current corresponding to 4 mA.

Correct operation of this circuit is demonstrated in Figure 3.29. To characterize the circuit, 850 nm light from a laser diode was focused onto the OPFET through a microscope objective, and the laser power was varied. With no applied light, the OPFET is in its off state and the DFET load pulls the optical inverter




Figure 3.29: Response of circuit shown in Figure 3.28 to an applied optical input.

output to the DCFL high level of 0.6 V. Following three inversions, the signal arrives at the gate of the pull-down EFET as a low. The drain of this off-state EFET rises to 3 V, and 2 V remains across the LED. There is no significant LED current and the optical output is off. When the incident power reaches just above 0.1 μW , the OPFET switches on. The optical inverter output is now pulled to a DCFL low level of 0.1 V. This signal cascades through the three inverter stages and turns on the pull-down EFET. The drain of this EFET now drops to 0.8 V. With 4.2 V across the LED, the EFET sinks 4.6 mA of current corresponding to an optical output power of 2 μW . This power level is 13 dB above the 0.1 μW optical input threshold. This regenerative effect, or “optical gain”, allows EoE-integrated LED-based OEIC to meet the realistic system requirements of many optical interconnect architectures. In fact, this 13 dB figure represents a lower bound on the optical gain achievable with these OEICs. Use of an 85 μm DGW LED would increase this optical gain to roughly 20 dB.

“Number array”

Figure 3.30 simultaneously demonstrates both the impressive capabilities of the EoE technology and the unfortunate limitations imposed by the large variability in OPTOCHIP LED efficiency. The circuit

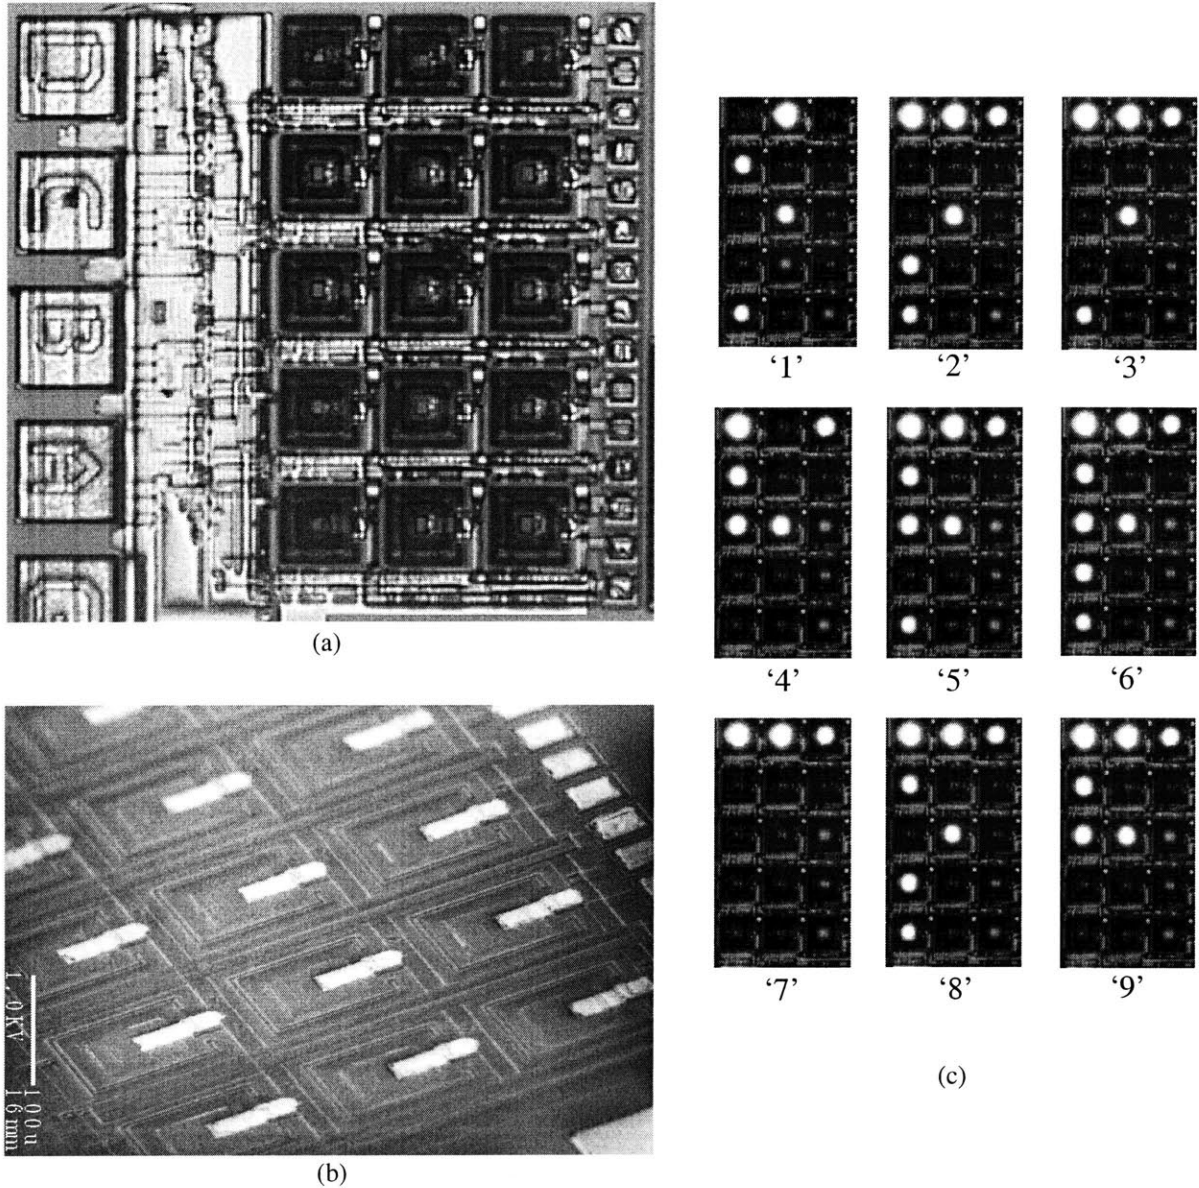


Figure 3.30: “Number array”

An EoE-integrated numerical display [147]: (a) A photomicrograph of the circuit. The decimal equivalent of a four-bit binary value, input through the four bondpads labeled ‘A’, ‘B’, ‘C’, and ‘D’, is displayed on a 3x5 array of standard OPTOCHIP LEDs. The combinational logic circuit is located to the left of the LED array, while LED drivers are placed immediately adjacent to the LEDs. The LED density in the array is nearly $6800/\text{cm}^2$. (b) An SEM view of a portion of the array. (c) The display output for the values one through nine as viewed by an IR sensitive CCD camera. Eight out of the fifteen LEDs are operating efficiently.

accepts a four-bit binary input and displays the corresponding decimal value on a 3x5 array of standard OPTOCHIP LEDs [147]. The combinational logic circuit which carries out the binary-to-display translation is to the left of the LED array in Figure 3.30(a), while driver circuits are placed immediately adjacent to each LED in the array itself. The LED density in the array is nearly 6800/cm². Figure 3.30(b) shows the display output for digits one through nine. The identity of the digits is obscured by the poor LED yield, although eight out of the fifteen LEDs are in fact emitting brightly.

“Smart pixel”

An additional demonstration of EoE’s functionality is given in Figure 3.31(a). This circuit, which will be examined in Chapter 5, is representative of a smart pixel designed as part of a 1 Tb/s/cm² array. It includes one 85 μm DGW LED, one OPFET-based receiver, and the complex digital electronics indicated in Figure 3.31(b). The intent of the circuit is to demonstrate dynamic optoelectronic circuit operation involving both optical input and output while slowing down the output in order to simplify characterization. As shown in Figure 3.31(b), an internally-loaded 63-stage ring oscillator generates a clock which is divided down by a chain of seven toggle-connected D-flip-flops. The output of one of the final two flip-flops is selected by a multiplexer on the basis of the optical input and is applied to the LED driver circuit detailed in Figure 3.31(c). The LED driver is, in fact, the OPTOCHIP drop-in cell shown in Figure 3.13. The electrical output of the driver and the optical output of the LED are shown in Figure 3.31(d) in the case of zero optical input. The response time of the LED output trace in Figure 3.31(d) is limited by the external receiver used in the measurement. The integrated LED oscillates at 476 kHz in the absence of optical input and at 943 kHz when an optical input is applied (These measured values differ from the expected 1:2 ratio by less than 1%; see Section 3.5). Internally, the DCFL electronics are operating at very high speeds. The gate delay extracted from the loaded ring oscillator in this circuit is 130 ps.

Summary

The integrated LEDs have fulfilled their reason for being. They have demonstrated the ability of the EoE technology to integrate efficient light emitters with high performance VLSI electronics. They have provided a vehicle to exercise EoE’s rich design capabilities. And, they have enabled the completion of the OPTOCHIP project which has helped to promote optical interconnect system development.

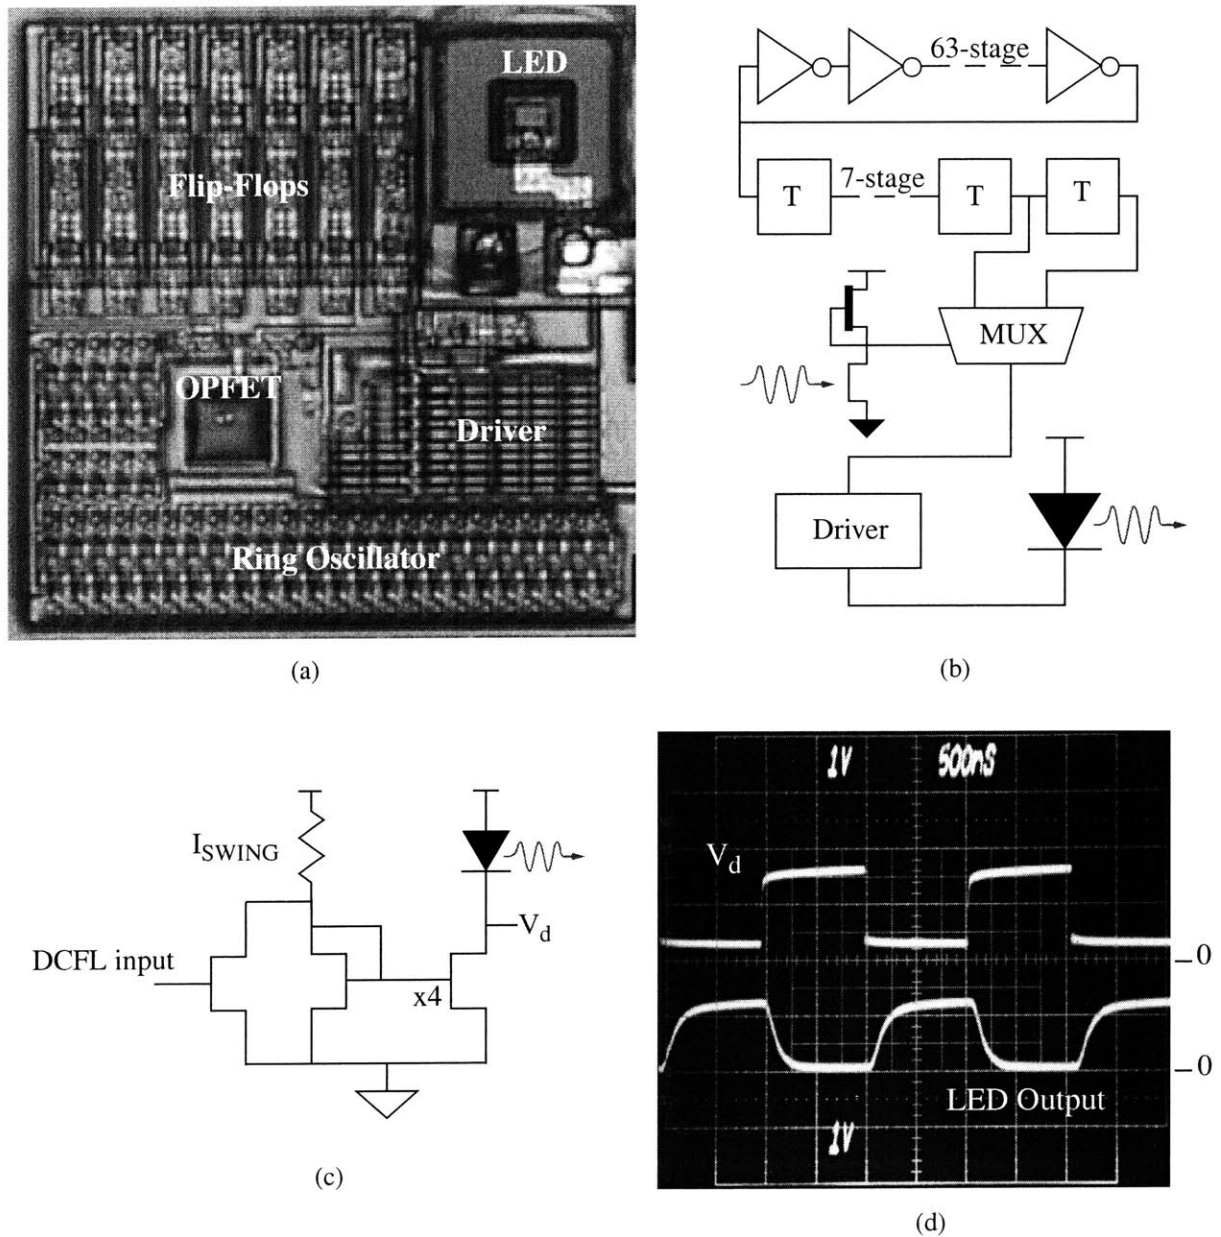


Figure 3.31: “Smart pixel”

A canonical tera-bit/s/cm² EoE smart pixel: (a) A photomicrograph of the completed 300 μm x 300 μm pixel. (b) A schematic description of the pixel’s logical functionality. The LED output oscillates at different frequencies depending on the optical input applied to the OPFET. (c) The LED driver circuit (same as Figure 3.14). (d) An oscilloscope trace photograph of the driver output voltage and the LED optical output in the absence of an optical input. The LED output trace rise/fall time is limited by the speed of the external detector used in the measurement.

3.5 Optical Crosstalk

In addition to standard electrical crosstalk issues, OEICs must deal with crosstalk involving optical signals. This includes effects which takes place entirely in the optical domain as well as optical-to-electrical crosstalk. This section is concerned with the latter topic. Purely optical crosstalk depends on the nature of the optical system being used and is not addressed in this thesis.

Surface shields: blocking stray external light

One possible source of unwanted optoelectronic interaction is due to stray external light. Naturally, the optical system being used should be designed to avoid this problem. Further protection is provided by ensuring that all non-optical elements on a circuit are covered by metallization which will block stray light. Since high speed circuits typically make use of entire metal sheets for power and ground distribution (Vitesse's metal-3 and metal-4 layers are intended for this very purpose), no addition cost is incurred in implementing this type of shielding.

Optical shields: blocking stray LED emission

Stray emission is also introduced directly on the OEIC by the LEDs. The OPTOCHIP LEDs emit light isotropically, and are thus expected to illuminate the circuitry in their vicinity. To block this emission, an optical shield is build around the DGWs by stacking metal layers up to metal-3. Metal-3 is high enough to block lateral emission from the LED core. For reasons to be explained shortly, the shield also includes a ring of n^+ source/drain implant and ohmic contact. An FIBE cross section of an optical shield is shown in Figure 3.32; this is a style #1 optical shield as used on the majority of the OPTOCHIP LEDs. This optical shield design violates H-GaAs III design rules and was implemented only after consultation with Vitesse. The large non-planarities and void in the overglass are undesirable consequences of this design. The style #2 optical shields attempt to reduce the non-planarities created by the shield by laterally displacing the vias. This shield design conforms to the design rules but consumes a greater area.

Both shield designs are expected, simply on the basis of their geometry, to be effective in blocking LED emission from coupling into nearby electronics. No direct evidence confirming or denying this is available, however, due to the existence of a much more significant optical crosstalk mechanism acting through the common, floating backgate present on H-GaAs III circuits.




Figure 3.32: FIBE cross section of an optical shield

SEM of the cross section of an optical shield, exposed by focused ion beam etching (FIBE) [137]. The shield forms a ring which surrounds an integrated LED. It consists of a source/drain implant and all metal and via layers up to metal-3, including the ohmic contact to the implant. The shield blocks stray emission from the LED from coupling into nearby electronics. It may also be electrically contacted in order to bias the implant, reducing optically induced crosstalk through the backgate.

✓ Backgate crosstalk

Vitesse's process begins with a semi-insulating GaAs wafer. The doping of these wafers is nominally n-type at 10^{16} cm^{-3} , but this low doping level is poorly controlled. In order to maintain device uniformity, and to implement isolation between adjacent MESFETs, a blanket p^- implant is applied to incoming wafers. This implant is roughly $0.8 \mu\text{m}$ deep and results in a p-type doping of around 10^{16} cm^{-3} . N-type implants are then used to build up the MESFETs and photodetectors. In the H-GaAs III process, the p^- layer constitutes a floating backgate which is common to all devices on the circuit. In Vitesse's high density digital circuits, the voltage on the backgate node is determined "democratically" by the MESFETs and is uniform enough for the types of circuits being implemented in H-GaAs III (most commonly gate arrays).

Consider, however, the situation shown in Figure 3.33 where a MESFET or OPFET (floating-gate MESFET) is located near another n-type region. The n-type region could belong to an OPFET, MSM, or




Figure 3.33: H-GaAs III cross section showing mechanism for backgate crosstalk

Vitesse H-GaAs III circuits are built on a semi-insulating GaAs substrate with a blanket p^- implant. The p^- layer constitutes a common, floating backgate beneath all of the devices in the circuit. Light incident on the n^+ implant belonging to an optoelectronic device (i.e. OPFET, MSM, or LED) effects the voltage on the backgate through the photovoltaic effect. This effect is felt by MESFETs and OPFETs elsewhere in the circuit.

LED. This implant forms an open-circuited p-n junction in conjunction with the backgate layer. When light strikes this region, the junction produces a photovoltaic self-bias; i.e. it's a solar cell. This changes the potential of the p^- region, and this bias change carries over to adjacent devices. The effect on nearby electronics is a threshold voltage shift similar to the electronic backgating effect. In digital logic circuits, this optically induced threshold shift may be acceptable in some cases. For instance, the smart pixel circuit above continued to function in the presence of light, with the ring oscillator frequency, or equivalently the gate delays, shifting by less than 1%. On the other hand, analog circuits tend to be more sensitive to backgate voltage variations. The effect on photodetectors, and especially the OPFETs, is very dramatic. Light incident on a backgate p^- -n junction, either from an external source or from an integrated LED (which is located directly above such a junction), can switch on OPFET photodetectors within several hundred microns of the light source.

In the case of the LED, it should be noted that since the emission is near the bandgap, its absorption length apparently exceeds the $4.5 \mu\text{m}$ thickness of the GaAs buffer layer below the LED active region. LEDs or lasers implemented at the more common 850 nm wavelength would not be expected to produce




Figure 3.34: Backgate isolation rings in H-GaAs III

A ring of source/drain n^+ implant may be used to isolate regions of the p^- layer in order to block the backgate crosstalk effect. A positive bias to the implant in order to increase the depth of the depletion region below it thereby pinching off the p^- layer.

this effect. The absorption length of light at this wavelength is $\sim 1 \mu\text{m}$ and would be expected to be absorbed by the GaAs buffer layer before reaching the p^- - n junction. $I = I_0 e^{-\alpha z}$

To mitigate optical backgate crosstalk problem on the OPTOCHIP OEICs, the photodetectors and LEDs were surrounded by n^+ source/drain implant rings as shown in Figure 3.34. The n^+ implant is intended to isolate the backgate region within the ring from the outlying p^- layer by depleting the material below the implant. The extent of depletion is increased by applying a positive bias to the n^+ implant. As indicated above, isolation rings of this type were incorporated into the optical shields surrounding the LED and were placed around the OPFET and MSM drop-in cells shown in Figure 3.15 and Figure 3.16. However, isolation rings of this type have been only partially successful in blocking the backgating crosstalk because of the limited depth of the source/drain implant. Significant backgate crosstalk effects were observed by some of the OPTOCHIP participants.

A robust solution to the backgate crosstalk has, fortunately, been provided by Vitesse in the latest revision of their process, H-GaAs IV. Motivated by the growing importance of mixed-signal electronics in their product line, H-GaAs IV provides both deep n^+ implants needed to form reliable isolation rings, and ohmic contacts to the p^- layer which allow the backgate voltage to be explicitly set. H-GaAs IV has been used in the circuits of Chapter 7 and Appendix G, and has been effective in eliminating optical backgating problems.

3.6 Lessons From OPTOCHIP

This chapter has reviewed the fabrication of integrated LED and collected a variety of device and circuit results from LED-based EoE OEICs fabricated as part of the OPTOCHIP project. Combined with the rich design capabilities outlined in Section 3.3, the above data identify the strengths of the EoE technology and point out areas of continued work. Foremost, the data show that with regard to electronic complexity and performance, optoelectronic density, and electronic/optoelectronic inter-operability, the EoE integration technology can fulfill the OEIC needs of advanced optical interconnects. The performance of the current LED-based OEICs is adequate to enable worthwhile exploration and demonstration of many optical interconnect approaches. With regard to the OPTOCHIP OEICs in particular, however, interconnect demonstrations have been hampered by the poor LED uniformity resulting from processing-related difficulties and the LED-to-detector and detector-to-detector backgate crosstalk. Neither of these problems is an inherent part of the EoE integration technology.

Circuits consisting of OPFETs, LEDs, and electronics are well suited to many optical interconnect demonstrations. It should be noted, however, that OPFETs provide high sensitivity at the cost of a low bandwidth of ~ 10 KHz. To achieve higher data rates, MSM photodetectors and transimpedance amplifiers, may be used instead of OPFETs. In this scenario, the interconnect performance will be limited by the LEDs. In particular, LEDs are ill-suited to most high-performance optical interconnect applications because of their low efficiency, directionality, and bandwidth. To take advantage of the high speed electronics available in EoE, the LEDs must be replaced with surface emitting lasers.

A further observation from the OPTOCHIP project is the need for sophisticated optoelectronic interface circuits. While the OPTOCHIP design groups exhibited great dexterity in digital circuit design, in spite of the non-CMOS-like DCFL logic which was involved, implementation of the mixed signal electronics needed to take full advantage of the optoelectronic components was not pursued. The task of providing an "end-to-end" solution thus falls to the technologist and is taken up in the latter portion of this thesis.

4

Vitesse MSM Photodetectors

An MSM photodetector is a planar device consisting of interdigitated fingers of Schottky metallization. By applying a bias across the fingers, the material between them can be fully depleted. Photogenerated carriers are then transported by drift, with the electrons collected at the higher potential electrode, and the holes at the lower potential. By fully depleting the material between the fingers, the capacitance of the device is reduced to that of the fringing fields alone [148-153]. The MSM photodetector's planar geometry makes it well suited for use in a monolithically integrated receiver [154-158], and its low parasitic capacitance results in high sensitivity [158,159].

Under license with IBM, Vitesse has incorporated an MSM photodetector into its technology [105,156-158]. The proprietary process includes additional masking steps needed to control the doping throughout the device, and is used in the fabrication of receiver arrays supplied to IBM as well as in individual receivers sold commercially by Vitesse¹. However, because of the additional cost associated with these devices, they are not available through the MOSIS service. This has prompted the investigation of alternate device structures implemented using only the standard VLSI process, or a minimally augmented version of it.

This chapter will begin by briefly describing MSM-like detector structures implemented in Vitesse's H-GaAs III process and used in the OPTOCHIP project. The majority of the chapter will then focus on a modified detector structures which includes one additional mask layer used at two points within the H-

1. The VSC7810, for example, combines a 100 μm detector with a transimpedance amplifier and has a typical bandwidth of 1.2 GHz.

GaAs IV process. These devices are found to have significant performance limitations. The physics underlying their operation is not fully understood, but a qualitative model will be put forth which is consistent with the experimental results. In addition, a “black-box” model of the detector will be constructed for use in SPICE simulations of the receiver examined in Chapter 7. Finally, an alternate detector structure will be proposed as a candidate for future investigation.

4.1 MSM Photodetectors in the Standard Vitesse H-GaAs III Process

A cross-sectional diagram of an MSM-like photodetector implemented using the standard H-GaAs III process is shown in Figure 4.1. Devices of this type were implemented in the OPTOCHIP OEICs. The p^- layer is a blanket implant made over the semi-insulating GaAs substrate at the beginning of the Vitesse process. The n layer at the surface is the EFET channel implant, and the electrodes are MESFET gates. As part of the self-aligned process flow, the regions in between the fingers receive the n^+ LDD and source/drain implants. The high doping does not allow the material between the fingers to be fully depleted. Rather, the Schottky diodes may be considered as being connected by a photoconductor. Nearly all of the voltage applied to the device is dropped across the reverse biased Schottky diode. As a result the capacitance between the fingers is essentially that of the reverse-biased electrode and is expected to be significantly




Figure 4.1: Cross section of MSM photodetector made with standard Vitesse H-GaAs III process

The Vitesse process begins with a blanket p^- implant over the semi-insulating GaAs Substrate. Using only the standard process steps used to build electronics, MSM photodetectors may be made from interdigitated fingers of gate metal placed over an EFET active region. The self-aligned n^+ source-drain implant then appears between the fingers. Furthermore, the p^- layer forms a floating backgate which is shared with adjacent devices.

larger than that of the fully depleted MSM described above.

A more pressing problem with this structure relates to the p^- layer. A p-n junction is formed between this layer and the n and n^+ implants at the surface. This junction plays a major role in the operation of the device, the discussion of which will be delayed to the next section. Aside from its role in the operation of the detector, however, this p-n junction takes part in the backgate crosstalk effect examined in Section 3.5. To briefly repeat the effect, light incident on the p-n junction generates a photovoltaic bias on the p^- layer. Since the p^- layer forms a floating backgate which is common to all proximate devices, this bias influences the operation of adjacent devices and circuits.

Static characteristics

Figure 4.2 shows the layout of an MSM detector test block from the MIT OPTOCHIP design (Figure 3.22). It includes four devices made with $1.0\ \mu\text{m}$ gate metal fingers over $40\ \mu\text{m}$ square EFET active regions. Finger spacings of 1.6 , 2.2 , 3.0 , and $9.0\ \mu\text{m}$ are implemented.




Figure 4.2: Layout of test block from OPTOCHIP featuring various MSM photodetectors

The device on the far right uses a $20\ \mu\text{m}$ square active region and $1.2\ \mu\text{m}$ gate metal fingers spaced $1.6\ \mu\text{m}$ apart. The remaining devices have $40\ \mu\text{m}$ active regions and $1.0\ \mu\text{m}$ fingers spaces at 1.6 , 2.2 , 3.0 , and $9.0\ \mu\text{m}$. The static characteristics of these four devices are shown in Figure 4.3 and Figure 4.4.




Figure 4.3: Dark current of MSM photodetectors made with standard Vitesse H-GaAs III process

The dark current was measured as a function of the voltage between the MSM-fingers. The devices with 1.0 μm gate metal fingers in Figure 4.2 were used.

The dark current of these devices is shown in Figure 4.3. The dark current is seen to rise exponentially with the voltage across the MSM electrodes. In addition, it is seen to increase as the finger spacing is made smaller. This is likely due to a combination of increasing the Schottky contact area and decreasing the resistance in series with the back-to-back Schottky diodes.

The static photoresponse of the detectors is shown in Figure 4.4. It is seen to be essentially linear, increasing slightly in efficiency at higher power levels. The efficiency is higher for devices with larger finger spacing.

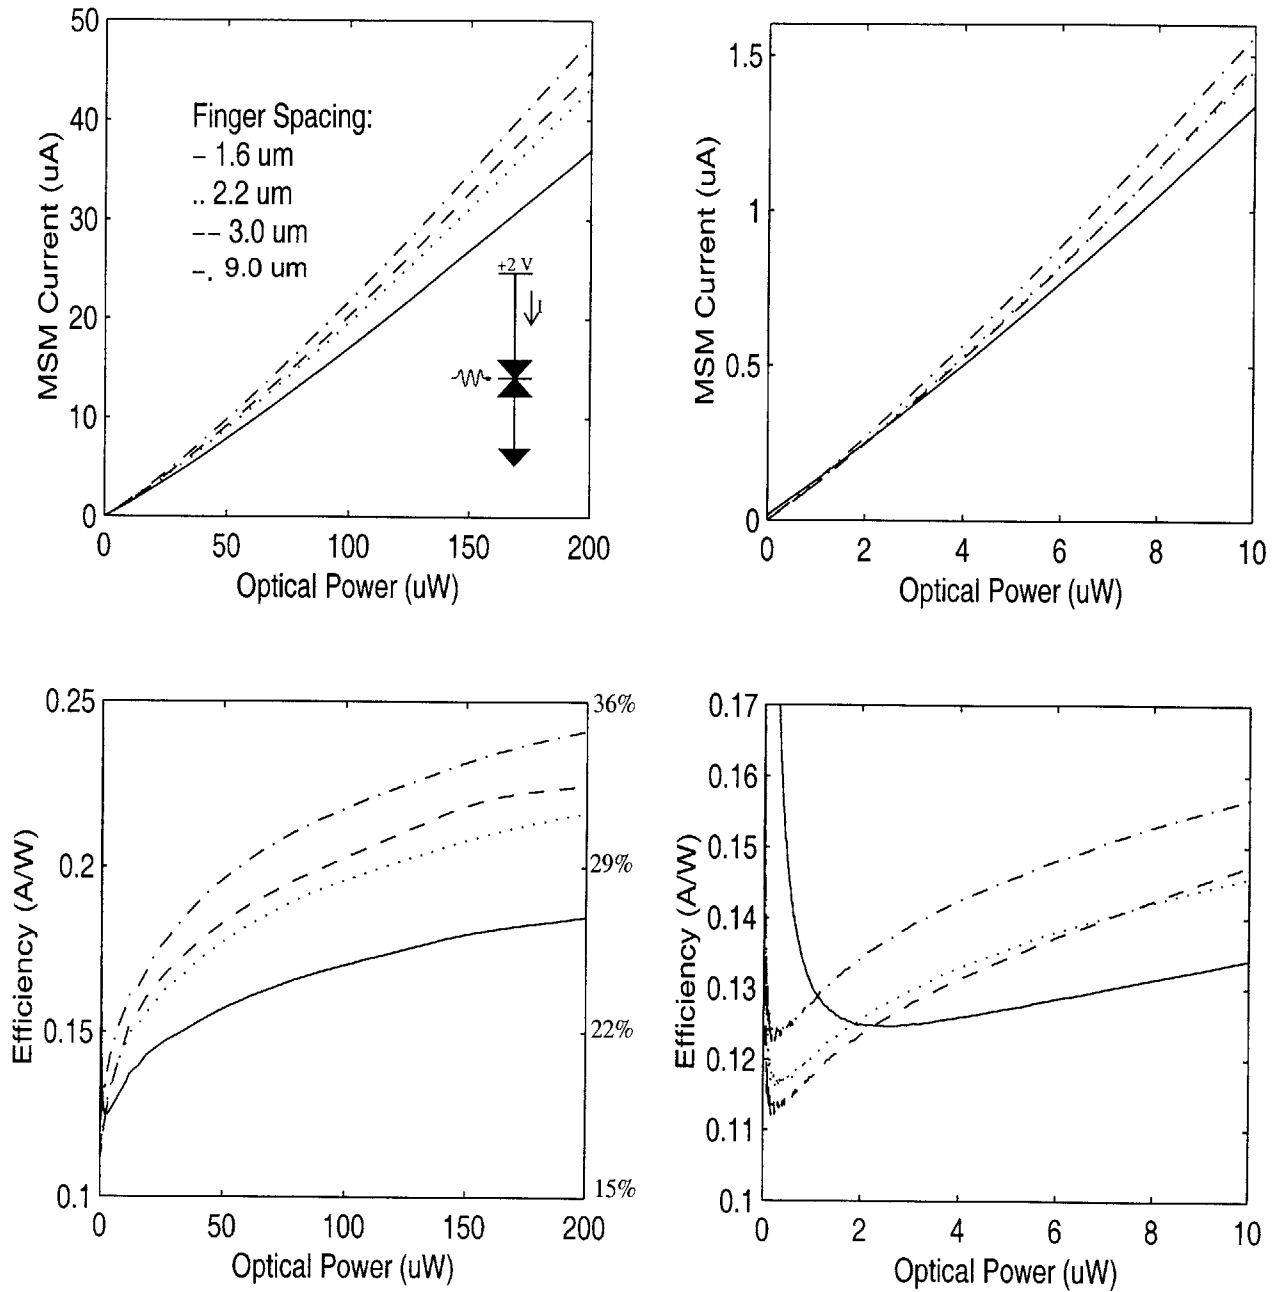


Figure 4.4: Static response of MSM photodetectors made with standard Vitesse H-GaAs III process

The $1.0\ \mu\text{m}$ gate metal finger devices of Figure 4.2 were measured at a bias voltage of 2 V. The measurements were made with 850 nm wavelength light. (a) and (b) show the photocurrent as a function of the optical power level, with (b) showing the low-power data in more detail. (c) and (d) give the conversion efficiency of the device, determined as the ratio of the photocurrent to the optical power. The corresponding quantum efficiencies are given on the right axis of (c).

Conclusive direct measurements of speed of these devices were not made. The device with $3.0\ \mu\text{m}$ finger spacing was included in the integrated receiver shown in Appendix F, and was operated at up to 200 Mb/s. However, the backgate crosstalk effects resulted in widely varying receiver decision levels depending on the bit rate of the finite-length pseudorandom data sequence. This was accompanied by up to 5 ns of timing variation in receiving isolated 1 and 0 data.

The biggest problem with these standard-process detectors was the common, floating backgate. It was also believed that the inability to fully deplete the devices would also ultimately limit the bandwidth of the detector. The next section describes an attempt at overcoming these limitations.


4.2 Modified MSM Photodetectors in the Vitesse H-GaAs IV Process

This section will describe an experimental detector structure which eliminates the heavy n-type doping between the MSM fingers and incorporates backgate isolation. The continued presence of the p^- layer below the detector will be found to severely impact the detector's performance. A conclusive physical model of the device has not been established, and the detailed numerical simulation which may be helpful in gaining such an understanding has been beyond the scope of this thesis. However, an explanation of the most overt aspects of the device's behavior will be presented. This basic description suggests that the observed performance limitations are an inherent result of the p^- layer, and that a lengthy investigation of the present device structure is not likely to overcome these limitations. Since these detectors have been incorporated into the optical receivers examined elsewhere in this work, an empirical, "black box" model will be demonstrated which can be used to account for their behavior in conjunction with the receiver electronics.

The device structure


The common, floating backgate has been eliminated by Vitesse in the H-GaAs IV process. Motivated by the growing importance of mixed-signal electronics in their product line, Vitesse introduced deep n^+ implants which could be used to form isolation rings. The new process also included ohmic contacts to the p^- layer so that the backgate potential could be fixed. In this way, various regions of the substrate could be isolated and maintained at different biases. These isolation structures have been found to be effective in eliminating optical crosstalk through the backgate.

The additional problem of high doping between the MSM fingers has also been addressed by Vitesse on an experimental basis. An additional mask layer has been introduced which would cover the EFET active region of the MSM, extending $1\ \mu\text{m}$ beyond it. This mask is used at two points in the Vitesse process to block the LDD and the source/drain implants. The resulting device is shown in Figure 4.5 along with the new isolation and contact structures. The layout shown in Figure 4.5(b) is of the device used in the receivers of the following chapters. It has a $75\ \mu\text{m}$ diameter active area and $0.5\ \mu\text{m}$ fingers spaced $2.1\ \mu\text{m}$ apart. This chapter will focus primarily on measurements made on this device, referred to as "MSM27".



(a)

Vitesse H-GaAs IV



(b)

Figure 4.5: Cross section and layout of modified MSM photodetector in the Vitesse H-GaAs IV process

(a) shows a cross sectional diagram and (b) the layout of a 75 μm diameter device with 0.5 μm gate metal fingers spaced 2.1 μm apart. The modified MSM photodetectors are made from interdigitated fingers of gate metal placed over an EFET active region. A non-standard mask layer is used to block the LDD and source/drain implants leaving only the EFET channel implant below and between the fingers. Using the new features of the H-GaAs IV process, the p^- layer, which begins as a blanket implant over the semi-insulating GaAs substrate, is contacted and isolated. Although the actual doping profiles are complex, they may be approximated in hand calculations as follows: the semi-insulating substrate has a nominal n-type doping of 10^{16} cm^{-3} , the p^- layer is 1 μm deep and doped at 10^{16} cm^{-3} , and the n-type channel is 80 nm deep and doped at $2 \times 10^{17} \text{ cm}^{-3}$. The Schottky barrier height may be taken to be 0.8 V and the barrier height due to surface pinning is around 0.6 V.

The MSM shown in Figure 4.5 is a three terminal device. In this text, the more positively biased MSM electrode will be referred to as the “drain”, and the lower biased MSM electrode will be called the “source”. The p^- layer is referred to as the “backgate”. The subscripts “S”, “D”, and “B” will be used to denote these terminals. The biasing configuration used in the receiver circuits in this work is shown in Figure 4.6. A positive supply is applied to the drain of the MSM while the source voltage is set by the input bias point of the front-end transimpedance amplifier. Since the input bias point tends to be nearer to the receiver ground, this configuration is chosen to ultimately allow operation of both the detector and receiver from a single positive supply. The p^- layer is connected to ground, the implications of which will be seen shortly. The configuration in which the MSM drain is connected to the transimpedance amplifier will also be considered later in this chapter.

Depletion

The doping concentrations and depths given in Figure 4.5 are a crude approximation to the complex, implanted doping profiles of the actual device. Using these approximations, basic electrostatic calculations indicate that the n-type channel is fully depleted in the unbiased device. This is as expected beneath the Schottky contacts since the n-implant is that of an enhancement-mode MESFET. In the regions between the fingers, the n-layer is depleted from above by surface pinning and from below by the p^- -n junction. The depletion region of the p^- -n “backgate” junction extends around 20 nm into the n-layer and 0.4 μm into the p^- layer at zero bias. Since the semi-insulating substrate is typically n-type, the p^- layer is also depleted from below by around 0.3 μm . Thus, there is around 0.2 μm of undepleted p^- material beneath the device.




Figure 4.6: Common biasing configuration of MSM photodetectors in an integrated receiver

The source voltage of the MSM is determined by the input bias point of the transimpedance amplifier while the drain is raised to a positive supply. The backgate connection, not shown, is set to 0 V.




Figure 4.7: Capacitance vs. voltage of modified MSM photodetectors

Measurements were made on $400\ \mu\text{m} \times 400\ \mu\text{m}$ devices with $0.5\ \mu\text{m}$ finger spaces as indicated. The two sets of fingers on each device were connected together, and the capacitance was measured at a frequency of 1 kHz as a function of the voltage applied between the fingers and the p^- layer. A 40-pin DIP package was used to hold the IC in these measurements, and the capacitance of the package has not been subtracted. The drop in capacitance between 0 and 2.5 V applied bias is ascribed to the change in depletion of the p^- layer. Normalized to the total length of gate metal used in the device, the capacitance change takes on a value of $0.33\ \text{fF}/\mu\text{m}$ for all three finger spacings.

A reverse bias of 2 V on the backgate junction would be adequate to deplete the remainder of the p^- layer. This corresponds to an applied bias of just over 2 V between the gate and the ohmic contact to the p^- layer. This calculation is consistent with the measurements shown in Figure 4.7 of the capacitance between the fingers and the p^- layer as a function of the voltage applied across them. The capacitance is seen to drop as the voltage is increased, and reaches a constant at around 2.5 V.

Dark current

With no light incident on the device, the currents flowing through each terminal of MSM27 are shown in Figure 4.8 as a function of the drain-to-source and backgate-to-source voltages (V_{DS} and V_{BS}). Provided the backgate diode is not forward biased, the dark currents are found to be around 0.1-1 nA. The backgate diode becomes forward biased with respect to the source at around $V_{BS}=0.5$ V. At this point, substantial current originating at the backgate begins to flow out of the source.

In addition, the drain current is seen to rise exponentially with increasing V_{BS} . To explain this, it is helpful to think of the p^- backgate and the n implant as forming an enhancement-mode n -channel JFET. As a result of the extent of the surface depletion, the JFET pinch-off voltage is actually greater than the junction turn-on voltage (it is not a very good JFET). The drain current seen in Figure 4.8 thus represents the sub-threshold JFET channel current, which depends exponentially on V_{BS} .

Most photogenerated holes are collected by the backgate

As seen in Figure 4.9, the absorption coefficient of GaAs at a wavelength of 850 nm is around 1.2×10^4 cm^{-1} . That is, the intensity of 850 nm light propagating in GaAs decays exponentially with a characteristic length of 0.8 μm . Around 10% of this light will be absorbed within the n -layer at the surface of the device, around 60% will be absorbed within the p^- region, and the remaining 30% of the light will be absorbed by the substrate and lost to the detector. As indicated by the qualitative band profiles in Figure 4.10, all photogenerated electrons drift towards the surface of the device, while most of the photogenerated holes are collected by the backgate. The holes generated within the n -layer are within a diffusion length of both the p^- - n junction edge and of the source and drain contacts. The band profile will tend to drive them to the surface, however, so the majority of them are likely to be collected by the surface contacts. In a conventional MSM photodetector, holes drift to the lower-biased electrode through which they exit the device. In the MSM-like structure being examined, the majority of the photogenerated holes go to the p^- layer. Consequently, the terminal current of this device depends strongly on what happens at the backgate. The resulting, undesirable static and dynamic behavior of MSM27 will be detailed in this chapter.

Static photocurrent with the backgate open-circuited

Measurements were made on MSM27 with its backgate connection open-circuited. In this case, the backgate junction may be thought of as a photovoltaic cell. The junction self-forward-biases so as to cancel

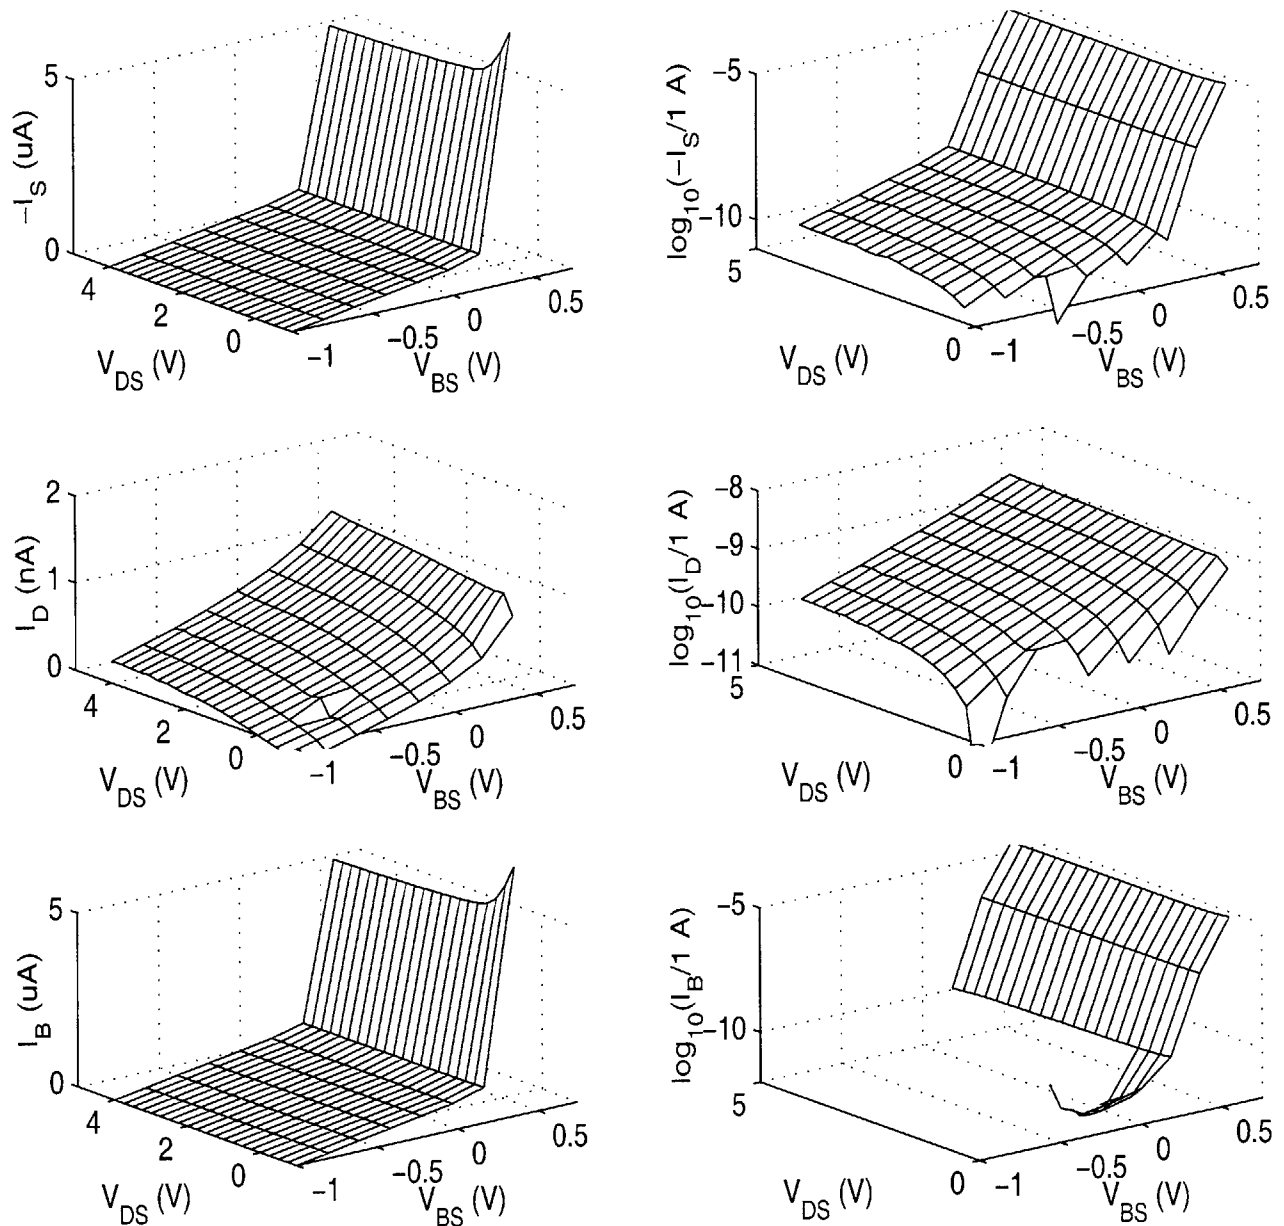


Figure 4.8: Dark current of MSM27

$-I_S$ is the current out of the source electrode of the MSM (lower biased electrode), while I_D and I_B are the currents flowing into the drain and backgate, respectively. V_{DS} is the drain-to-source voltage and V_{BS} is the backgate-to-source voltage. The left-hand plots are given on a linear scale, while the right-hand plots show the log of the data. The dark current into the drain varies between 0.1-1 nA and is shared between the source and the backgate. Provided the backgate diode is not forward biased with respect to the source, the dark current flowing out of the source remains at around this level. At around $V_{BS}=0.5$ V, the backgate diode becomes forward biased and flows a substantial amount of current into the source. Note, the plotted backgate current was calculated as $I_B=-(I_S+I_D)$.




Figure 4.9: Optical absorption coefficients of various semiconductors

The absorption coefficient of GaAs at a wavelength of 850 nm is around $1.2 \times 10^4 \text{ cm}^{-1}$. At this wavelength, the intensity of light propagating in GaAs would decay exponentially with a characteristic length of $0.83 \mu\text{m}$ (the reciprocal of the absorption coefficient). The above plot was taken from [160].

the reverse-flowing photocurrent. The resulting backgate voltage is shown in Figure 4.11, and $I_D = -I_S$ is plotted in Figure 4.12 as a function of V_{DS} and optical power level. In this experiment, V_S was set to 0 V and V_D was swept from -1 V to 4 V. Light from an 850 nm fiber-pigtailed in-plane laser diode was collimated by a graded index lens and focused onto the $75 \mu\text{m}$ detector using a microscope objective. The terminal current rises linearly with V_{DS} up to a point and then saturates. Reflection at the SiO_2/GaAs interface allows 84% of the incident photons to enter the GaAs¹. The gate metal blocks all but 81% of the incident photons. As discussed above, 70% of the photons that enter the GaAs are available to the device.

1. Assuming $n(\text{SiO}_2)=1.5$ and $n(\text{GaAs})=3.5$




Figure 4.10: Vertical band profiles under various bias conditions

(a) shows the vertical structure of the device for reference. (b) is a qualitative, unbiased band profile under either a Schottky contact or the GaAs surface. (c) and (d) are profiles under the Schottky contacts when the p^- - n junction is reverse- and forward-biased, respectively. As indicated by the arrows, photogenerated electrons are always swept to the surface by the built-in field of the junction, while holes always drift towards the substrate.




Figure 4.11: Open-circuited backgate voltage of MSM27

V_S was set to 0 V while V_D was swept. The p^+ contact, n^+ isolation rings, the substrate were shorted together and the voltage on this node, V_B , was measured both in the dark and under 850 nm illumination at the power levels indicated.

The compounded theoretical efficiency is 48%. The saturated MSM27 efficiency shown in Figure 4.13 is consistent with this calculation. The initial V_{DS} dependence of the photocurrent is explained next.

As the backgate junction becomes forward biased, the extent of the depletion region in the n-layer is reduced. A similar effect reduced the depletion width at the surface of the device as well. As a result, a conductive channel is formed through the un-depleted n-type material which connects the source and drain. Current flow through this channel increases linearly with the voltage drop across it¹.

1. The channel current will increase up to the point at which the drain-end of the channel becomes pinched off. This is the normal JFET drain-current saturation mechanism. This is not believed to be the saturation mechanism in MSM27.

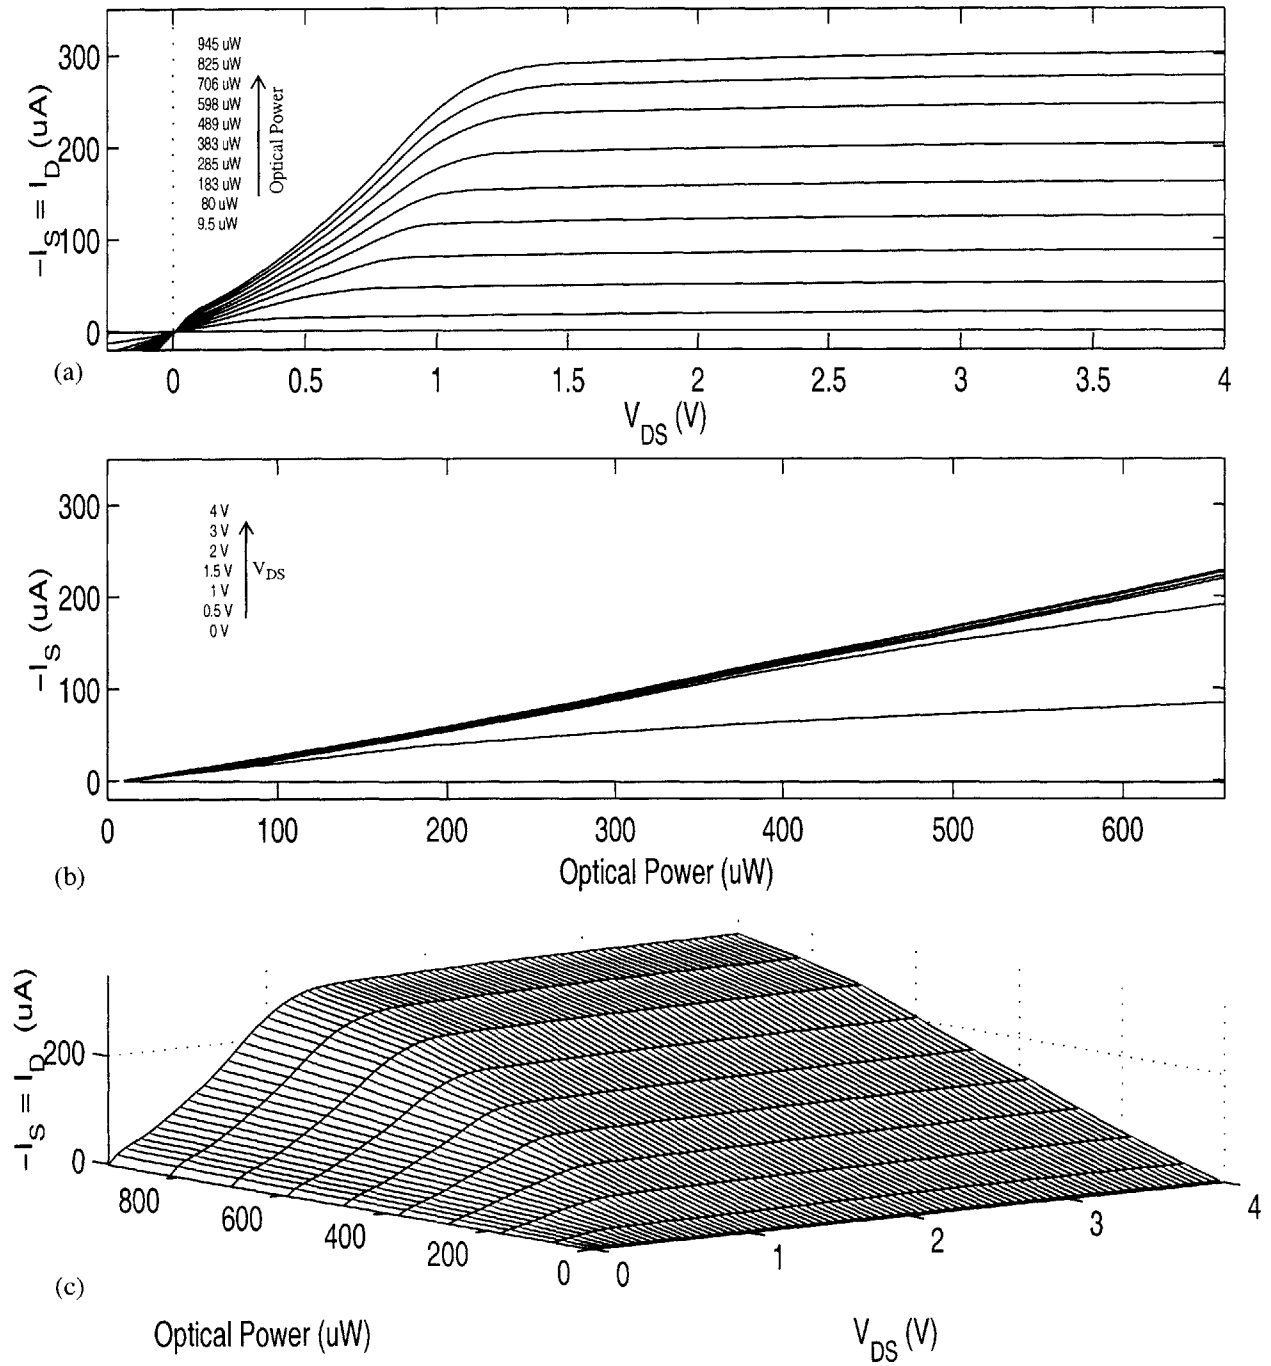


Figure 4.12: Static MSM27 photocurrent with backgate open-circuited. V_S was set to 0 V, V_D was swept, the 850 nm light was used to excite the device.

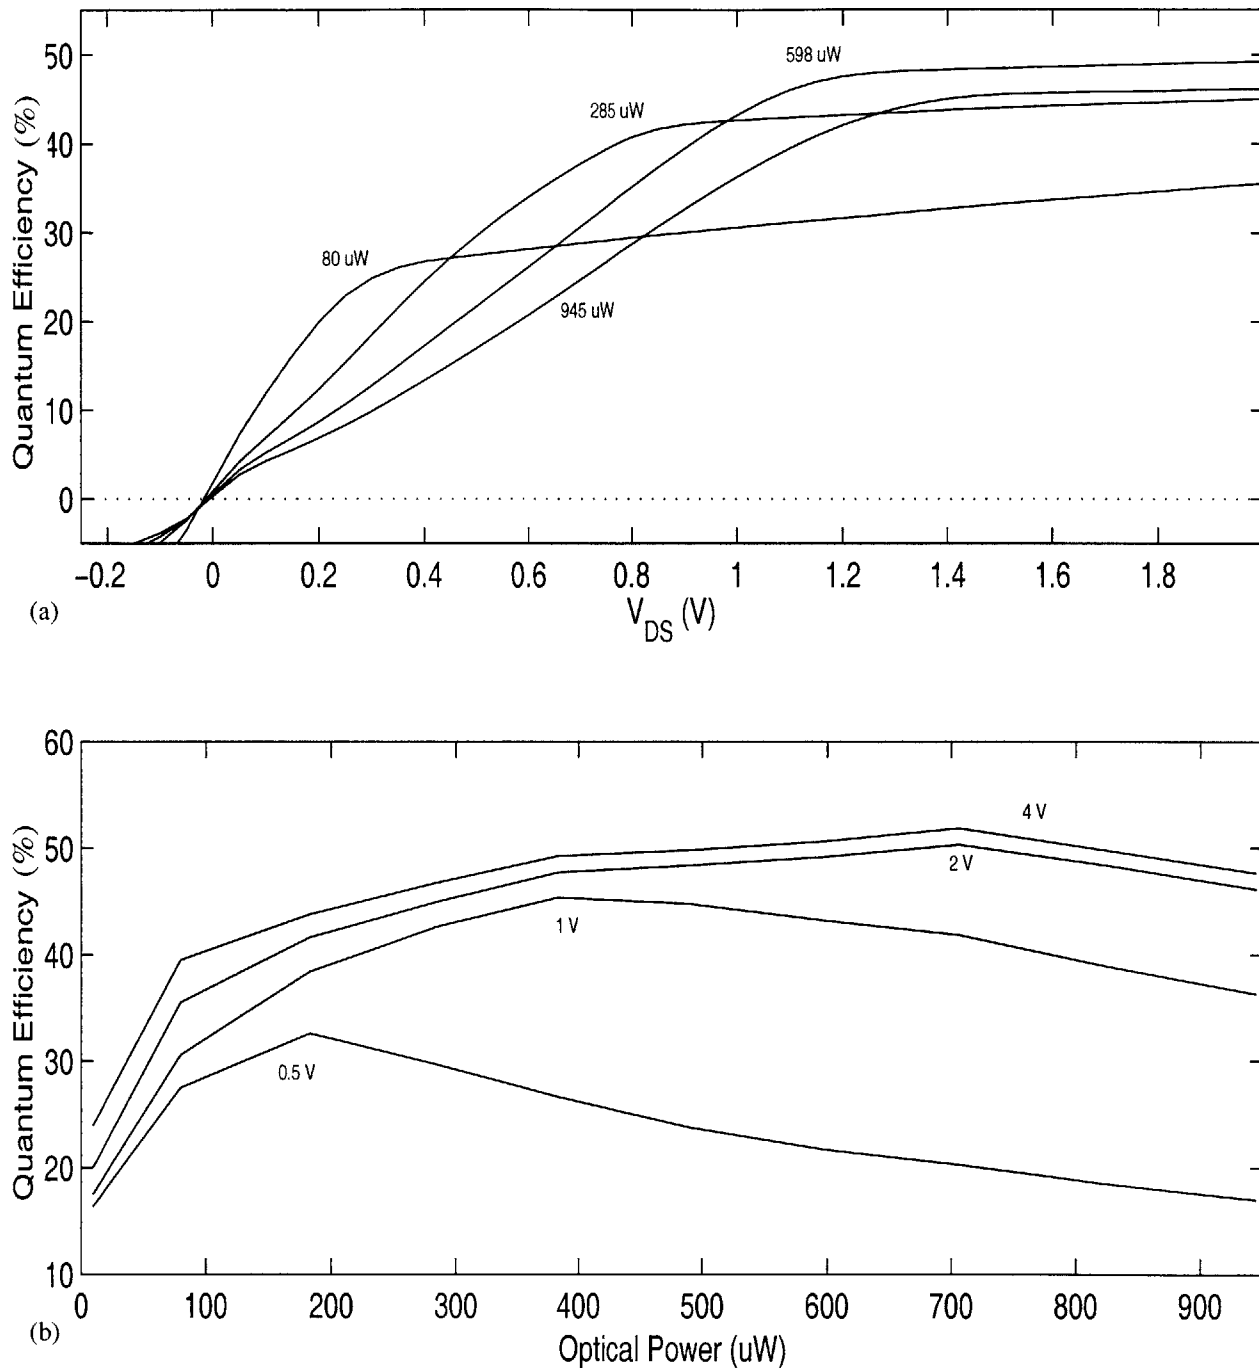


Figure 4.13: MSM27 quantum efficiency with backgate open

The quantum efficiency was calculated as $(I/q)/(P/h\nu)$ from the data of Figure 4.12 assuming 850 nm light.

When $V_{DS}=0$, both the source and drain Schottky diodes are forward biased and the potential at each end of the channel is roughly one Schottky diode drop below the voltage applied to the corresponding electrode. This is explained in Figure 4.14. Holes which are back-injected from the p^- layer, and those which diffuse to the contact after being photogenerated near the surface, are collected by the Schottky contacts. Since, with $V_{DS}=0$, there is no terminal current, the Schottky diodes become forward biased to allow a flow of electrons to counter the holes that are being collected.

As the drain voltage is increased, a voltage drop appears across the channel. The channel current increases roughly linearly with the applied voltage. The channel current diverts electrons from the source contact. As it does, the voltage drop across the source Schottky diode decreases. Once the terminal current has risen to equal the photocurrent, no additional electron current is available at the source. As V_{DS} increases further, the source Schottky diode becomes reverse biased and the potential at the source end of the channel rises with the drain. The terminal current is thus saturated at a value equal to the photocurrent.

Static photocurrent with the backgate grounded

With the backgate contact now grounded, the terminal currents of MSM27 are shown in Figure 4.15 as a function of the drain-to-source voltage and the incident optical power. The drain voltage was varied while the backgate contact was held at 0 V and the source at 1 V. The data of Figure 4.15 is recast in Figure 4.16 as families of curves to show the dependence on V_{DS} and optical power more distinctly. The corresponding quantum efficiencies are plotted in Figure 4.17.

Unlike the open-backgate case, there is now significant current flow through the backgate. With $V_{DS}=0$, this current is shared equally between the source and drain. At relatively low light levels, the backgate collects all of the holes that are generated within the p^- -region. The resulting quantum efficiency is around 40% (based on the same calculation as carried out above for the open-backgate case, but with only the 60% of the photogenerated carriers in the p-layer collected by the backgate). In Figure 4.17(b), the backgate quantum efficiency is seen to rise to this level once the optical power reaches around 100 μW . It remains at this level until around 300 μW beyond which it begins to drop. The initial rise in efficiency may be due to a recombination process which is saturated at the higher levels.




Figure 4.14: Biasing mechanism for the source and drain Schottky diodes

The band structure below the source or drain contact is shown in (a) with no applied bias or incident light. When light is incident on the device, the backgate junction becomes forward biased, as in (b). Photogenerated electrons drift towards the surface while photogenerated holes drift towards to the backgate. The forward-biased backgate diode produces a back-injected hole current to cancel the photocurrent. The back-injected holes and holes which have been photogenerated near the surface and have diffused to the Schottky contact are swept into the contact. To prevent the unbounded build-up of electrons below the contact, the Schottky diode becomes forward biased, allowing electron current to flow. When $V_{DS}=0$, the hole and electron flows into the contact are equal and the net current is zero. When $V_{DS}>0$, the two flows differ by the terminal current. However, the Schottky diode remains forward biased at both the source and drain. Once the terminal current has risen to equal the photocurrent, the electron current at the source will have dropped to zero and the voltage drop across the source Schottky diode is zero. The terminal current is saturated at this point.

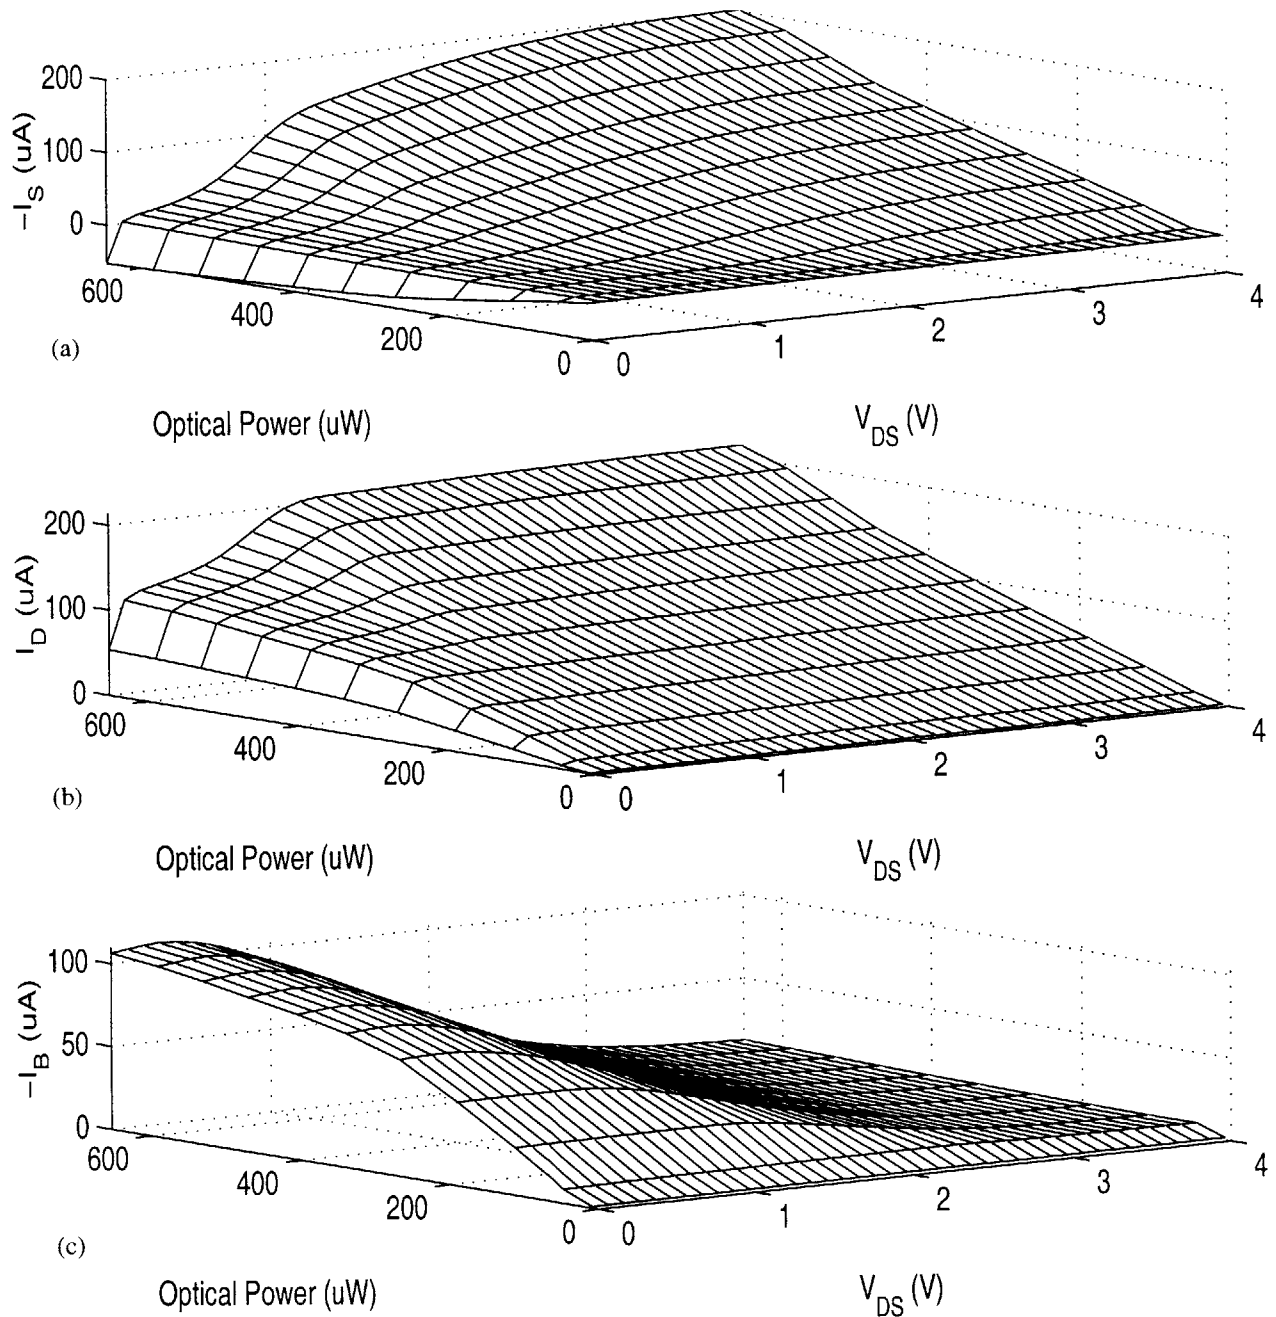


Figure 4.15: Overview of static MSM27 photocurrent with backgate grounded

The backgate connection was set to 0 V, V_S was held at 1 V, and V_D was swept. An 850 nm light source was used to excite the device. The drain current, I_D , and the source current, I_S were measured, and the backgate current, I_B was calculated at $I_B = -(I_D + I_S)$.

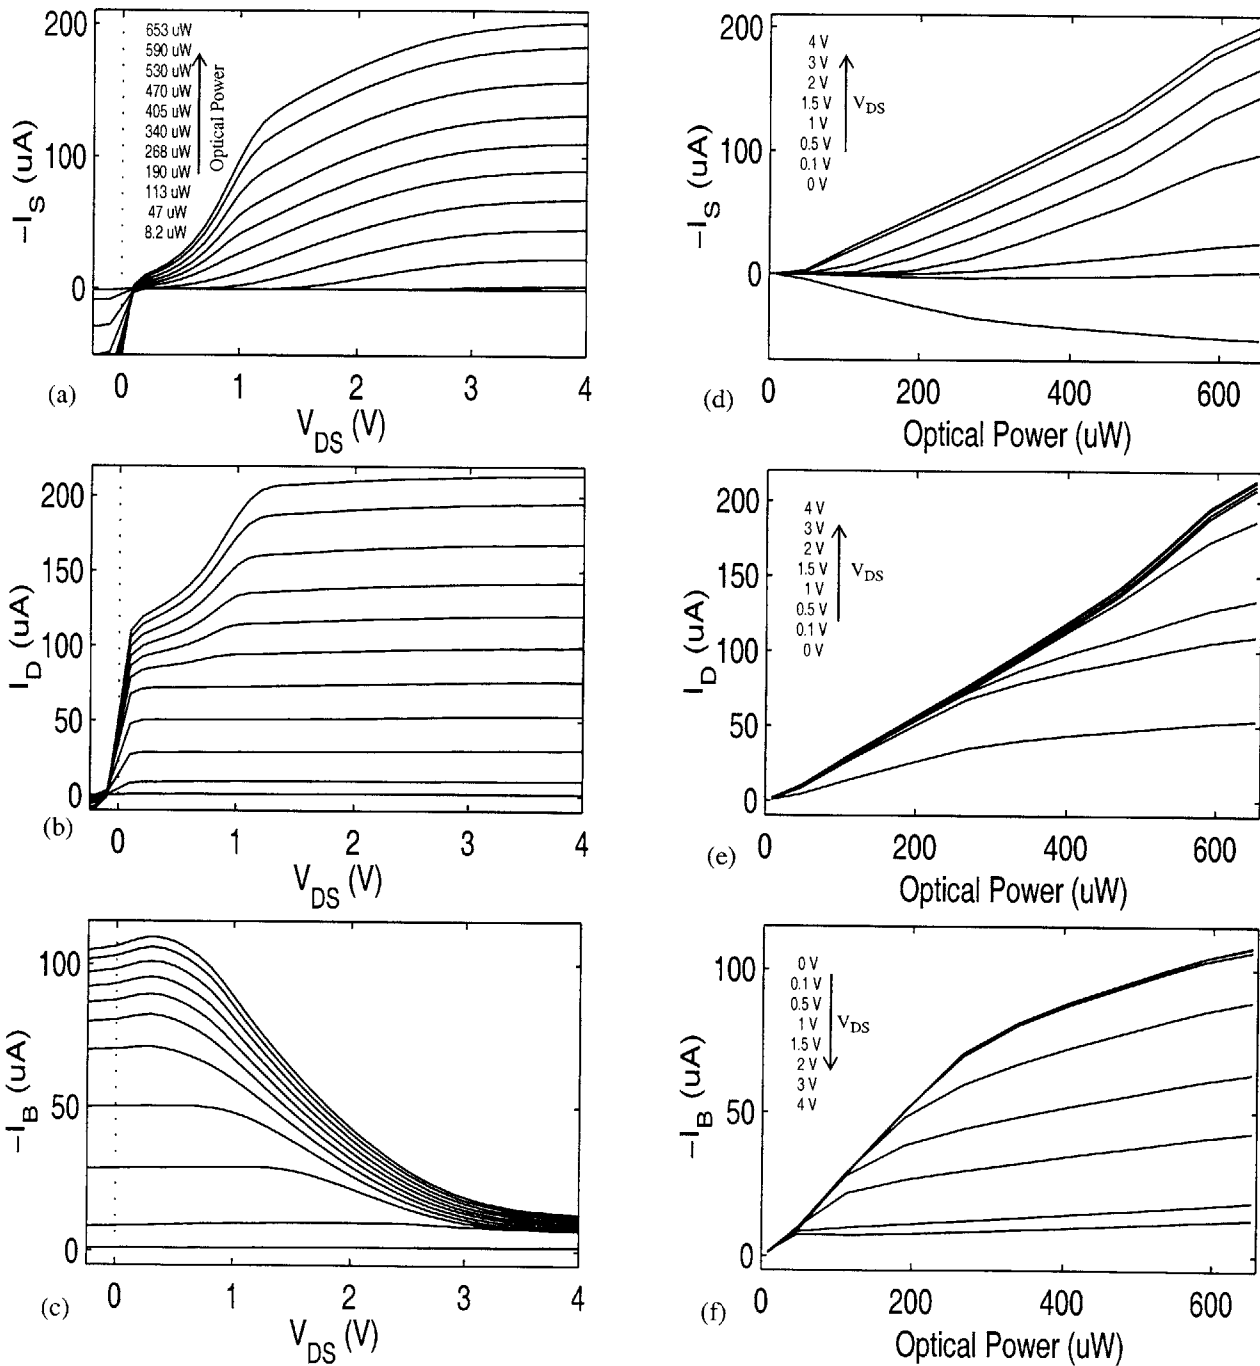


Figure 4.16: Static MSM27 photocurrent with backgate grounded

The data from Figure 4.15 has been recast here to show the dependence of V_{DS} and optical power more distinctly.




Figure 4.17: MSM27 quantum efficiency with backgate grounded

The quantum efficiency was calculated as $(I/q)/(P/h\nu)$ from the data of Figure 4.15 assuming 850 nm light.

The drop at higher power levels is a result of the finite resistance of the backgate connection. The ground connection to the backgate node is made through the resistance of the partially depleted p⁻-layer. The undepleted p⁻-layer has a sheet resistance of ~30 k Ω /square, so that a backgate resistance of several tens of k Ω 's is likely in the partially depleted case. As the current through the backgate increases in response to an increased optical power level, the voltage at the p-side of the junction increases. By the same reasoning as above (Figure 4.14), the source and drain are forward biased, setting the n-side of the junction to around 0.5 V (with $V_S = V_D = 1$ V). The backgate junction thus becomes forward biased when the backgate current is 50-100 μ A, and this is seen in Figure 4.16(f) to occur at a power level of around 300 μ W. The forward current of the diode now begins to counter the photocurrent, thus reducing the rate of increase of backgate current with increasing light level. That is, the backgate quantum efficiency drops with increasing light level.

As V_{DS} is increased slightly above zero, all of the photogenerated electrons which were going to the source are shifted to the drain. The drain current quickly doubles, then remains flat as V_{DS} is increased further. This behavior is observed in Figure 4.16(b) for power levels below 300 μ W. For higher powers, the drain current rises from the value it has when V_{DS} is just above zero. It rises gradually at first then becomes linear with V_{DS} , strongly resembling the open-backgate case studied above.

The key observation to be made here is that as the backgate diode begins to be forward biased, it has the same effect as the open-circuited backgate discussed above. The drain current thus goes from simply collecting photogenerated electrons, to a combination of photogenerated electrons and channel conduction. This process again saturates once the drain current is equal to the total photocurrent.

The bias configuration in which the MSMs are used (Figure 4.6) relies on the source current as the input to the transimpedance amplifier. The source current is determined by two mechanisms when the backgate is grounded. A threshold behavior is observed whereby, below a certain optical power level, all current flow is between the drain and the backgate (see the threshold in I_S in Figure 4.16(d)). But, as the backgate current increases, the backgate voltage increases, and the backgate diode becomes forward biased. This injects holes from the p⁻ layer towards the surface of the device, and these holes are collected by the source. This is independent of the second mechanism, the flow of current directly between the source and drain via the conducting channel.

The optical threshold level which controls both the forward injection of holes to the source and the creation of the conducting channel is determined by the backgate resistance. The backgate resistance is increased by raising the drain voltage to widen the p^- depletion width. This explains the bias dependence of the threshold level seen in Figure 4.16(d). Another way to see this dependence is to look at devices of differing areas. Figure 4.18 compares three structures with 2.1 μm finger spacing but with areas of 400 $\mu\text{m} \times 400 \mu\text{m}$, 40 $\mu\text{m} \times 40 \mu\text{m}$, and the 75 μm diameter MSM27. Compared to MSM27, the larger device has a much more distinguished transition to open-backgate-like behavior. The smaller device, on the other, exhibits no such behavior at all. The source current is produced entirely by direct injection from the backgate (characterized by the gradual, sub-linear V_{DS} dependence.)

The bottom line

The theory presented above is not complete and in all likelihood contains some inaccuracies. The general concepts, however, are very likely to be valid. Most importantly, the flow of source current requires the backgate diode to be driven into forward bias. This results in strongly bias-dependent, non-linear static characteristics, and has grave implications for the dynamic characteristics. Namely, there is a large time constant associated with charging and discharging the backgate node, which has a large capacitance to the substrate, isolation rings, and drain, and a large resistance to its ground connection. Reducing the device size reduces these parasitics at the expense of increasing the optical threshold level. (Because of the nonlinearities involved in switching process, its not clear if this will actually speed up the device.) The increased threshold means that the device may not be used in low-optical-power applications. There are additional dynamics associated with the transit time across the conductive channel.

What about using the drain current?

If the detector is configured with its drain attached to the transimpedance amplifier, and provided the optical power level is low enough not to forward bias the backgate, then the pathologies discussed above are not encountered. Equivalently, the source could be left floating, or connected to the drain (i.e. just have one contact). Basically, the structure being used in this case is a reverse biased p-n junction, with the Schottky contact made to the n-side of the detector instead of an ohmic (this is acceptable since it only has to work in forward bias). The difficulty with this lies in the capacitance. From Figure 4.7, the capacitance of this device is expected to be around 400 - 700 fF for a 75 μm device. A capacitance of this scale is a sig-

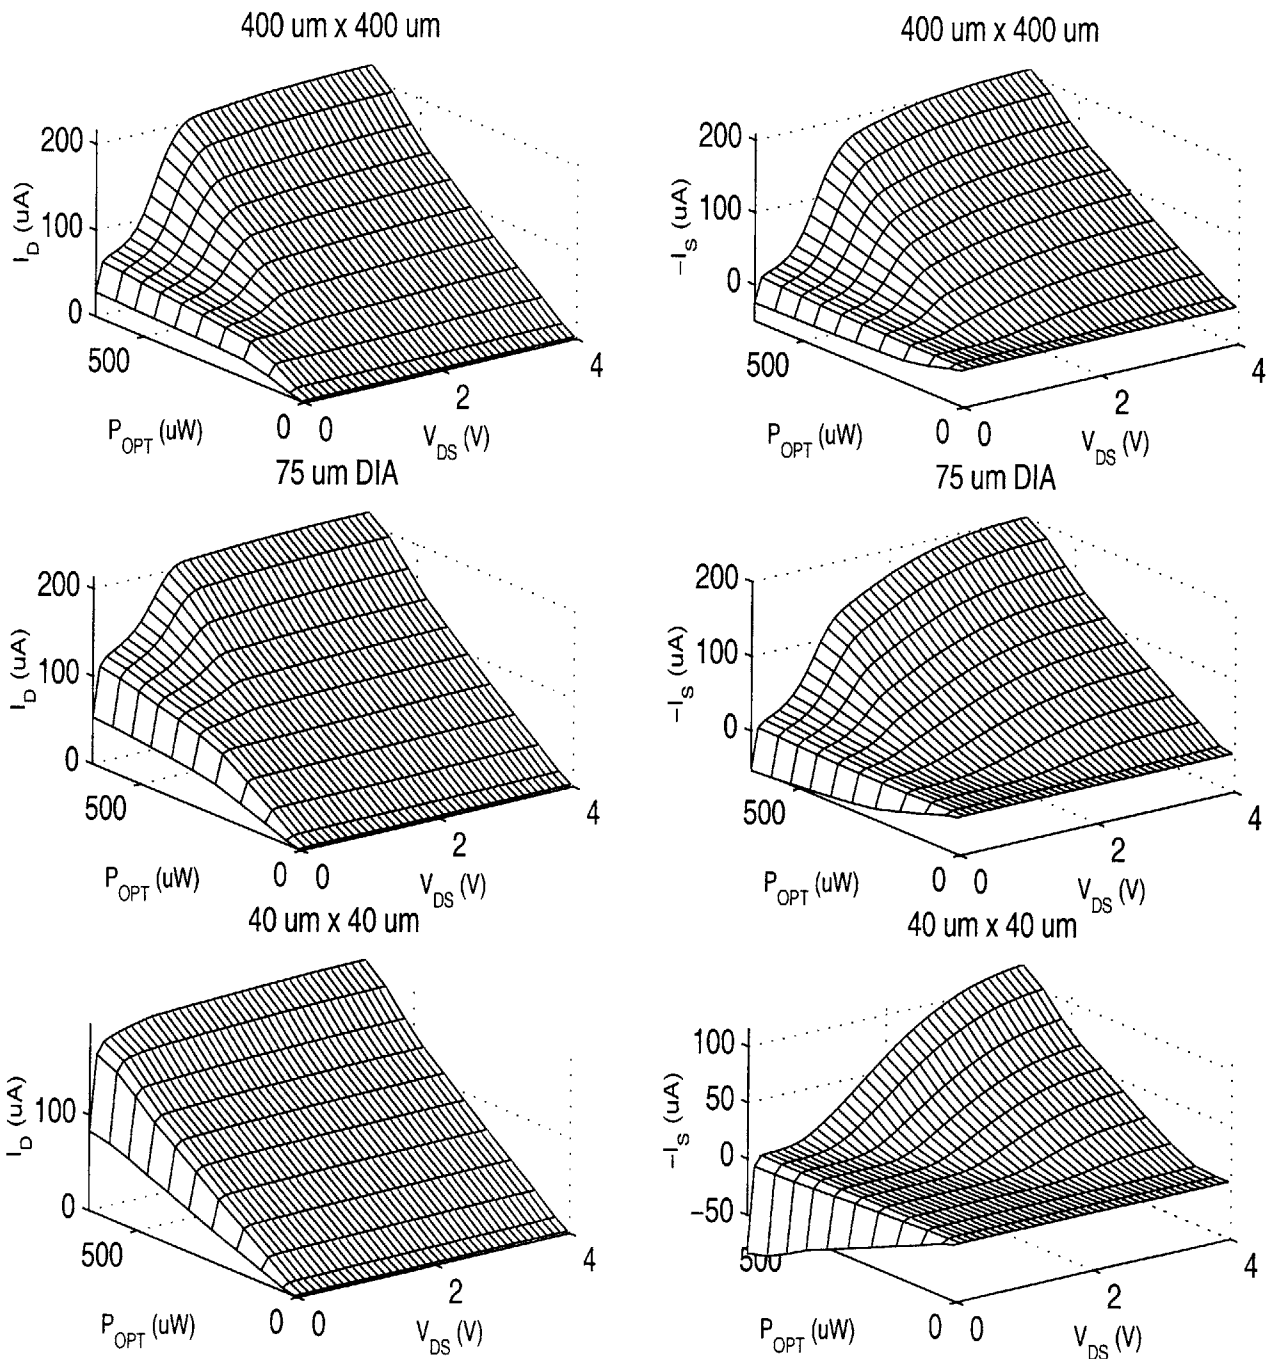


Figure 4.18: Comparison of devices with varying active areas

I_D and I_S of three devices have $0.5 \mu\text{m}$ gates spaced at $2.1 \mu\text{m}$ were measured with their backgate connection grounded, $V_S=1 \text{ V}$, and V_D and the 850 nm optical power level swept. The three devices included a $400 \mu\text{m} \times 400 \mu\text{m}$ square, the $75 \mu\text{m}$ diameter MSM27 device, and a $40 \mu\text{m} \times 40 \mu\text{m}$ device.

nificant performance constraint in the receiver design (though its certainly better that using the source current). If the device is scaled down to $40\ \mu\text{m} \times 40\ \mu\text{m}$, which is likely to be useful free-space optical interconnects, the capacitance drops to 150 - 300 fF. The use of this detector configuration will be revisited in Section 4.3.

Empirical modeling approach

MSM27 has been incorporated into optical receivers discussed elsewhere in this thesis, and will be found to limit their performance. A simple black-box model of the detector is valuable in explaining these results. To this end, the source current may be considered to be produced by an equivalent circuit consisting of a current source in parallel with an impedance. This is similar to the standard practice of representing a photodetector as a current source and capacitor. For example, in a p-i-n photodetector, the current source may be considered a filtered copy of the optical signal, where the impulse response width corresponds to the transit time of carriers across the i-region, and the capacitance is the depletion capacitance. Likewise, in a properly functioning MSM detector, the current source impulse response arises from the drift of carriers between the Schottky fingers and the capacitance is due to fringing fields between the fingers.

Because of the idiosyncrasies of the detector structure being considered here, the model used to describe it must be somewhat more general. The source current depends on the optical input signal and on the terminal voltages (V_D , V_S , and V_B). It may not necessarily evolve in a linear, time-invariant manner. Since the backgate does not directly appear in the receiver circuits, V_B will not explicitly appear in the black-box model; the backgate dynamics will be implicit in the remainder of the model. Furthermore, V_D will be a low impedance supply and will be assumed to be fixed, while $V_S(t)$ will be treated as the sum of its mean value, V_S , and a small signal variation, $v_s(t)$. Similarly, the optical input signal is $P(t) = P + p(t)$. As diagrammed in Figure 4.19(a), $I_S(t)$ can be thought of as a function of the optical power signal, $P(t)$, and the source voltage, $V_S(t)$. From the static characteristics above, it is seen that the dependence of I_S on V_S is relatively weak provided V_S is a few tenths of a volt above V_B and V_{DS} is greater than ~ 2.5 V. Thus, $I_S(t)$ can be broken up into two components, $I_{S,opt}(t)$ and $I_{S,elect}(t)$, as shown in Figure 4.19(b). $I_{S,opt}(t)$ is due explicitly to the optical signal, and is a function of $P(t)$ that is parameterized on the bias points V_S , V_D , and V_B . $I_{S,elect}(t)$ arises from the source voltage variations, $v_s(t)$, and is thus a function of $v_s(t)$ that is parame-

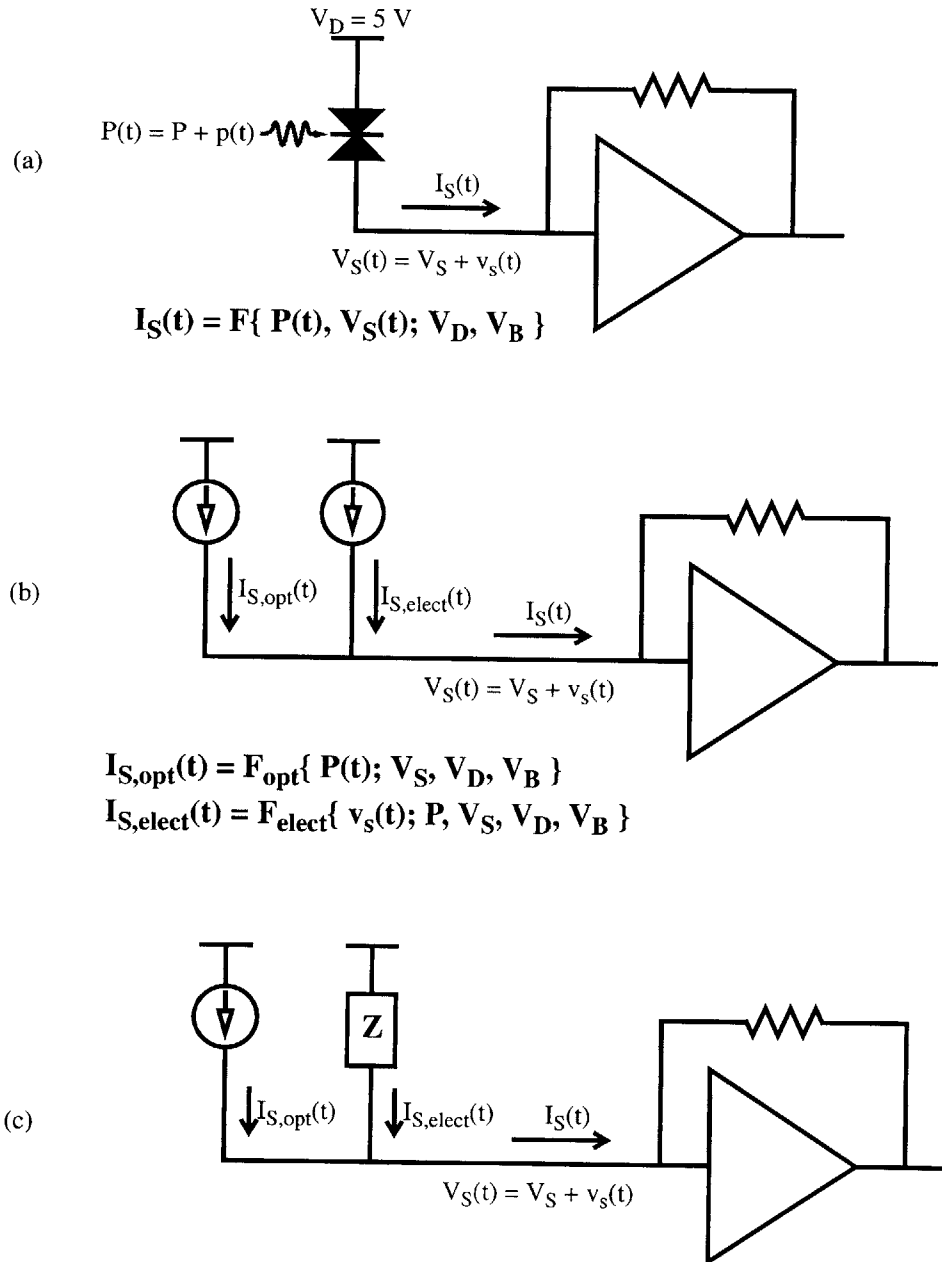


Figure 4.19: Evolution of a black-box model for MSM27

(a) In general, the MSM27 source current has a non-linear and non-time-invariant functional dependence on the optical input signal, $P(t)=P+p(t)$, and the source voltage $V_S(t)=V_S+v_s(t)$. (b) The total source current is broken into terms which originate from $p(t)$ and $v_s(t)$ separately, having only a time-average dependence on each other. Finally, the small-signal nature of $v_s(t)$ is used to simplify the model to that shown in (c). $I_{S,opt}(t)$ must, in general, be directly measured since its dependence on $P(t)$ is still not linear or time invariant.

terized on the bias points V_S , V_D , and V_B , and on the average power P . The latter statement will be found to be a reasonable approximation based on measurements presented below¹. Finally, since $v_s(t)$ is a small signal, the $I_{S,elect}(t)$ function may be linearized resulting is the impedance block shown in Figure 4.19(c). Note that the dependence of $I_{S,opt}(t)$ on the input optical signal is still not necessarily linear or time invariant. $I_{S,opt}(t)$ must be explicitly measured for a given optical input signal.

Measuring the detector dynamics

To measure the detector current source, $I_{S,opt}(t)$, and the terminal impedance, Z_{MSM} , a ground-signal-ground co-planar probing strategy has been employed. This eliminates parasitic components that would otherwise obscure the measurement by attaching the source and drain of the detector directly to the end of a 50 Ω transmission line. A portion of the test chip MIT-OEIC-7 was dedicated to ground-signal-ground (GSG) probable detectors, and is shown in Figure 4.20. The GSG probes may be used with the HP 8510B network analyzer in order to determine the terminal impedance from S_{11} measurements, or they may be used with an oscilloscope to directly view the photocurrent step or impulse response (the measured voltage waveform is related to the detector current by a 50 Ω characteristic impedance). In either case, the ground probes are connected to earth ground through the instrument. Thus, the ground connection is used as the MSM drain, and the source (signal connection), backgate, and isolation are set to negative voltages.

In carrying out the S_{11} measurements, the on-chip open, short, and 50 Ω load structures were used to calibrate the 8510B. In this way, the parasitic effects of the probe pads were removed from the measurements. Since no such calibration is possible when observing the photocurrent step/impulse response, the pads themselves were also measured to determine the equivalent circuit shown in Figure 4.21 (in this case, the 8510B was calibrated using HP-supplied calibration structures). In measuring the photocurrent response, both the pad parasitics and the terminal impedance of the MSM will appear in parallel with a transmission line characteristic impedance of 50 Ω . In this case, the effects of both the pads and the terminal MSM impedance, which will be shown momentarily, are found to be inconsequential. That is, the bandwidth limitation resulting from the combination of the MSM terminal impedance, the GSG probe pad

1. Measurements of the equivalent input impedance which gives rise the $I_{S,elect}(t)$ show that it depends minimally on the static optical power level. Thus, as far as $I_{S,elect}(t)$ is concerned, $p(t)$ may be considered a small signal so that only the bias point of the optical signal, P , is used in determining $I_{S,elect}(t)$.

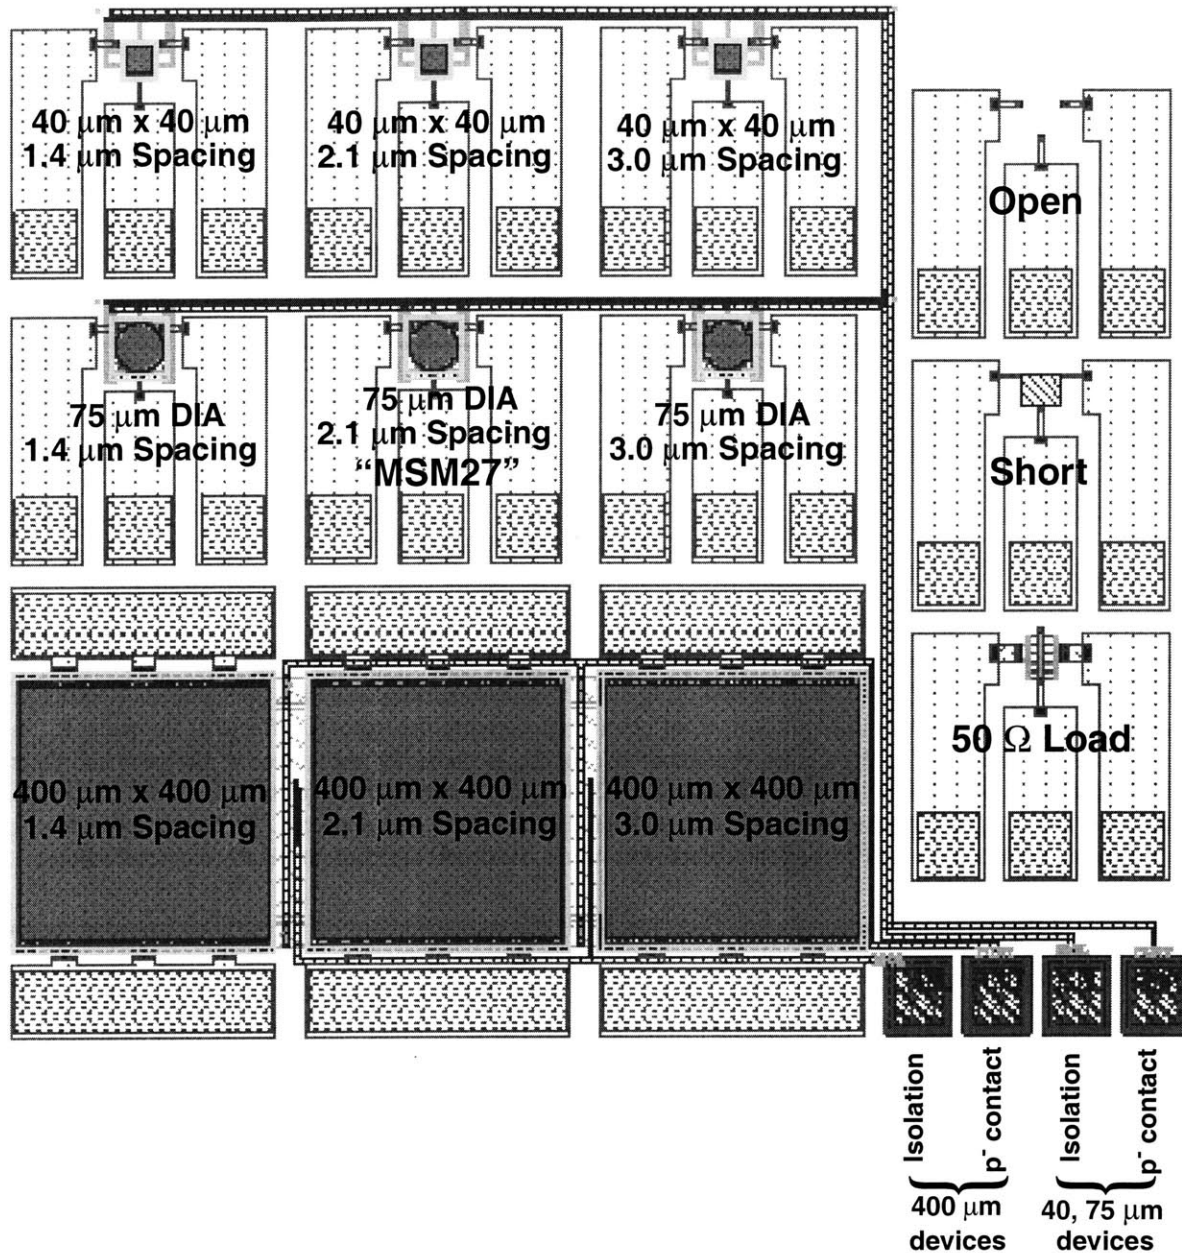


Figure 4.20: Layout of detector test structures on MIT-OEIC-7

Large area (400 $\mu\text{m} \times 400 \mu\text{m}$) devices are provided for static photocurrent and C-V measurements. 75 μm diameter and 40 μm square devices with ground-signal-ground probe pads are used for static and dynamic photocurrent measurements and for S_{11} measurements. The ground pads connect to one set of MSM fingers and the signal connects to the other. Open, short, and 50 Ω calibration structures are provided for use in S_{11} measurements.

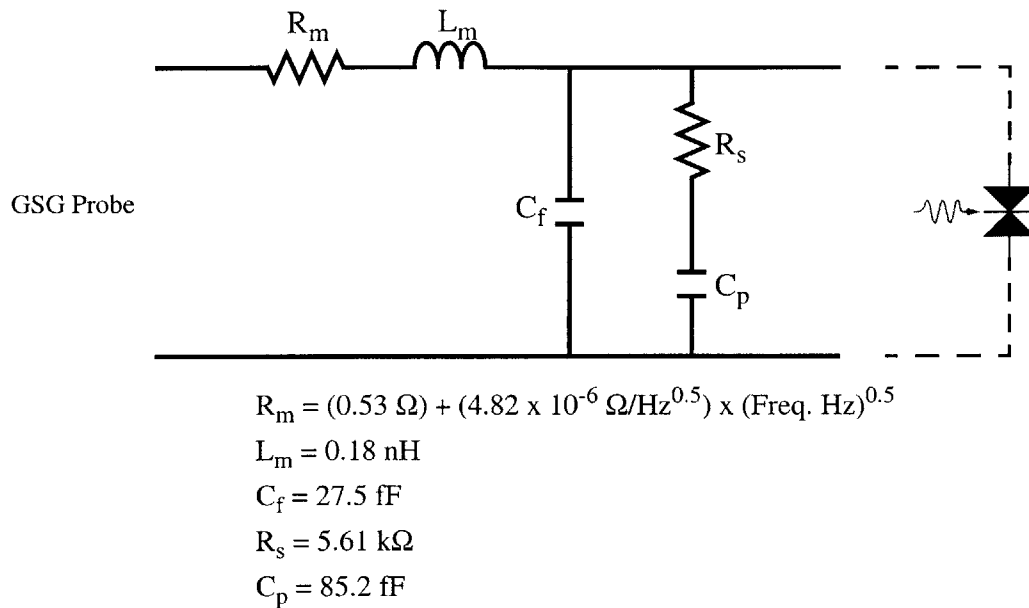


Figure 4.21: Equivalent circuit for ground-signal-ground pads seen in Figure 4.20

The HP 8510B network analyzer was calibrated using an HP standard substrate and the open and short structures in Figure 4.20 were used to determine this equivalent circuit representing the parasitic effects of the pads. R_m is the series resistance of the pads, and includes skin-depth effects. L_m is the series inductance of the pads and C_f is the fringing capacitance between the ground and signal pads. C_p accounts for the parallel-plate capacitance of each pad to the substrate, and R_s is substrate resistance through which these capacitors are connected. This model provides a good fit to the measured data up to 8 GHz, and the component values are consistent with physically expected values.

parasitics, and the 50 Ω transmission line does not limit the measurement of the step or impulse response of the MSM photocurrent.

The design of Figure 4.20 does not take into account the dynamics of the MSM backgate. There are significant parasitics associated with the backgate connections. These include the capacitance of the long on-chip interconnects which lead to the p^- contacts and the off-chip bond-wire and cable inductances. The latter tends to defeat the grounding of the backgate at high frequencies, while the former can act as a bypass capacitor on this node. The overall effect of this test structure on the MSM measurements is not certain. However, similar wiring is used to connect to the backgate of the MSM27 incorporated into the receiver of Chapter 7, so the MSM measurements made here may be applied to understanding the receiver characteristics. On the other hand, the detectors used in the array chips discussed in Appendix G have direct backgate-to-ground connections and thus may behave somewhat differently.

Terminal impedance

The terminal impedance of MSM27 was determined from measurements of S_{11} made at the bias conditions summarized in Table 4.1. The S_{11} measurements, taken from 45 MHz to 20 GHz, are plotted in Figure 4.22. There is a clear change between the unilluminated S_{11} characteristics and the others. However, S_{11} does not change significantly as the optical power level is increased from below the I_S threshold (all current flow is between drain and backgate; light level is only a few μW) up to very high levels ($I_S=265 \mu\text{A}$; based on Figure 4.15 the light level is near 1 mW). It is this observation that allows the simplification of the black box model discussed above.

The S_{11} data may be transformed to give the detector's terminal impedance using the formula, $S_{11}=(Z_{\text{load}} - Z_0)/(Z_{\text{load}} + Z_0)$, for the reflection from a load of impedance Z_{load} at the end of a transmission line of characteristic impedance Z_0 . Figure 4.23 plots, versus frequency, the real and imaginary parts of MSM27's terminal admittance (reciprocal of the impedance).

$-I_S$	$-I_B$	I_D
~ 0	6.3 μA	6.3 μA
5.1 μA	10.4 μA	15.5 μA
15.2 μA	12.0 μA	27.2 μA
91.7 μA	14.1 μA	106 μA
265 μA	14.3 μA	279 μA

Table 4.1: Summary of bias conditions for MSM27 S_{11} measurements

V_D was grounded while the backgate contact, isolation, and substrate were set to -5 V. V_S was set, through a bias-tee internal to the HP 8510B, to -4.675 V. This mimics the bias condition of the MSM27 used in the receiver of Chapter 7 in which the receiver input is at 0.325 V, the backgate is grounded, and the drain is set to 5 V. S_{11} was measured with no incident light, and with 850 nm light resulting in the terminal currents shown above. I_S and I_B were measured while I_D was determined from their sum. An offset current of 4.7 μA was present in the I_S measurements (was present with no device attached) and has been subtracted from the figures above.

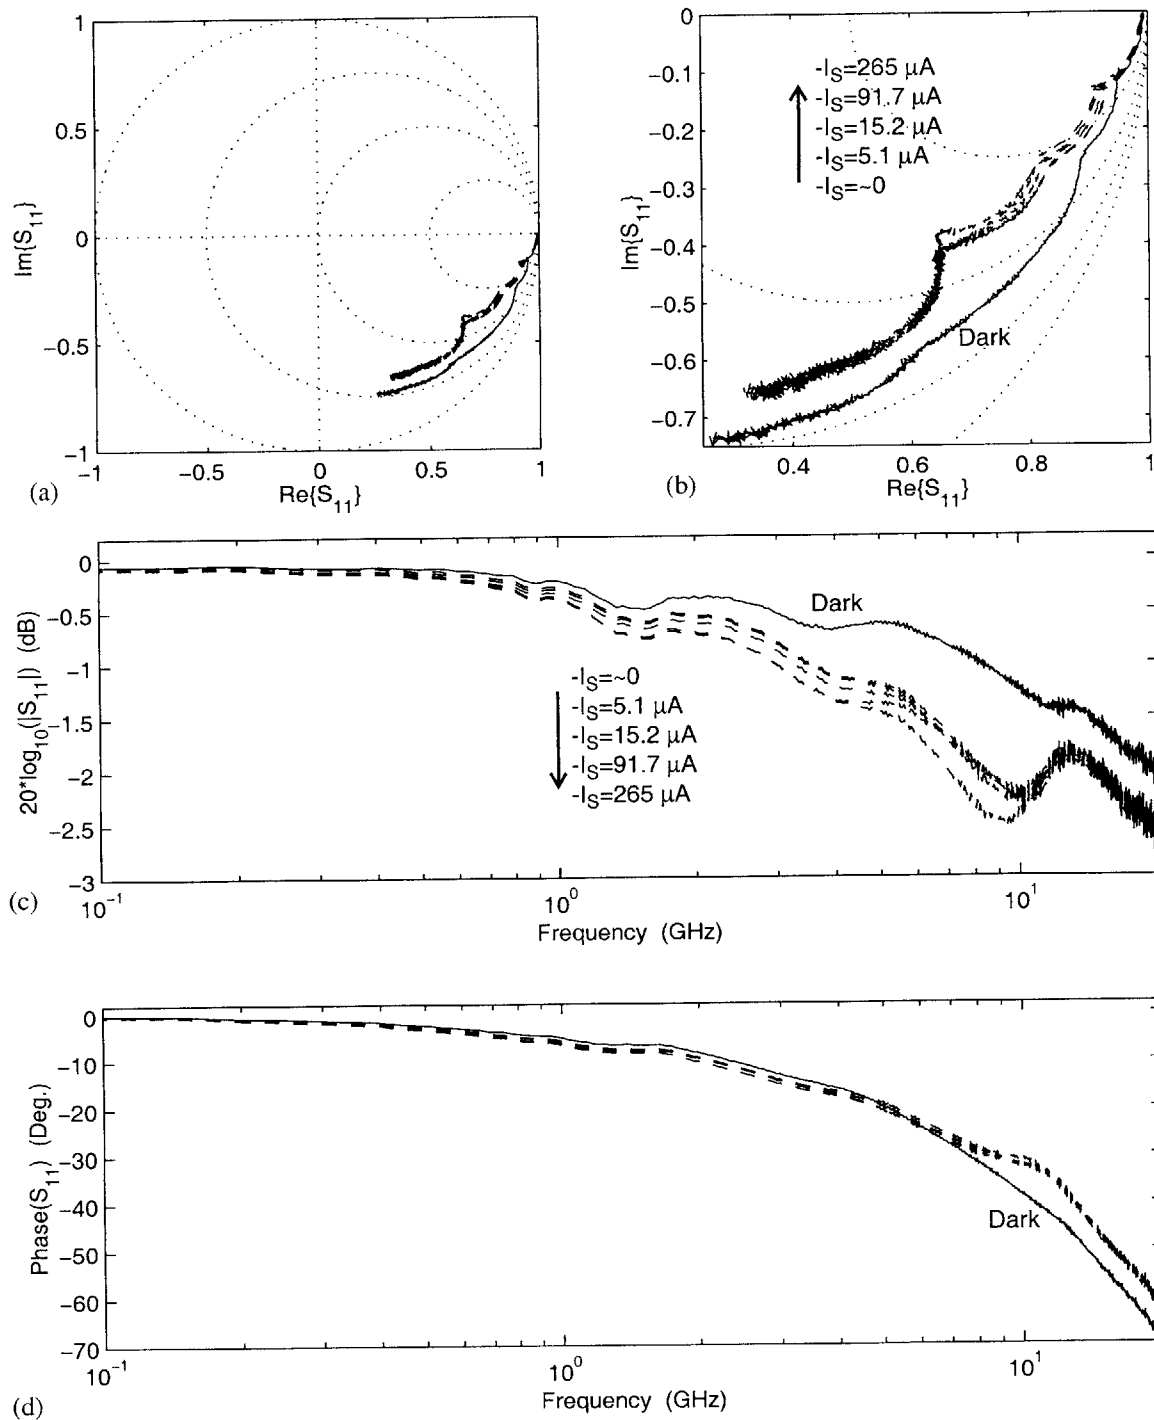


Figure 4.22: Measured S_{11} of MSM27

S_{11} was measured from 45 MHz to 20 GHz under the bias conditions summarized in Table 4.1. The measurements are plotted on a Smith chart in (a) and (b), while (d) and (e) show the magnitude and phase of S_{11} versus frequency.

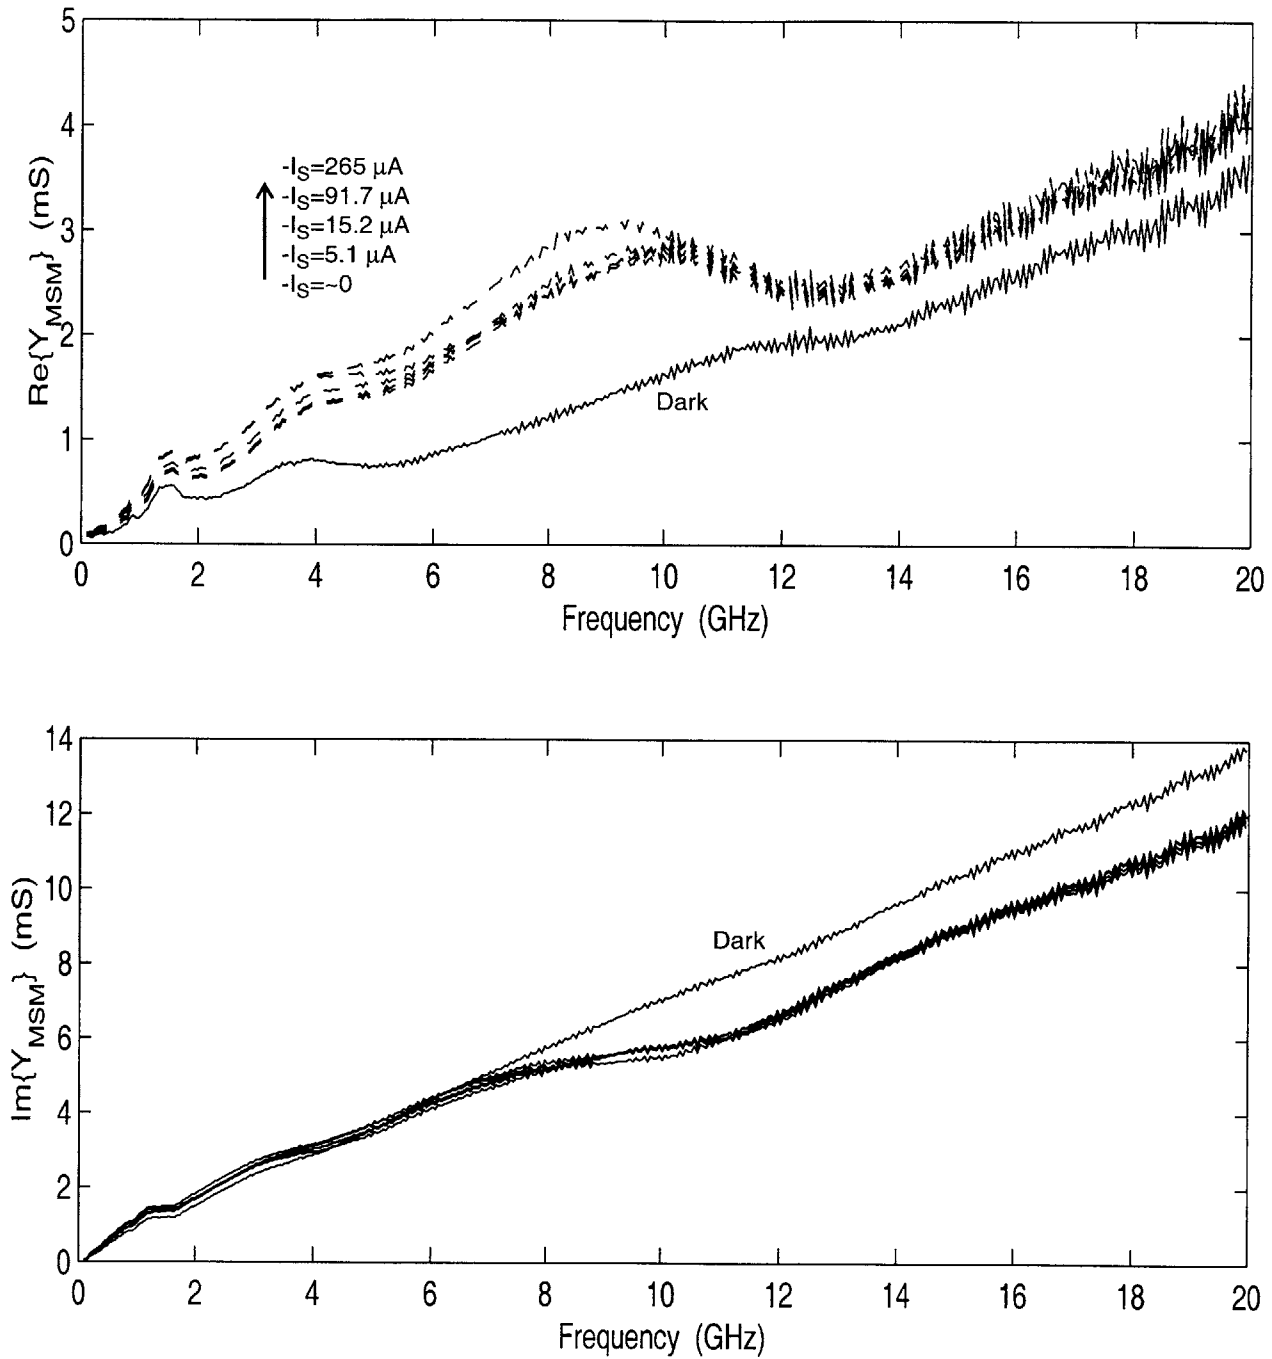


Figure 4.23: Terminal admittance of MSM27

The terminal admittance of MSM27 is determined using the formula $Y_{\text{MSM}} = 1/Z_{\text{MSM}} = (1/50 \Omega) \times (1-S_{11}) / (1+S_{11})$ to transform the S_{11} data of Figure 4.22.




Figure 4.24: Equivalent circuit for MSM27 terminal impedance

(a) Summarizes the black-box MSM model, which consist of the photocurrent source, $I_{S,opt}(t)$, in parallel with a terminal impedance, Z_{MSM} . $Y_{MSM}=1/Z_{MSM}$, determined from S_{11} measurements, is plotted in Figure 4.23. (b) Shows a simple circuit used to approximate the measured terminal impedance. This is not a physically meaningful model, but does reproduce Z_{MSM} accurately enough for the present needs. Component values are given in Table 4.2.

The measured terminal impedance or admittance may be used directly in calculations. Alternately, a circuit having a similar impedance may be constructed for use in SPICE simulations. A simple RC network which gives an adequate fit to the data is shown in Figure 4.24 and best-fit component values at each bias level are given in Table 4.2. The S_{11} and admittance data corresponding to “dark” and “ $I_S=15.2$ mA” bias conditions are compared with the equivalent circuit results in Figure 4.25 and Figure 4.26, respectively. Although the equivalent circuit model does not have a structural basis, it reproduces the measured data adequately for the purpose of understanding the effect of MSM27 on receiver performance.

$-I_S$	R_1	C_1	R_2	C_2
~ 0	15.3 Ω	90.6 fF	520 Ω	111 fF
5.1 μA	14.2 Ω	91.3 fF	497 Ω	113 fF
15.2 μA	15.6 Ω	91.1 fF	524 Ω	125 fF
91.7 μA	16.4 Ω	91.0 fF	526 Ω	149 fF
265 μA	13.6 Ω	90.5 fF	449 Ω	148 fF
Dark	15.8 Ω	112 fF	1.50 k Ω	103 fF

Table 4.2: Component values for the MSM27 terminal impedance equivalent circuit

For each of the bias conditions summarized in Table 4.1, the component values of the model shown in Figure 4.24 were optimized, using the simplex algorithm, on the basis of a least-squares fit to the S_{11} data.




Figure 4.25: Comparison of measured and modeled S_{11} of MSM27

For the "dark" and " $-I_S = 15.2 \mu\text{A}$ " bias conditions, the measured S_{11} data (dashed lines) is compared with simulated S_{11} (solid lines) based on the equivalent circuit shown in Figure 4.24.

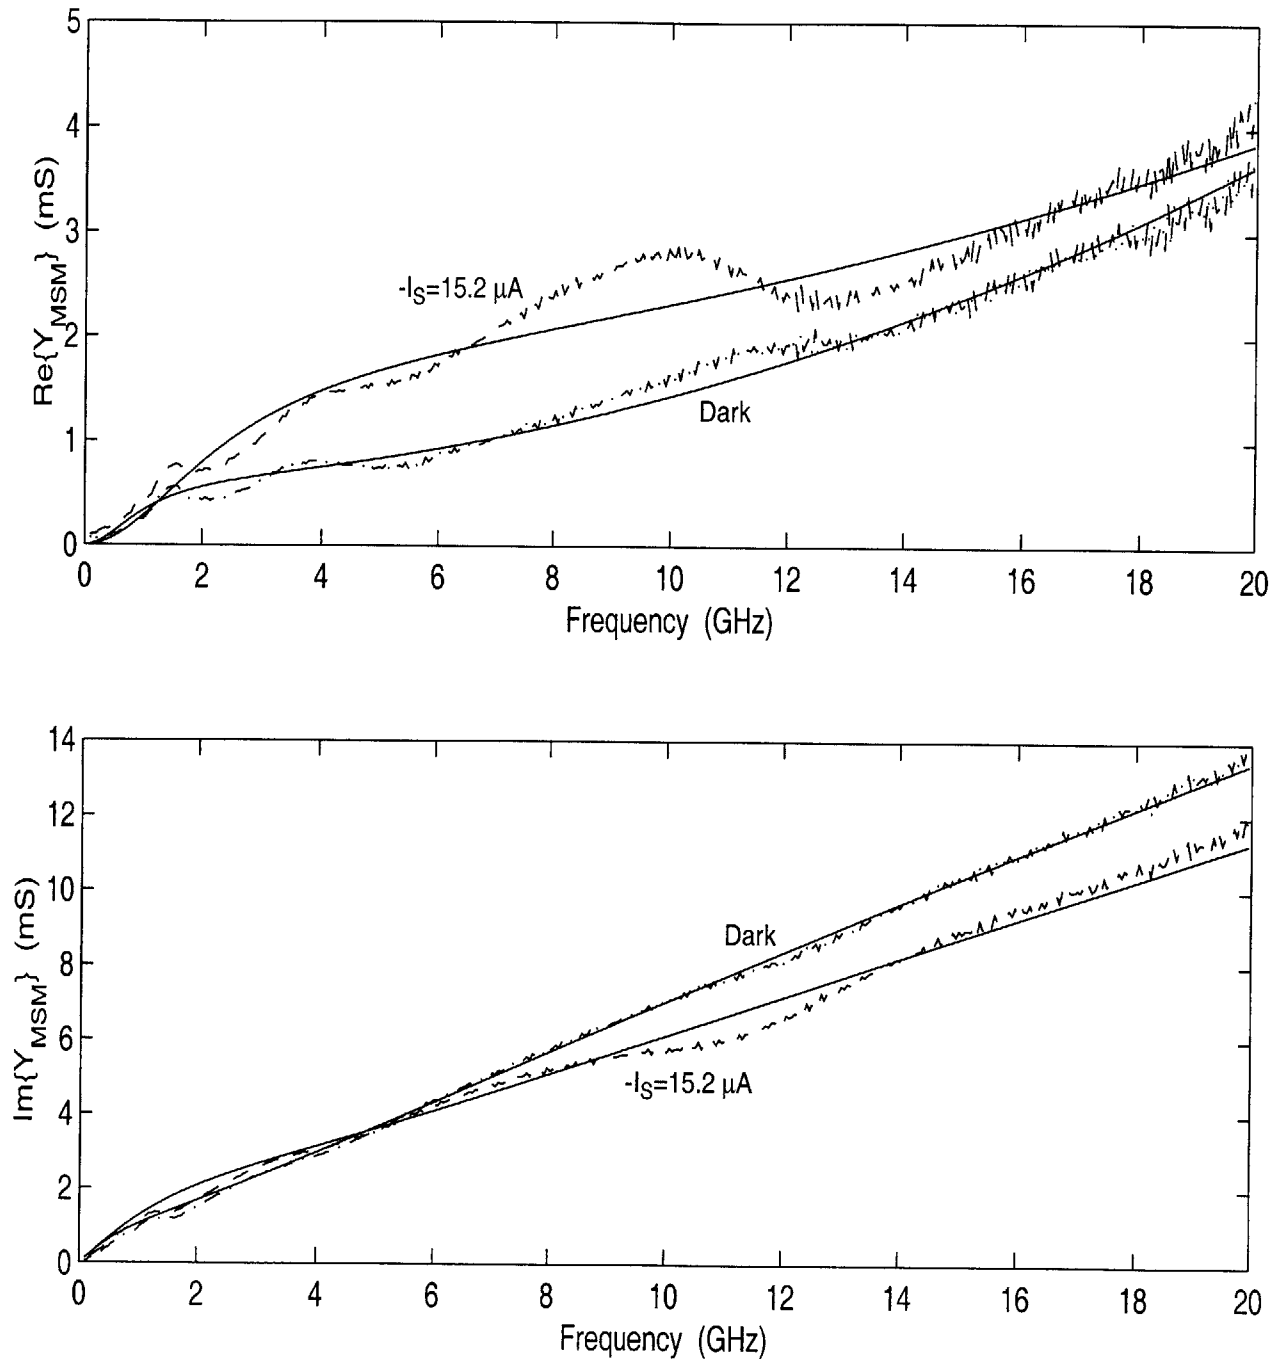


Figure 4.26: Comparison of measured and modeled conductance of MSM27

For the "dark" and " $I_S=15.2 \mu\text{A}$ " bias conditions, the measured admittance data (dashed lines) is compared with simulated admittance (solid lines) of the equivalent circuit shown in Figure 4.24.

Dynamic photocurrent

It is common practice to characterize the dynamic behavior of a photodetector by either its impulse response, excited by a mode-locked laser, or by its small-signal frequency response, measured using an sinusoidally modulated laser [161]. This assumes that the detectors photocurrent is well modeled by a linear, time-invariant system. Unfortunately, this assumption does not hold for the detector structures being examined here. Instead, the photocurrent, $I_{S,opt}(t)$, must be directly measured for the actual optical signal that is of interest. In a digital optical link, the optical signal is stepped between high and low levels. Here, the dynamics of MSM27 will be characterized in terms on its step response¹.

To excite the MSM step response, an 850 nm VCSEL was modulated using an ECL signal. The experimental setup used for this purpose is described in Appendix C. Rising and falling laser output waveforms from this setup are shown in Figure 4.27. Reflection in the signal path leading to the VCSEL appear in the output signal for around 8 ns after the initial, rapid transitions. Fortunately, the optical signal is still faster than MSM27, and thus may be used to characterize it.

Since the input signal to the VCSEL was of a fixed amplitude, the magnitude of the optical step was varied by attenuating the optical signal. Figure 4.28 and Figure 4.29 show the measured step response of MSM27 to various optical input step heights. The average terminal currents and step heights are summarized in Table 4.3. As in the S_{11} measurements, the MSM drain was set to 0 V and the source to -4.675 V. These measurements were carried out both with the backgate set to -5 V and with the backgate open-circuited. The isolation contact and the substrate were set to -5 V.

The waveforms in Figure 4.28 and Figure 4.29 take from 10 ns to 20 ns to settle, depending on the optical step height. These waveforms may further be summarized by a number of observations. 1) At low optical power levels, the backgate-grounded data differs greatly from the open-backgate data. At high incident power, the two cases are nearly identical. 2) At low power levels, the rising and falling waveforms are very different from each other. The rising waveforms include a secondary, delayed step. The time delay to this step decreases with power level. At high power levels, the rising and falling edges are quite similar.

1. To assess the detector performance in a data link, a pseudorandom bit stream should be used to modulate the light source. This will be done in Chapter 7 in the context of the understanding the behavior of an optical receiver. In the present chapter, the step response will give a good general indication of the MSM's dynamic characteristics.




Figure 4.27: Step waveforms of ECL-driven VCSEL used to produce Figure 4.28 and Figure 4.29

The experimental setup used to generate this optical signal is described in Appendix C. The rising and falling edges of fiber-coupled VCSEL output are shown here. Reflection at the input of the VCSEL effect the output for around 8 ns after an initial, rapid transition, but the optical signal is still faster than MSM27.

This is true of both the backgate-grounded and backgate-open data, but the delayed step vanishes at much lower power levels in the open-backgate data. 3) Aside from the secondary step, the transitions are faster at higher power levels. An exception to this is the falling edge of the backgate-grounded data. These transitions initially become slower with increasing power level, up to the point where they are similar to the open-backgate waveforms, then become slightly faster as the power level is increased further.

These observations are consistent with the physical description given above wherein the source current is controlled by the backgate. The delayed secondary step is likely related to the time required, at a given backgate current level, to charge the backgate node to the point where the backgate diode becomes forward biased. Once the backgate diode is forward biased, the grounded-backgate and open-backgate data look similar.




Figure 4.28: MSM27 step response at various modulation levels -- 1 of 2

The plots on the left were made with the backgate grounded. The backgate was open circuited to produce the waveforms on the right. The optical input waveform is included in each plot (dashed line) for reference. The time-average source current, $\langle -I_S \rangle$, indicated in the plots corresponds to the summary in Table 4.3. The waveforms have been normalized and their steady-state high and low values noted along the right axis. Continued in Figure 4.29.

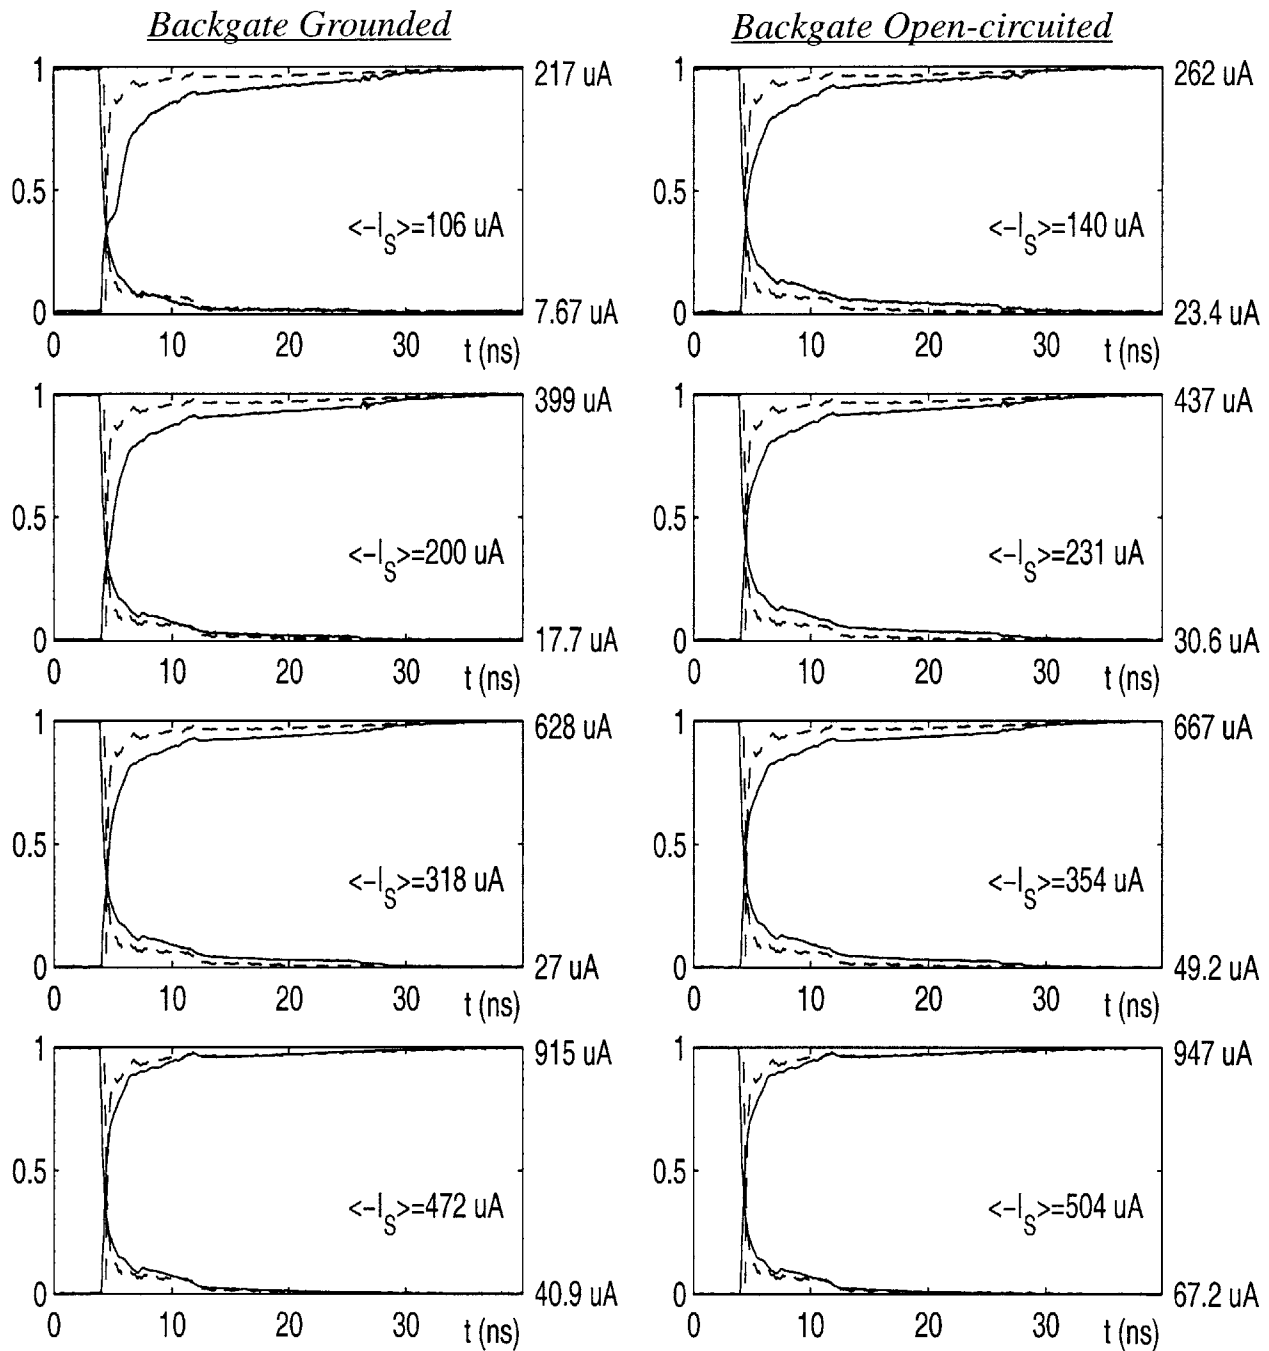


Figure 4.29: MSM27 step response at various modulation levels -- 2 of 2
Continuation of Figure 4.28

<i>Backgate Grounded</i>				<i>Backgate Open-circuited</i>	
$\langle -I_S \rangle$ (μA)	$\langle -I_B \rangle$ (μA)	$\langle I_D \rangle$ (μA)	Step Height (μA)	$\langle -I_S \rangle$ (μA)	Step Height (μA)
2.0	9.6	12	4.8	12	17
13	17	30	31	28	48
31	22	53	67	54	88
74	24	98	145	102	171
106	27	133	209	140	238
200	34	234	380	231	406
318	39	357	600	354	617
472	32	504	874	504	880

Table 4.3: Summary of MSM27 step response measurement parameters

The MSM27 step response data was obtained by ground-signal-ground probing the device and recording the resulting waveform on an oscilloscope. V_D was grounded while the backgate, isolation, and substrate were set to -5 V. V_S was set, through a bias-tee, to -4.675 V. This mimics the bias condition of the MSM27 used in the receiver of Chapter 7 in which the receiver input is at 0.325 V, the backgate is grounded, and the drain is set to 5 V. For comparison, measurements were also made with the backgate open-circuited. The step response was measured with 850 nm light at various optical step heights. The values of $\langle I_S \rangle$ and $\langle I_B \rangle$ reported above were measured while $\langle I_D \rangle$ was determined from their sum. The step height is measured between the steady state low and high values.

The low power, backgate-grounded waveforms can not be described in terms of a linear, time-invariant system response. It is for this reason that the empirical model of the MSM uses a direct measurement of the photocurrent. However, the similarity of the rising and falling edge data at high power levels opens the possibility of implementing a more predictive, linear model for the photocurrent in this operating regime. To this end, the impulse response of MSM27 was measured using an 850 nm mode-locked laser. A static 850 nm light source was coupled with the pulse signal to provide an optical bias level. A representative impulse response from MSM27 taken with 55 mA of average source current is shown in Figure 4.30. At high power levels, this impulse response may be used to predict the current from the optical signal waveform. This is demonstrated in Figure 4.31 which shows measured optical and I_S waveforms along with the result of filtering the optical waveform with the impulse response of Figure 4.30.




Figure 4.30: MSM27 impulse response

MSM27 was ground-signal-ground probed and excited by a mode-locked laser emitting 850 nm pulses of roughly 1 ps duration. The mode-locked laser had an 80 MHz repetition rate and an average power of 127 μW . A static optical bias was also applied, and the device was biased with $V_{\text{DS}}=4$ V and $V_{\text{BS}}=1$ V. The average source current in this case was 55 μW .

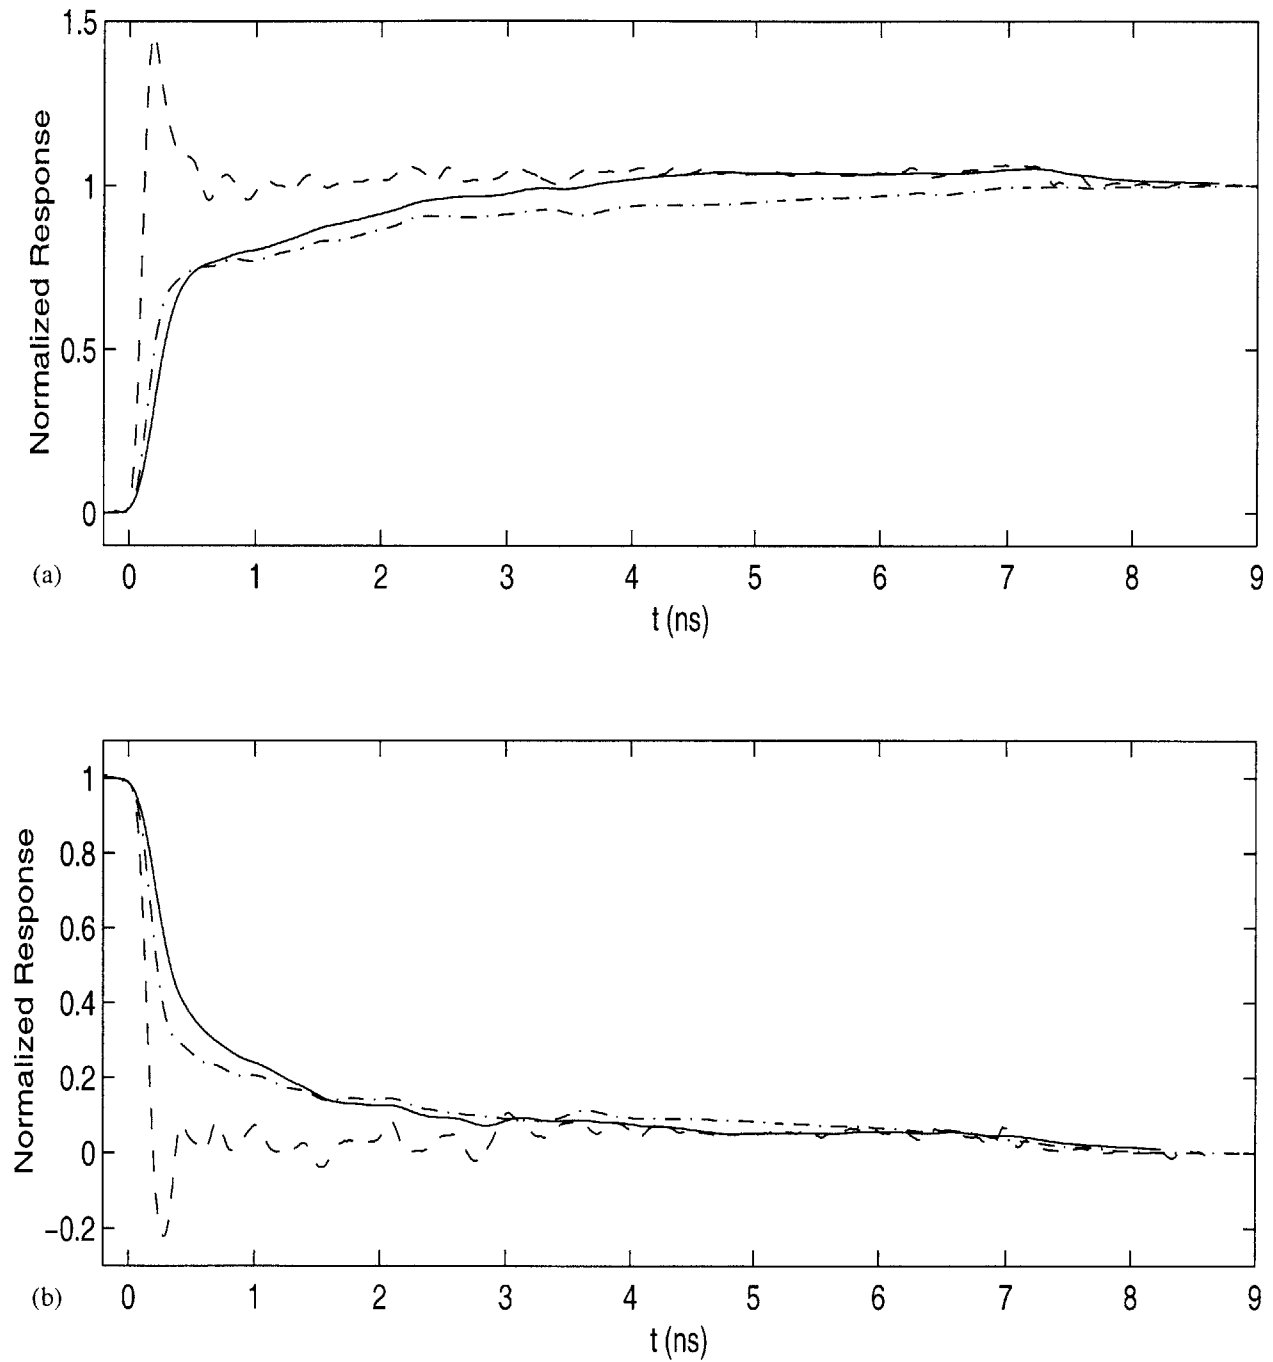


Figure 4.31: Linear-system model of MSM27 photocurrent at high optical power levels

Normalized rising and falling steps are shown in the (a) and (b), respectively. The dashed curves are the measured optical signal and the dashed-dotted curves are the measured MSM27 source current. The solid curves are produced by convolving the optical signal with the impulse response shown in Figure 4.30.

Summary

This section has examined a novel photodetector structure implemented in Vitesse's H-GaAs IV process. Despite its considerable drawbacks, this structure represents a step forward in the development of a high performance photodetector for use in large scale optoelectronics circuits. Aside from the details of the detector itself, an important aspect of this evolution is the introduction of the backgate isolation structure which eliminates the crosstalk effects observed on previous OEIC generations.

A physical model of the experimental MSM-like photodetector was presented which captures much of its functionality. At the core of this model is the understanding that the source current of the detector is generated as a consequence of the backgate p⁻-n junction being driven into forward bias by the photocurrent which it collects. This leads to the undesirable static and dynamic properties which have been observed. Although this model is not complete, and in all likelihood contains inaccuracies, the issues surrounding the essential behavior of the backgate appear to be well established.

While an alternate detector structure will need to be developed in the future, the current devices are included in the integrated receivers discussed elsewhere in this thesis. To help in assessing the effects of the detector on the receiver's performance, an empirical model of the detector has been constructed. This model consists of a linear impedance in parallel with a current source. In general, the photocurrent is not well represented by a linear, time invariant system. The photocurrent resulting from a given optical input must thus be measured directly and coupled with the terminal impedance. In the case of a large optical input signal (large enough to strongly forward bias the backgate), the photocurrent response does become linear, and may be predicted from the optical signal by making use of the detector's impulse response.

4.3 A Proposed New Photodetector Structure


As pointed out at the beginning of this chapter, Vitesse maintains an MSM photodetector process in which the p^- layer is not present below the detector structure. Since there is a significant additional cost associated with implementing this process, developing a photodetector structure which can coexist with the p^- layer continues to be an important aspect on the OEIC development effort. As seen in the previous section, the p^- layer will collect photocarriers if it is given the slightest opportunity to do so. This suggests making the backgate junction an integral part of the device.

An example of this was given above: using one set of Schottky fingers in conjunction with the grounded backgate. Ironically, if this is done with the structure shown above the backgate can not be fully depleted if current is to flow to the p contacts at the perimeter of the device. To avoid this problem, and to reduce the device capacitance over a wide range of optical power levels, p-contacts may be interspersed with Schottky contacts made to n-type materials. An example of this is given in Figure 4.32.

A lateral pin photodiode

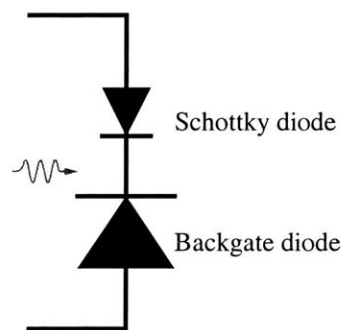
A drawback of the structure shown in Figure 4.32 is that it can only be used with the Schottky contacts connected to the transimpedance amplifier input and the p contact set to a lower potential. As seen above, there is substantial current flow out of the Schottky contacts when the p contact is biased around 0.5 volts above the source (see dark current in Figure 4.8). Also, it still uses a non-standard mask layer at two points in the process in order to block the LDD and source/drain implants. However, aside from a possible, small decrease in capacitance, there is no longer any need to do this.

If the source/drain implant is put back in place and an n-type ohmic contact is made to it, the result is a lateral p- p^- -n diode. It will be referred to here as simply a pin diode, and a possible implementation of it is shown in Figure 4.33. To maximize the detector's fill factor (percent of device area that is not blocked by metal), it is laid out as a grid of minimum-width, source/drain implanted active area and an isolated p contact is placed at the center of each grid element. The p-contacts are connected by minimum-width gate metal wires which are routed through breaks in the n^+ grid. The device may be laid out by forming a unit cell, as in Figure 4.33(b), which is arrayed to fill the desired shape. A 75 μm diameter example of this is shown in Figure 4.34.



Vitesse H-GaAs IV


(a)



(b)

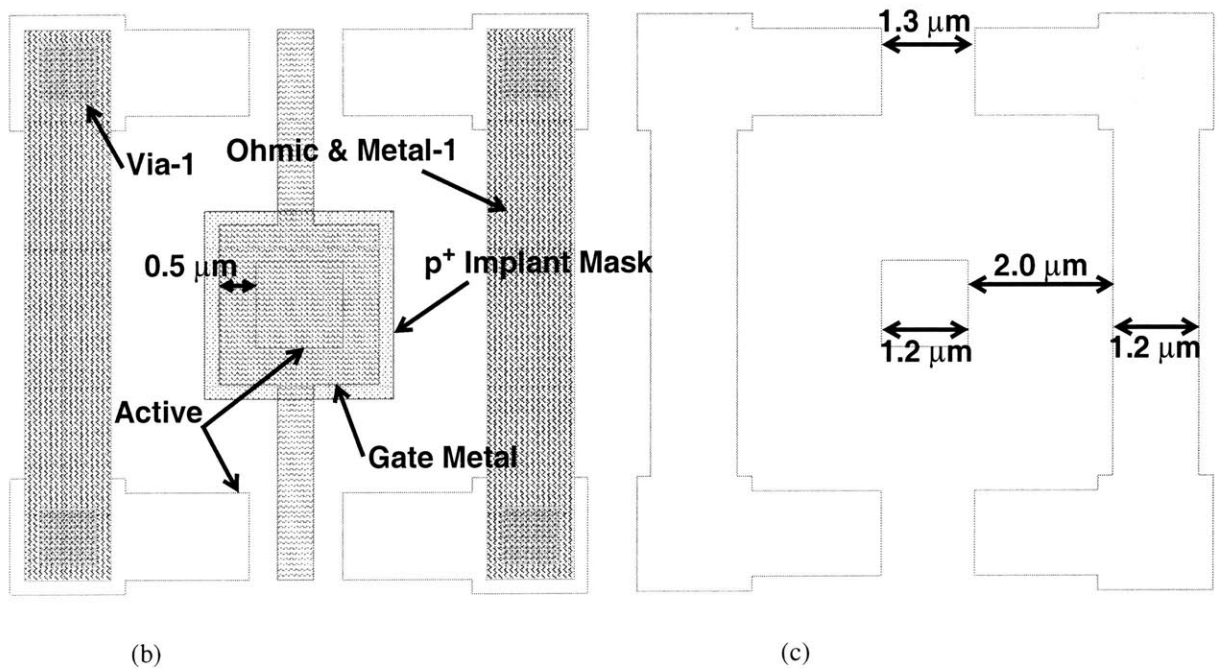
Figure 4.32: Modification of MSM-like structure to use backgate diode as photodetector

The structure of the MSM-like detector shown in Figure 4.5 may be modified to use only one set of Schottky fingers interspersed with p contacts. (a) shows a cross section of this type of device while (b) shows a circuit representation.



(a)

Vitesse H-GaAs IV



(b)

(c)

Figure 4.33: Lateral pin photodiode in the standard Vitesse H-GaAs IV process

(a) Shows the cross section of a lateral pin photodiode implemented using the standard H-GaAs IV process. The diode is formed using p^+ backgate contacts and n^+ source/drain implants. The unit cell layout shown in (b) may be arrayed on a $6.4 \mu\text{m} \times 6.4 \mu\text{m}$ grid to form a detector of a desired size and shape (see Figure 4.34). (c) Shows only the active layer of the layout in (b). The gap in the active pattern is provided so that the p contact can be routed in minimum width ($0.5 \mu\text{m}$) gate metal. Using minimum allowed dimensions, this layout achieves a fill-factor of 64% (i.e. 36% of the surface is covered by metal).

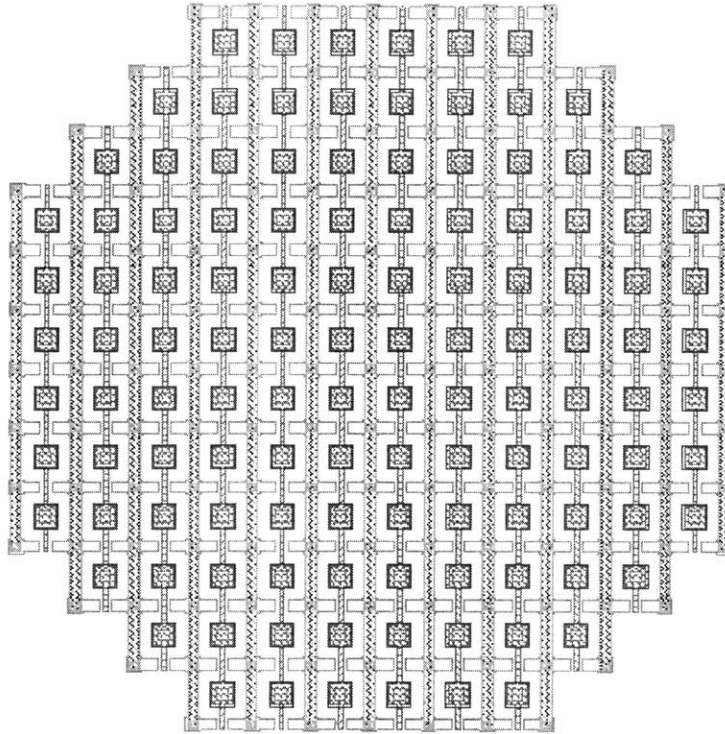


Figure 4.34: 75 μm diameter lateral pin photodetector

A 12 x 12 array of the unit cell in Figure 4.33(b) was formed and 6 cells were removed from each corner to form this 75 μm diameter photodetector. To complete the layout, the vertical metal-1 (n-contact) and gate metal (p-contact) interconnects must be contacted at the perimeter. An isolation ring also needs to be placed around the device. If the n^- contact is to be connected to a positive supply, it may be possible to merge the isolation ring into the detector.

Figure 4.35 qualitatively shows band profiles which may be encountered in the device. Profiles are given with the device unbiased and connected to a transimpedance amplifier. Both positive- and negative-supply configurations are shown. The substrate is assumed to be grounded in all cases. However, the substrate resistance varies from several 10 $\text{k}\Omega$'s for a large die ($\sim\text{cm}^2$) to several $\text{M}\Omega$'s for a small die ($\sim\text{mm}^2$).

In most H-GaAs IV circuits, all p^- contacts are set to ground and the substrate is grounded. However, the H-GaAs IV design rules state that the substrate should be left floating in a circuit which contains p^- regions that are to be maintained at a potential higher than ground. This is to ensure that the p^- - n^- diode formed by the backgate and the substrate is not forward biased¹. It is not clear if there are negative conse-

1. The receivers of Chapter 7 and Appendix G violate this rule (it was not known in time), but the resulting substrate current is small because of the large resistance of the substrate.




Figure 4.35: Band profiles in lateral pin photodetector

Qualitative band profiles are shown along the path indicated in (d). (a) Shows the device with no applied bias. (b) and (c) show the device biased according to the inset figures. The substrate is treated as n^- material and is assumed to be grounded.

quences to leaving the substrate floating in the presence of photogenerated carriers. When the pin photodetector is biased with the p-contact connected to the input of a transimpedance amplifier and the n-contact set to a positive supply, as shown in Figure 4.35(c), the p⁻-n⁻ junction may be forward biased if the substrate is grounded. The resulting current will be limited by the resistance of the substrate and may or may not be a problem. It will, however, depend on the die area.

In Figure 4.33, the p⁺ and n⁺ implants are spaced at 2.0 μm. This gives a fill-factor of 64%. (By comparison, MSM27, with 0.5 μm fingers spaced at 2.5 μm, has a fill-factor of 83%.) At zero bias, the n⁺-p⁻ depletion region extends 0.4 μm into the p⁻ material. Vitesse estimates the capacitance of this junction at 0.1 fF per micron of edge. A reverse bias of 4 V across the pin diode increases the depletion width to 0.8 μm and should drop the capacitance to 0.05 fF per micron of n⁺ implant edge. The 6.4 μm x 6.4 μm unit cell in Figure 4.33(b) has an n⁺ edge length of 20.6 μm. Thus, under a 4 V bias it has around 1 fF of capacitance, or 0.024 fF/μm². For a 75 μm diameter detector, the total capacitance adds up to 120 fF (implemented as a 12 x 12 array with 6 units removed from each corner--total area is equivalent to a 79 μm diameter device).

A significant fraction of the photocurrent collection in the above device occurs in the undepleted n⁺ and p⁻ regions. While carriers generated within the depletion region are transported by drift within ~10 ps, the carriers generated in the undepleted regions must diffuse to the depletion region edge before being collected. The diffusion of electrons in the undepleted p⁻ region is expected to be the slower of the two diffusion processes because of the larger width of this region as compared to the n⁺ region. With the n⁺ and p⁺ implants spaced by 2 μm and a 0.8 μm lateral depletion width (at 4 V reverse bias), photogenerated electrons in the exposed 0.7 μm of the 1.2 μm p⁻ region will respond with a bandwidth of 6 GHz. This is shown in Figure 4.36, which plots the frequency response found by solving the diffusion equation in the frequency domain. The bandwidth of this current component may be increased by reducing the n⁺-to-p⁺ spacing. For a spacing of 1.7 μm, the diffusion time of electrons in the 0.4 μm of undepleted p⁻ material would be roughly comparable to that of the holes in the n⁺ region. The corresponding bandwidth would be 17 GHz, the fill factor would drop to 60%, and the capacitance (per unit area) would increase to 0.031 fF/μm². Most likely, the parallel combination of the detector capacitance with the input resistance of the receiver circuit will limit the response below this point. In practice, the n⁺-to-p⁺ spacing should be chosen in concert with the design of the receiver in order to equalize the RC and transport response times.

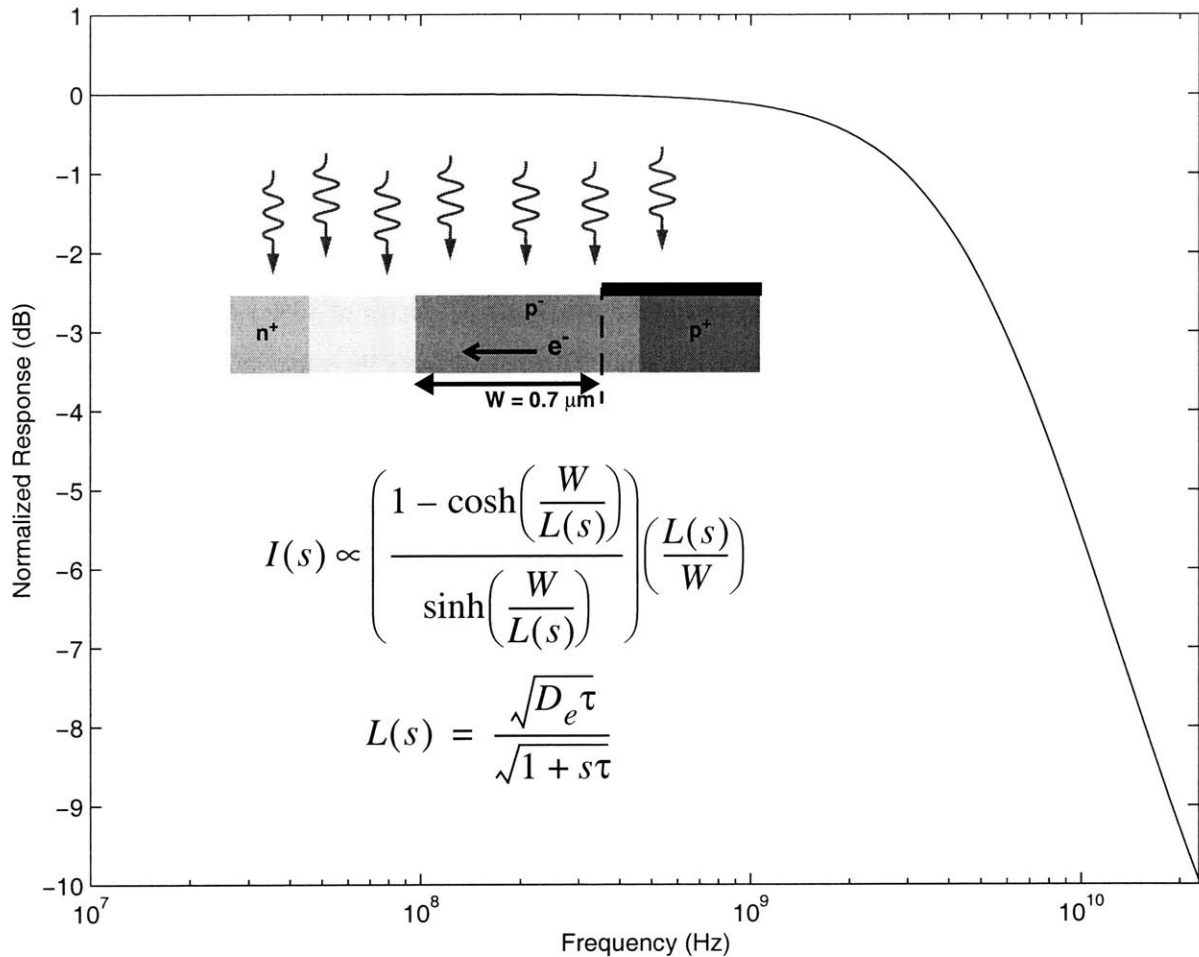


Figure 4.36: Frequency response of photogenerated carrier diffusion in lateral pin photodetector

With the n^+ and p^+ implants spaced by $2 \mu\text{m}$ and a $0.8 \mu\text{m}$ lateral depletion region, photogenerated electrons in the $0.7 \mu\text{m}$ of exposed p^- material must diffuse to the edge of the depletion region. Solving the diffusion equation in the frequency domain gives the frequency response of this component of the photocurrent. Assuming an electron mobility of $1000 \text{ cm}^2\text{V}^{-1}\text{s}^{-1}$ and a lifetime of 1 ns gives a bandwidth of 6 GHz . Note that diffusion of holes in the n^+ region is expected to be faster than this since the width of this region is much shorter.

Another possible bandwidth limitation may result from surface states. Surface states may slow the detector response by trapping photogenerated carriers and releasing them on a long time scale. In this regard, there does not appear to be any difference between the lateral pin detector and the MSM detectors used in Vitesse's product line operating at over 1 GHz .

4.4 Conclusion

The planar construction of the MSM photodetector makes it an ideal candidate for incorporation into a VLSI electronics process. Vitesse and others have taken advantage of this to make monolithically integrated optical receivers commercially available. In the case of the Vitesse process, implementing an MSM detector has required the elimination of the p-type doping within the detector area. This requires several masking steps in addition to the standard process. Thus, there is a motivation to demonstrate a detector structure which can coexist with the p-type backgate doping so that the processing cost may be reduced or even eliminated. Unfortunately, the MSM-like structures that have been studied thus far have been impaired by the action of the backgate junction.

Undeterred, a new detector structure has been proposed. The new structure uses the p^- material as part of a lateral p- p^- -n photodiode which is implemented using only the standard process steps. Although the doping of the p^- layer is too high to allow the structure to be fully depleted at a reasonable bias, the results of the preliminary analysis above indicate that the performance of this structure (its impulse response and capacitance) is consistent with the needs of EoE-based digital optical interconnects.

Another point to keep in mind is that EoE can be used to integrate photodetectors in addition to (or, in some applications, instead of) emitters. Two topics that have been investigated in the past include the integration of a telecommunication-wavelength photodetector based on lattice mismatched growth [119,162-166] and the fabrication of resonant cavity photodetectors from a VCSEL heterostructure [167,168].

5

Digital Optical Interconnect Systems

Chapter 1 gave an introduction to the role of optical interconnects in high speed digital systems, and the Epitaxy-on-Electronics integration technology, which is well suited to an important subset of optical interconnect applications, was described in Chapter 2. The integration of optical emitters with VLSI electronics was demonstrated in Chapter 3 with the help of LED structures, and one path to the integration of photodetectors was examined in Chapter 4. Technology development in the areas of emitter and detector integration is ongoing, with the former focused on the vertical-cavity surface-emitting laser (VCSEL) and the latter emphasizing the implementation of photodetector structures using the commercial VLSI electronics process. However, as observed in Chapter 3 based on experiences from the OPTOCHIP project, the technology development effort needs to encompass more than device integration. In order to make the OEIC technology accessible to system level users, the details of the circuit-level electrical-to-optical interface must be made transparent to them. That is, rather than providing optical emitters and detectors, the OEIC technology should include optical transmitters and receivers which interface directly with standard digital electronics.

Ultimately, both the characteristics of the optical channel and of the transmitter and receiver may be standardized, much in the way that logic families are. In the present developmental stage, however, the transmitter and receiver must be made flexible enough to function within dramatically different optical systems. At the same time, their performance level must be adequate to implement convincing demonstrations.

This chapter examines a variety of system-level issues preceding the design of the laser drivers in Chapter 6, and the receivers in Chapter 7, Appendix F, Appendix H, and Appendix I. It begins with a look at representative systems which could benefit from EoE-integrated OEICs. Typical characteristics of the link components are then reviewed, the structure of the digital optical link is examined, and its performance is studied analytically.

5.1 Epitaxy-on-Electronics OEIC Applications

Chapter 1 (Table 1.2) identifies two important optical interconnect applications in which EoE OEICs may play a role: smart pixel arrays for telecommunication systems and parallel optical interconnects for data-communication. This section looks at how EoE OEICs may be used in these applications and identifies a number of design guidelines for the transmitter and receiver interface circuits.

Tera-bit-per-second smart pixel arrays

OEICs will have a dramatic impact on telecommunication systems if they are able to provide an information processing capacity in excess of 1 Tb/s per cm² of die area. These OEICs may be used in either a free-space optical interconnect architecture or in conjunction with a planar lightwave circuit. In practice, an OEIC may include an arbitrary arrangement of electronic and optoelectronic functionality. However, it is convenient to conceptualize the OEIC as a “smart pixel” array in which each pixel is composed of one optical transmitter, one optical receiver, and some amount of digital processing circuitry.¹ In this context, the OEIC may be analyzed in terms of issues such as pixel density, processing complexity, electronic and optoelectronic bandwidth, and thermal management.

The data throughput density of a smart pixel array is the product of the pixel density and the bit-rate of the individual pixels. The pixel density and speed must be optimally chosen in light of the complexity of the resulting optical system, the difficulty of the electronic design, the extent of processing required within the pixels, and the packaging expense associated with managing the OEICs thermal environment. Consider a 1 cm² die containing 100 pixels operating at 10 Gb/s each, 1,000 pixels operating at 1 Gb/s each, and 10,000 pixels operating at 100 Mb/s. The first choice leads to very difficult electronics design while the last

1. The term “smart pixel” originated in conjunction with optical information processing and optical computing research to describe an imaging device in which each picture element possesses a degree of computational capability. The term is still in use in the context of optical interconnects in which case it simply refers to a VLSI circuit with a large number of optical inputs and outputs.

choice leads to very difficult optical system design. On the other hand, the choice of 1000 pixels/cm² operating at 1 Gb/s each minimizes optical system complexity and may be implemented with relative ease in GaAs VLSI electronics.

The power dissipation, and thus the processing complexity, are limited by the need manage the OEIC's thermal environment while minimizing packaging cost. Commercial GaAs VLSI devices dissipating up to 20 W/cm² are comfortably packaged in thermally enhanced plastic ball-grid arrays [169]. If the power dissipation of the 1 cm² OEIC die is limited to 20 W, each of the 1,000 pixels may dissipate 20 mW. For a pixel containing one emitter and one detector, this 20 mW may be distributed as 5 mW for the laser and laser driver, 5 mW for the receiver, and 10 mW for the processing electronics. Using direct coupled FET logic (DCFL) electronics in the Vitesse H-GaAs IV process capable of 1 Gb/s operation¹, the 10 mW budget allows for a 60-gate processing circuit. This level of complexity qualifies the smart-pixel as "smart" [170]. Sixty gates per pixel also equates to around 180 thousand transistors on the die (assuming an average fan-in of three). This transistor count is similar to that of existing electronic implementations of the targeted telecommunication components².

Figure 3.31 of Chapter 3 describes a demonstration smart pixel designed according to the above pixel density guidelines. A photograph of the cell is repeated in Figure 5.1(a). The cell is 300 x 300 μm, allowing 1,000 such pixels to be arrayed on a 1 cm² die. It contains one LED within an 85 μm square DGW, LED drive electronics, one OPFET-based receiver, and well over 60 gates. The large gate count points out that the above guidelines are limited by power density rather than by area. This canonical smart pixel exemplifies EoE's ability to provide the necessary level of electronic and optoelectronic complexity. To take advantage of this capability in a 1 Tb/s application, the LED and OPFET must be replaced by a vertical cavity surface emitting laser (VCSEL) and high speed photodetector, each accompanied by a suitable interface circuit. Figure 5.1(b) shows the elements of a future Tb/s/cm² EoE smart pixel array. A VCSEL in a 50 μm square DGW³, a 40 μm square MSM detector, and a laser driver and receiver are shown along

1. Conservatively ratioed 0.6 μm gate-length DCFL logic suitable for a 1 Gb/s design in the H-GaAs IV process dissipates around 160 μW per gate.

2. Many of these functions are implemented on Vitesse gate arrays which have raw gate counts ranging from 10,000 to 350,000.

3. A VCSEL process flow and layout were devised in order to verify that VCSELs in 50 μm square DGWs can be implemented on EoE OEICs.

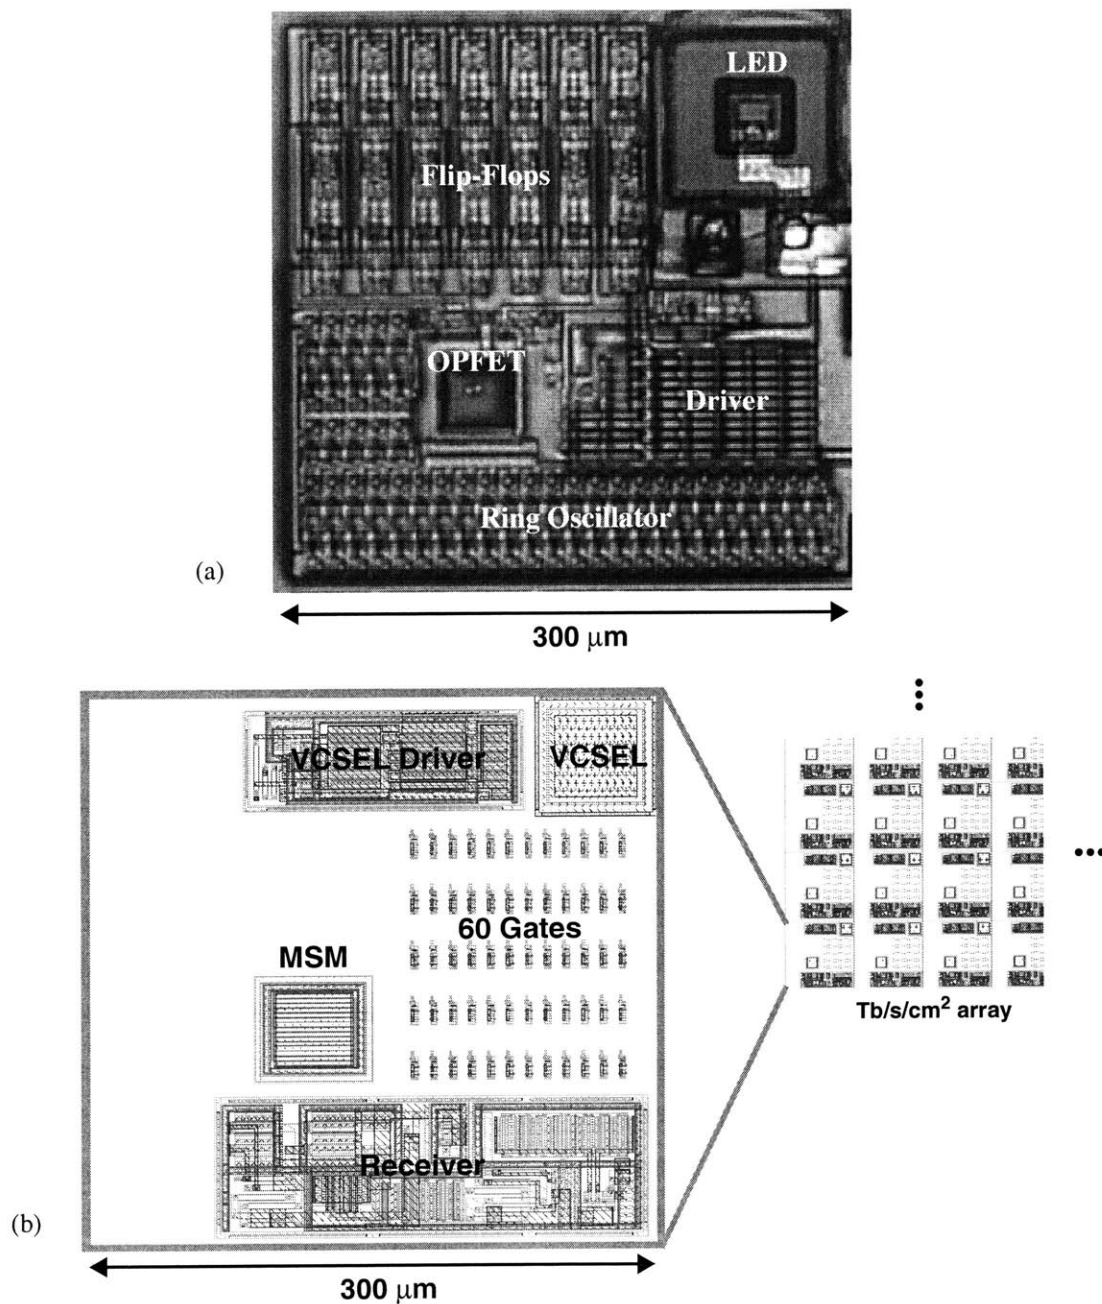


Figure 5.1: EoE Tb/s smart pixel OEICs

(a) A photomicrograph of a 300 μm x 300 μm demonstration circuit consisting of an LED and driver circuit, OPFET-based receiver, and complex digital electronics. (b) shows an equivalent cell area containing 1 Gb/s transmit and receive components and 60 DCFL logic gates. These building blocks satisfy the performance, power, and area constraints needed to build a 1 Tb/s/cm² smart pixel array.

with 60 DCFL logic gates. The interface circuits in this figure are the actual designs implemented in Chapter 6 and Chapter 7.

Parallel fiber-optic data links

Chapter 1 described the commercial parallel optical interconnects which have been implemented by packaging linear VCSEL arrays, detector arrays, and electronic ICs within hybrid modules. The role of a monolithic OEIC in this market would be to reduce cost and increase yield and reliability by simplifying package requirements, reducing part count and inventory overhead, reducing the number of critical alignments (since the integrated devices are automatically aligned to each other), and taking advantage of on-wafer testing of complete transceivers [105].

Figure 5.2 shows a mock-up of a monolithic ten channel bi-directional parallel optical transceiver OEIC. VCSELs using 75 μm square DGWs and 75 μm diameter MSM photodetectors have been interleaved on a 250 μm pitch to match that of standard multimode fiber ribbon. The laser driver and receiver circuits of Chapter 6 and Chapter 7 have been included, along with I/O drivers supporting differential ECL inputs and outputs¹.

Summary

The power consumption and area constraints of the parallel optical interconnect are not as stringent as those of the smart pixel array. The constraints derived from the smart pixel example are thus used to guide the design of the interface circuits of Chapter 6 and Chapter 7. This sets a 5 mW power consumption target for the transmitter (laser driver plus VCSEL combination) and for the receiver. The choice of assigning equal levels of power consumption to each components is somewhat arbitrary, but it does serve as a good starting point. Also, a specific area constraint has not been established. The designs should attempt to minimize area after meeting power and performance requirements. With regard to the latter, a 1 Gb/s data rate target is appropriate for both the smart pixel and the parallel link applications. Additional performance guidelines are determined later in this chapter.

1. Using the laser driver and receiver circuits developed in Chapter 6 and Chapter 7, the transceiver OEIC dissipates 90 mW per transmit/receive pair. Of this, 80 mW are consumed by the ECL I/O drivers. Although the monolithic implementation already dissipates less than half the power of a typical hybrid unit, further system-level power efficiencies may be achieved in the future by incorporating higher levels of functionality onto the OEIC.

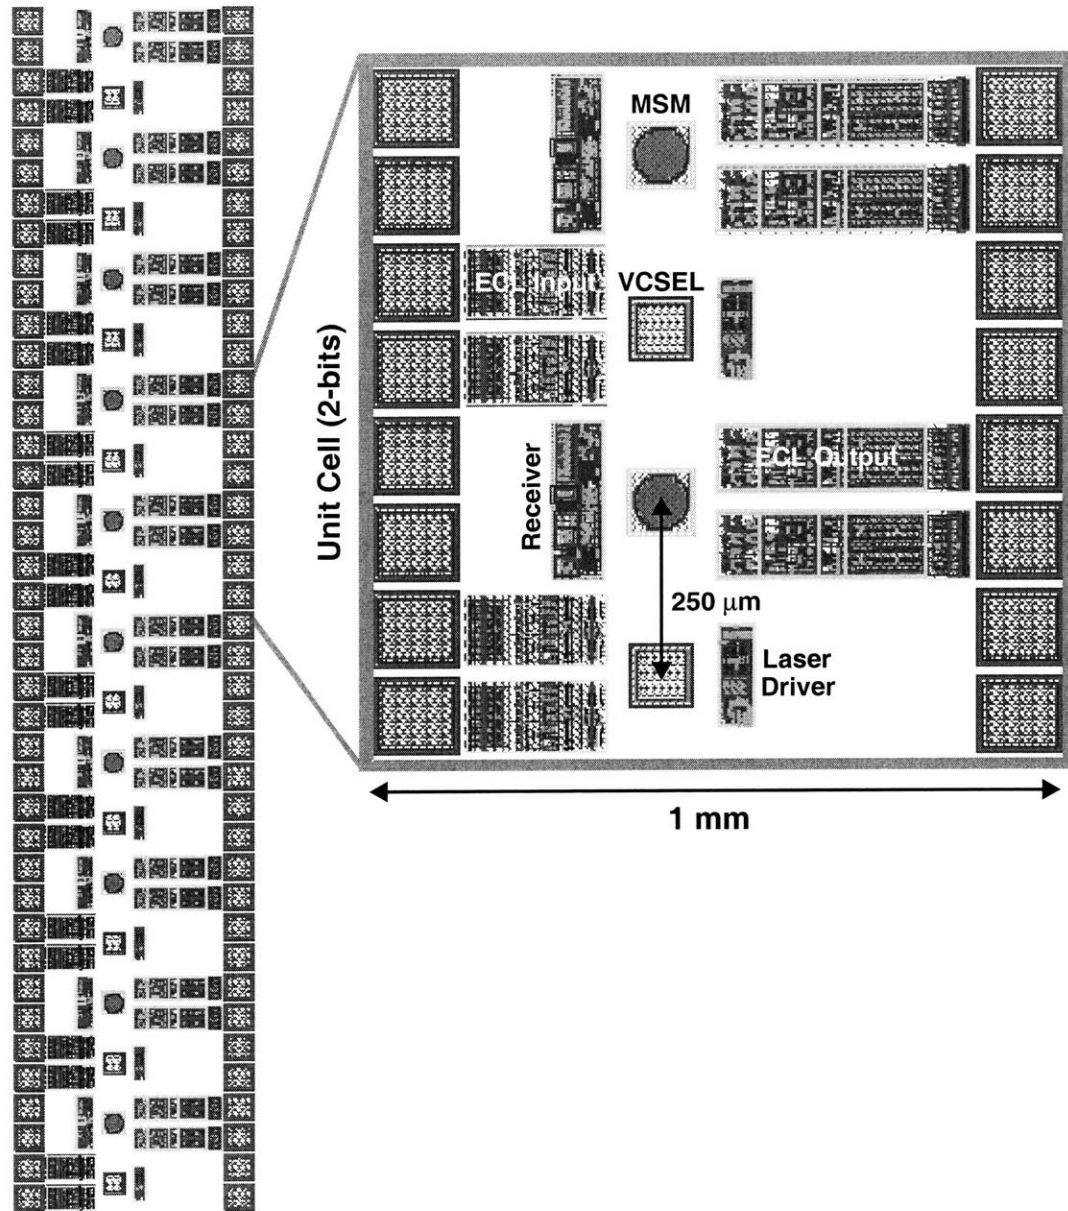


Figure 5.2: EoE parallel fiber-optic data link OEICs

This mock-up shows how EoE may be used to implement a monolithic parallel optical transceiver. Ten pairs of VCSELs in $75\ \mu\text{m}$ DGWs and $75\ \mu\text{m}$ diameter MSM photodetectors have been arrayed on a $250\ \mu\text{m}$ pitch matching that of a multimode optical fiber ribbon. Differential ECL I/O are provided for each channel. To balance the number of power and ground pads relative to the number of I/Os, a $0.5\ \text{mm}^2$ unit cell is composed of two transmit/receive pairs. Each transmit/receiver pair dissipates $90\ \text{mW}$ of which $80\ \text{mW}$ is due to the ECL I/O. This is $<45\%$ of the commercial unit shown in Figure 1.3. The entire 10-bit array occupies $5\ \text{mm}^2$ and dissipates $900\ \text{mW}$. This translates into a power density of $18\ \text{W}/\text{cm}^2$, allowing the use of low cost thermally-enhanced plastic packaging.

5.2 Characteristics of Optical Link Components

In addition to the interface circuits, the essential ingredients of the digital optical interconnects considered in this thesis are the vertical-cavity surface-emitting laser (VCSEL), the MSM photodetector, and an optical signal path. This section briefly summarizes typical characteristics of each of these elements.

Operating wavelength

Most optical interconnect systems are implemented at a wavelength of 850 nm for compatibility with low cost GaAs VCSELs and Si¹ or GaAs photodetectors. The standard optical communication wavelengths of 1.3 μm and 1.55 μm , which minimize loss and dispersion in single mode glass optical fiber, are not desirable in optical interconnects because of the high cost of the InP-based emitters and detectors that are needed.

The 850 nm VCSEL

VCSELs have benefited from an intensive development effort in the last decade, and are now finding widespread use in optical interconnect systems. They are compact, compatible with standard IC processing techniques, on-wafer testable, and their output is efficiently coupled into both fiber and free-space optical systems. They may be operated over a wide temperature range (e.g. -50° - 150°C) with a low threshold current (typically below 1 mA, and as low as 40 μA) and high power conversion efficiency (as high as 57% at an 850 nm output power level of around 1 mW) while simultaneously achieving a high modulation bandwidth (typically several GHz, and as high as 21 GHz) [85,117,172-187]. Figure 5.3 shows two possible implementations of an EoE-integrated VCSEL.

Measured characteristics of a typical high performance VCSEL are given in Chapter 6 in conjunction with the characterization of the transmitter interface circuit (laser driver). Figure 5.4 summarizes a number of key device parameters around which the interface circuits are based. The laser driver needs to switch the laser between two operating points corresponding to LO and HI logic states. The LO state may be just above the lasing threshold while the HI state may correspond to an output power of around 1 mW. For use

1. The Si photodetectors that have been commonly used are discrete pin diodes. Integration of photodetectors into Si VLSI electronics processes has been difficult because of the long optical absorption length ($\sim 10 \mu\text{m}$ for 850 nm light), though some promising approaches have been reported [see Ref. 171 and the citations therein].

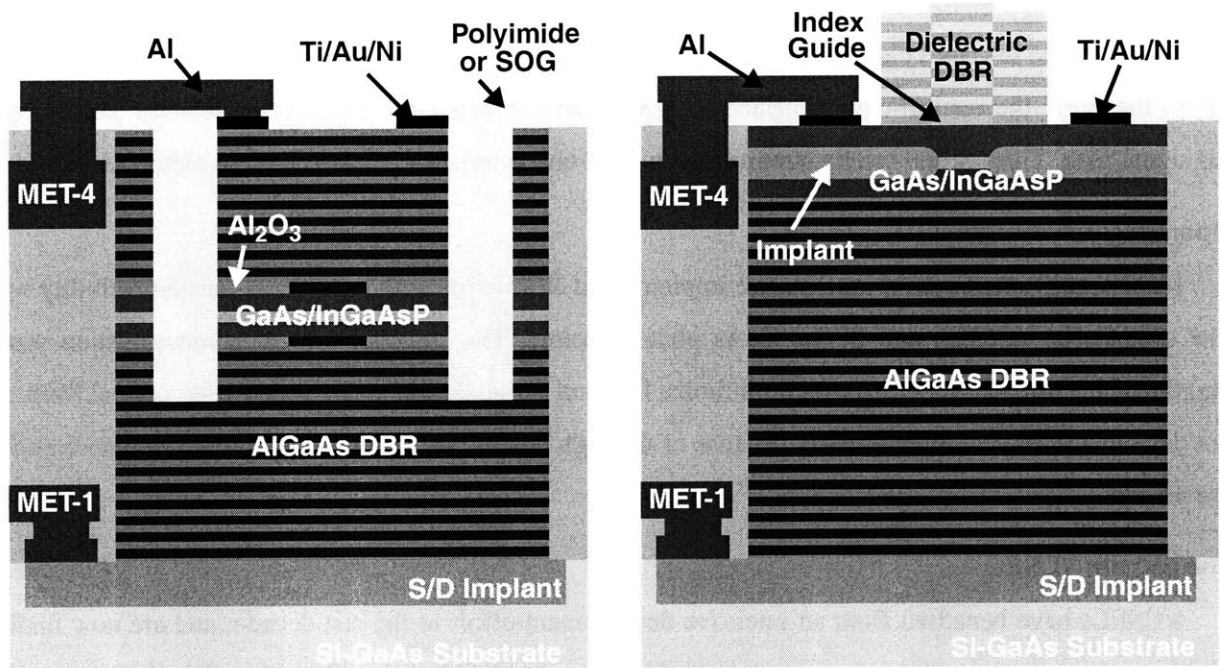


Figure 5.3: Possible integrated VCSEL structures

A vertical-cavity surface-emitting laser (VCSEL) is a pin diode structure in which the i-region (which may be lightly p-doped) forms a short optical cavity (a small integer multiple of half the optical wavelength in the material, $k \cdot \lambda / 2n$) between distributed Bragg reflectors (DBRs) consisting of a stack of quarter-wavelength-thick layers of alternating dielectric constant material. A number of quantum wells placed at the center of the cavity produce optical gain. To achieve a low threshold current and high power conversion efficiency, overlapping lateral confinement of both the injected current and of the optical mode are required. The structure in (a), which is modeled after Lear, *et. al.* [85, 178-187], simultaneously forms a current aperture and index guide by selectively oxidizing layers of $\text{Al}_x\text{Ga}_{1-x}\text{As}$ ($x \sim 0.98-1$). The structure in (b), introduced by Chirovsky, *et. al.* [117], forms a current aperture using a buried fluorine implant which, following a sub-500°C anneal, increases the resistivity of the material. A separate index guide is created by patterning a layer of material just above the current aperture. Both structures achieve excellent results, with (a) currently delivering the highest performance and (b) favored from a manufacturing perspective because of its planarity and lack of a difficult-to-grow p-type DBR. Current implementations of both structures use the AlGaAs material system throughout. The high index contrast provided by this material system is important to the success of the DBRs, but Al-containing active layers are not desirable in the EoE context. To produce an integrated device, the Al-free laser active layers demonstrated in Chapter 2 may be combined with AlGaAs DBRs, as shown above. The incorporation of non-radiative recombination centers in the reduced-temperature-grown AlGaAs should not effect the device since there is no significant minority carrier population in the DBRs (i.e. they are only used to conduct current to the active region). EoE compatible AlGaAs DBRs grown by the author and by others [2,9] have confirmed that the reflectivity and morphology of the DBRs is not degraded by the reduced growth temperature. High DBR conductivity remains to be demonstrated, but is not expected to depend significantly on growth temperature. If DBR conductivity is found to be a problem, (b) may be modified to use top-side contacts for both the n- and p-connections, eliminating DBR conduction altogether (this, in fact, was Chirovsky's original design).




Figure 5.4: L-I characteristics of a VCSEL

A generic VCSEL light-versus-current (L-I) curve is shown above. It has a lasing threshold current of below 1 mA and an output power of around 1 mW.

in experimental systems, the laser driver should have adjustable LO and modulation (HI-LO) current levels of a few milliamperes each.

The integrated MSM photodetector

The most promising detector structures for use in EoE OEIC is the metal-semiconductor-metal (MSM) photodetectors, as discussed in Chapter 4¹. For the purpose of the current discussion, the various non-idealities of the existing detector structure may be set aside. The MSM detector is characterized, for the most part, by its efficiency and capacitance. At 850 nm, it is reasonable to assume an efficiency of 0.2 A/W. The capacitance of an MSM sized for use with a multimode fiber (~75 μm diameter) is likely to be around 100 fF [148]². In a free space interconnect architecture, the photodetectors may be made smaller (~40 μm diameter) resulting in a lower capacitance. The larger capacitance associated with the multimode fiber system will be assumed in this thesis.

-
1. Another possibility is to modify an integrated VCSEL for use as a resonance-enhanced pin photodetector [167,168]. Additional device-level research is required to make this approach practical.
 2. The fiber core diameter of 62.5 μm is common. A detector diameter of 75 μm is used in this thesis to allow for easy alignment.
-

The optical system

Detailed knowledge of the optical system has not been found to be necessary in the designs presented in this work. The primary performance parameter is the coupling efficiency of the system, that is, the fraction of the VCSEL output power that is collected by the photodetector. The coupling efficiency of the index-guided VCSEL discussed above into multimode fiber is typically 50-75% [188,189]. For a free-space optical system, a coupling efficiency of around 10% is often assumed [190,191].

Summary

For applicability in a wide variety of experimental optical interconnect systems, the laser driver should support several milliamperes of off state and modulation current. In calculating the transmitter (laser driver + VCSEL) power consumption, a threshold of 1 mA and a power efficiency of 20-40% may be assumed for the VCSEL. The receiver should operate with an assumed detector capacitance of 100 fF and should produce valid logic output levels, at 1 Gb/s, in response to an input current amplitude of roughly 10-50 μ A.


5.3 Digital Interconnects and Communication Channels

Optical interconnects are often advocated as an extension of lightwave communication technology to shorter length scales. Optical communication has revolutionized telephony and data networking by taking advantage of the low loss and dispersion of optical fiber. These motivations do not necessarily apply at the shorter length scales targeted by optical interconnects. Rather, a number of practical motivations discussed in Chapter 1 may justify the use of optics at various levels of the interconnect hierarchy. Similarly, the lightwave communication analogy does not recognize different system-level requirements of digital interconnects. These system-level requirements are the topic of the current section.


Figure 5.5 shows simplified diagrams of communication and interconnect systems. A communication system is conceptually a discrete-time entity. Data symbols modulate pulses which, after suitable conditioning, are sampled in order to recover the transmitted data. The timing for the sampling process is recovered from the pulse stream itself [192-198]. By contrast, a digital interconnect is a continuous time system and may be used to transmit either data or timing.

The signal fidelity of both communication channels and interconnects may be represented by an “eye diagram”. The eye diagram is a superimposed plot of the link’s response to multiple, random bit sequences. Figure 5.6(a) is a simulated eye diagram for a typical binary communication channel taken just prior to the sampling stage. A generalization of $\text{sinc}(t)=\sin(t)/t$ known as the “raised cosine” pulse shape has been used to represent an optical channel with 1 GHz bandwidth transmitting data at 1 Gb/s [192]. The signal in Figure 5.6(a) would be sampled at the time instants 0 ns, 1 ns, 2 ns, ..., at which points it takes distinct values corresponding to high and low bits. The vertical eye opening relates to how well the receiver’s decision circuit will be able to distinguish the two input values in the presence of noise.

The design of a communication link must take two important factors into account: intersymbol interference (ISI) and signal-to-noise ratio (SNR). ISI refers to the distortion during a given bit interval resulting from the superposition of the responses from all other bit intervals. That is, the sample taken at time nT in Figure 5.6(a) includes an ISI contribution of $\sum_{k \neq n} A_k p_{rx}(nT - kT)$, where $p_{rx}(t)$ is the received, filtered pulse shape just prior to the sampling stage. The sinc-like raised-cosine pulse shapes pass through zero at multiples of the bit-time (i.e. $p_{rx}(kT)=0$ for $k \neq 0$) so that in an ideal system using these pulse shapes there is no ISI at the sample points. In principle, such pulse shapes can be constructed with a bandwidth of one-half



(a) Simple communication link using



(b) Digital interconnects carrying data and clock signals

Figure 5.5: Comparison of communication links and digital interconnects

Basic structure of a communication link (a), and a digital interconnect (b). In (a), the transmitted signal $s(t)$ is a train of pulses whose amplitudes reflect the data symbols A_k . The receive filter corrects for distortion in the channel, and the automatic gain control (AGC) function corrects for variable loss in the channel. The timing recovery circuit extracts the time instants at which the received signal is to be sampled to find the data symbols. Finally, the decision circuit estimates data symbols by comparing samples of the received signal to a decision level. At the conceptual level, such a communication link is a discrete-time entity. The digital interconnect, on the other hand, is a continuous time system. In (b), data and timing are carried on separate interconnects and the instantaneous value of each signal has significance. The transmitter (Tx) and receiver (Rx) interface between the logic circuit environment and the transmission medium.

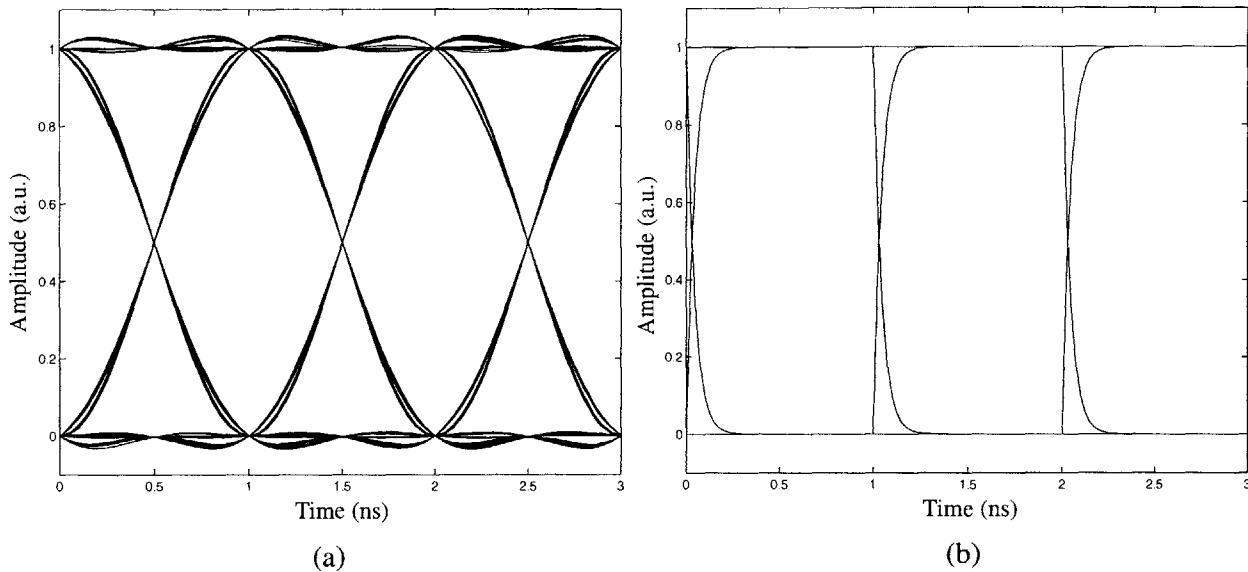


Figure 5.6: Simulated eye diagrams

Eye diagrams are formed by superimposing waveforms corresponding to the link's response to multiple, random bit sequences. (a) Represents a typical optical communication channel transmitting 1 Gb/s through 1 GHz of bandwidth using a raised cosine pulse shape with 100% excess bandwidth [192]. (b) Is a pattern for an interconnect within a digital system operating with a 1 GHz clock. This signal is modeled as a single-pole response with a bandwidth of 3.7 GHz. See Appendix D for more detailed examples.

the bit rate (i.e. 1 Gb/s transmission requires a bandwidth of 500 MHz). In real systems, additional bandwidth is required to allow for error margins in the timing recovery and sampling functions. In general, ISI is improved by increasing the system bandwidth.

The SNR, on the other hand, degrades with increasing bandwidth. That is, the narrower the bandwidth of the system, the more noise is filtered out of the signal before it is sampled. In an ISI-free system, the SNR of the sampled values entirely determines the probability that a bit-error is made. Thus, subject to ISI-related constraints, the system bandwidth should be minimized.

Now compare the communication system eye diagram with that in Figure 5.6(b) which represents the output of the receiver on the data path in Figure 5.5(b). The underlying pulse shape in this example is based on an elementary RC response, as would be found in on-chip wiring¹. Unlike the gradual, “just-in-

1. A bandwidth of 3.7 GHz has been used in order that the signal reach 90% of its final value in 0.1 ns, 10% of the 1 ns period of a 1 GHz clock. This type of responses is required at the output of the receiver, but may not exist internal to the receiver (see Section 5.4).

time” transitions of the communication system, the digital transitions are sharp, allowing them to accurately clock digital memory stages. And, while the communication signal only takes on valid values near the sampling instants, the digital interconnect signal maintains its logic value for the duration of the bit period following each transition.

Aside from the issues discussed above, important differences between communication channels and interconnects include bit-error rate requirements, the characterization of noise immunity, and delay and delay variation constraints. In practical implementations of interconnects at different length scales and in different applications, aspects of both types of links will be important¹. In the current discussion, however, an “interconnect” refers specifically to a link within a digital system carrying data or timing.

Bit-error rate

The required bit-error-rate (BER), i.e. the probability that a transmitted data bit is received incorrectly, is an important difference between communication channels and interconnects². A typical BER in an optical communication link running at several Gb/s is 10^{-9} . If the channel is supporting a voice or video session, this BER is practically undetectable by the user. If computer data is being transmitted, then the network protocol will request retransmission of corrupt data packets. For a 64 Kbyte packet, the maximum packet length used in the current internet protocol, a 10^{-9} BER results in retransmission of 1 out of 1906 packets, resulting in an overhead of only 0.05%. Thus 10^{-9} BER is quite acceptable in a common communication systems. Consider, on the other hand, an optical interconnect in a computer system operating at 1 GHz. A BER of 10^{-9} in this case means that the computer would crash after 1 second of operation. For the system to operate for one week the BER must be around 10^{-15} . Error-free operation for one year at 1 Gb/s would require a BER below 10^{-17} . The BER specification on an interconnect product is usually determined by the duration of time for which it has been tested.

1. For instance, if the information being carried by an interconnect originated from an error-prone source, such as a communication channel, then the interconnect’s bit-error rate may not need to significantly exceed that of the data stream. This differs from the general case in which the interconnect is expected to operate with essentially no error.

2. The term bit-error-rate (BER) is not really appropriate for an interconnect system. It would be more accurate to talk about the probability that a signal is not in the correct logic state at any given instant in time. Since this probability is analytically equivalent to the BER, however, the term will be applied to both types of systems in this thesis.

Signal and noise

The dramatic difference in BER for communication versus interconnect application leads to different analytical perspectives in the design of each type of system. In a communication channel, the probability of a bit-error is related to the SNR. The link is designed, sometimes including error-correction coding, to maintain an SNR consistent with the target BER. A similar probabilistic viewpoint is also important in understanding digital interconnects at a fundamental level. However, specifying interconnect performance in terms of SNR is awkward. The question in interconnect design is not “what SNR is required to achieve a given BER”, but rather “what is the minimum signal level and the maximum noise level which allow for error-free operation”. Consequently, digital connections are specified by a range of high and low signal levels and noise margins. The design of the interconnect then aims at limiting the degradation of the signal levels below the minimal specified values and at limiting noise susceptibility.

Delay and jitter

Aside from allowing adequate margins for timing recovery and sampling, time delay and timing accuracy (i.e. delay variation, or jitter) are not significant factors in a communication link. They are very important in digital interconnects. As discussed, the sample timing in a communication receiver is derived from the transmitted signal itself. As long as the clock recovery circuit can track timing variations in the received signal, the actual delay and jitter are not important. Digital systems, however, require synchronous interaction of multiple signal paths. Interconnect delay may directly contribute to the required cycle time¹. Likewise, uncertainty in the interconnect delay resulting from process and temperature variations and from data pattern dependencies must also be built into the timing margins. Thus, controlling both delay and delay variation are key elements in the design of a digital interconnect.

Summary

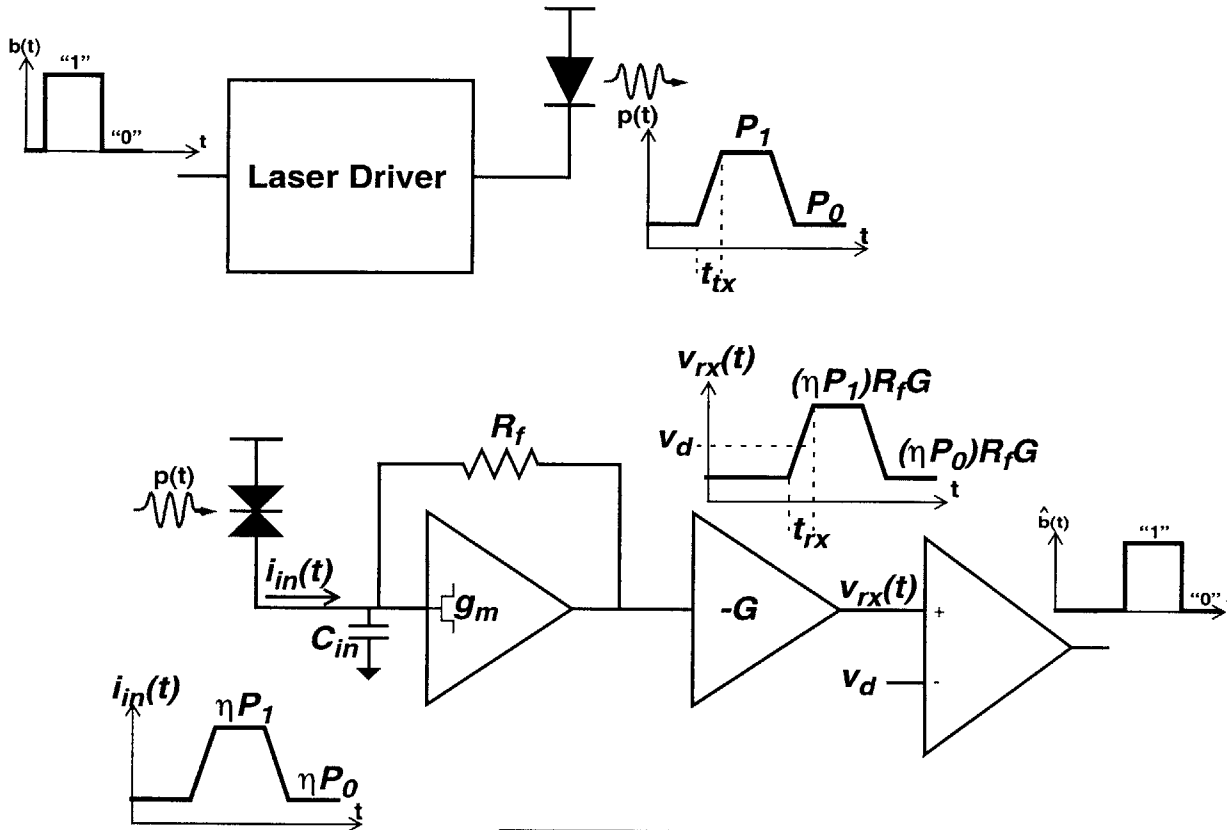
In a communication system, the goal is to send a serial bit stream over a great distance with a moderate error rate. In a digital interconnect system, the goal is to send many continuous-time binary signals over a relatively short distance with no errors. How these goals influence the design of the driver and receiver circuitry in digital optical interconnect is taken up in the next section.

1. In some systems, long delays, or latencies, are provided for by the use of pipelining. In these cases, it is the delay variation that is critical.

5.4 Digital Optical Interconnect Design

This section examines some of the issues encountered in the design of a digital optical interconnect. This process is analogous to the design of a logic family in that device-level consideration (e.g. laser threshold current and efficiency, coupling efficiency, detector efficiency and capacitance, ...), circuit topologies, and signal level definitions (e.g. minimum high output, maximum low output, maximum low input, minimum high input, ...) must be jointly optimized to meet a desired set of performance goals (e.g. delay, delay variation, power consumption, ...). It would be premature at this stage in the development of optical interconnect technology to attempt this level of specialization. It is useful, however, to clarify the relationships among some key design parameters and performance metrics. This section develops a analytical framework which models these relationships. While capturing the essence of the operation of the digital optical interconnect, this development limits the extent of the quantitative detail that is invoked so that the resulting model may be used to guide analysis and design at an intuitive level. The principle that are outlined here may also be used in a more rigorous, numerical context.

Figure 5.7 describes the make-up of a prototypical digital optical interconnect. A continuous time logic signal appears at the input of a laser driver which switches the current through a laser between two prescribed levels. The light output is coupled into an optical system (e.g. fiber, planar waveguide, free space, ...) and arrives at a photodetector with power levels P_0 and P_1 corresponding to logic low and high levels, respectively. The time for the optical signal to transition between these power levels is defined as t_{lx} . The photodetector, with efficiency η (A/W) and capacitance C_{det} , produces a photocurrent i_{in} which is input to a transimpedance amplifier (TIA) with feedback resistance R_f . The input FET of the TIA has a transconductance g_m and gate-to-source and gate-to-drain capacitances C_{gs} and C_{gd} , respectively. The total capacitance at the input of the TIA, for the purpose of noise calculation, is $C_{in} = C_{det} + C_{gs} + C_{gd}$. The TIA is followed by one or more additional gain stages ("post-amplifiers") having a total gain of G . The cascaded linear system consisting of the detector/TIA and the post amplifier has a bandwidth of B . The output of this "linear receiver", v_{rx} , switches between voltages $\eta P_0 R_f G$ and $\eta P_1 R_f G$ in a time t_{rx} . This signal is applied to a comparator which produces a logic high output if v_{rx} is greater than the "decision level", v_d , and a logic low output if v_{rx} is less than v_d .



t_{tx}	Transmitter rise/fall time
t_{rx}	Receiver (linear) rise/fall time
η	Photodetector efficiency (A/W)
$M=P_0/P_1$	Extinction ratio
R_f	Transimpedance amplifier feedback resistor
G	Post-amplifier gain
C_{in}	Total input capacitance ($C_{det} + C_{gs} + C_{gd}$)
g_m	Transconductance of input FET
B	Bandwidth of (linear) receiver
v_d	Decision level of comparator

Figure 5.7: Anatomy of a digital optical interconnect

In this prototypical digital optical interconnect, a continuous time logic signal appears at the input of a laser driver which switches the current through a laser between two prescribed levels. The light output is coupled into an optical system and arrives at a photodetector which delivers a photocurrent, i_{in} , to a transimpedance amplifier (TIA). The TIA is followed by one or more additional gain stages. The output of this "linear receiver", v_{rx} , is applied to a comparator along with the "decision level", v_d , to produce a logic output.




Figure 5.8: Example of timing and logic errors caused by signal distortion

Noise and ISI distort the signal which appears at the input of the comparator in Figure 5.7, resulting in errors in the logic output and in the output switching time. The latter is referred to as jitter.

The signal applied to the comparator, v_{rx} , is distorted by ISI and additive noise. Figure 5.8 gives an example of how this distortion can lead to errors in the output logic value and the output timing. Since, by construction, the output of the receiver (i.e. output of the comparator) is binary valued, the design of the transmitter, optical channel, and receiver are driven by two principles: **(I) OUTPUT THE CORRECT LOGIC VALUE**, and **(II) CHANGE THE OUTPUT AT THE CORRECT TIME**.

The optical channel design is not considered in this thesis. With regard to the transmitter (laser driver) and receiver, other design goals include power consumption, delay, and area. These may appear as constraints on the optimization of **I** and **II** above, and are of great practical importance in determining the realizable range of the variables appearing in Figure 5.7. The goal of this section is to quantify how these variables interact to determine **I** and **II**.

Gaussian noise and error probability

As shown in Figure 5.8, the presence of noise on v_{rx} as it ramps past the decision level of the comparator causes a random displacement in the time at which the output switches. There will thus be some uncertainty in the receiver output around the time of the switching event. The time duration of this jitter interval will be discussed below. At a time instant, t , outside of the jitter interval, the receiver is expected to produce a static high or low output depending on the value of $v_{rx}(t)$ relative the decision level, v_d . More specifically, the comparator output should indicate whether the mean value of the comparator input, $\bar{v}_{rx}(t)$, is greater or less than v_d . If the noise on v_{rx} is assumed to be Gaussian, the probability that the comparator produces the wrong value at a given instant in time is given by [193]:

$$Pr\{Error\}_t = Q\left(\frac{\bar{v}_{rx}(t) - v_d}{\sqrt{\langle v_{rx}^2 \rangle}}\right) \quad (5.1)$$

where the Q function is based on an integral of the normal distribution:

$$Q(x) \equiv \frac{1}{\sqrt{2\pi}} \int_x^{\infty} e^{-x^2/2} dx = \frac{1}{2} \operatorname{erfc}\left(\frac{x}{\sqrt{2}}\right) \quad (5.2)$$

If the noise variance, $\langle v_{rx}^2 \rangle$, is independent of \bar{v}_{rx} , then the optimal location of the decision level is midway between the high and low mean values. This is not true in an optical system since the noise variance associated with the optical signal is proportional to the optical power. However, taking the optical noise to be constant at its peak value represents a worst-case analysis and is appropriate here [194]. Furthermore, as a practical matter, v_d needs to be set near the mid-point of the transition in v_{rx} in order to produce uniform rising and falling delays through the interconnect. Assuming constant noise variance and placing v_d midway between the high and low value of \bar{v}_{rx} allows (5.1) to be written as

$$Pr\{Error\} = Q\left(\frac{\sqrt{SNR}}{2}\right) \quad (5.3)$$

With the signal-to-noise ratio, SNR , defined as the ratio of the signal power (squared amplitude) to the noise power (variance):

$$SNR = \frac{(\overline{v_{rx,high}} - \overline{v_{rx,low}})^2}{\langle v_{rx}^2 \rangle} \quad (5.4)$$




Figure 5.9: Error probability vs. SNR

Assuming the noise on v_{rx} is Gaussian and independent of time, and that the decision level is set midway between the high and low signal mean values, the probability that the output of the comparator in Figure 5.7 is incorrect at a time instant outside a transition interval is determined from the signal-to-noise ratio by (5.3).

Figure 5.9 plots the error probability as a function of SNR. Around 25 dB SNR is needed to meet the “vanishingly small” error probability requirement of 10^{-15} - 10^{-20} discussed above. At this SNR, the amplitude of v_{rx} is around 18 times the standard deviation of the noise. As a rule-of-thumb for signals with Gaussian noise, the “peak-to-peak” noise amplitude may be related to the standard deviation by a factor, α :

$$(Noise)_{peak-to-peak} \sim \alpha \cdot \sigma_{Noise} \quad (5.5)$$

A factor of $\alpha=20$ is convenient for analytical use in this section and corresponds to SNR=26 dB or an error probability of $\sim 10^{-20}$. An empirically determined value of $\alpha=10$ (SNR=20 dB, error probability $\sim 10^{-7}$) is more appropriate for use with laboratory measurements and reflects the limited duration of the measurement.

Receiver SNR and error probability

It is convenient to express the noise appearing in v_{rx} as an equivalent noise current at the receiver input.

$$\langle i_{in}^2 \rangle = 2q(\eta P_1)B + \frac{4kT}{R_f}B + 4kT\left(\frac{4\pi^2\Gamma}{3}\right)\frac{C_{in}^2}{g_m}B^3 \quad (5.6)$$

The first term in this expression is the (worst case) shot noise due to the optical signal. The remaining terms are due to the receiver and are discussed in Chapter 7 (page 358). The two dominant receiver noise sources appear as the second and third terms of (5.6). The second term is the thermal noise due to the TIA feedback resistor and the third term is the thermal noise of TIA's FET input. Γ is the FET noise factor and has a value of 1.8 for GaAs and 0.6 for Si.

The noise in v_{rx} can now be written in terms of the input noise as

$$\langle v_{rx}^2 \rangle = \langle i_{in}^2 \rangle (R_f G)^2 \quad (5.7)$$

The factor $(R_f G)$ is the gain of the linear segment of the receiver.

As indicated in Figure 5.7, the high and low mean values of v_{rx} are $(\eta P_1)R_f G$ and $(\eta P_0)R_f G$, respectively. Substituting these values into (5.4) along with (5.7) gives the *SNR* as

$$SNR = \frac{[(\eta P_1)R_f G - (\eta P_0)R_f G]^2}{\langle v_{rx}^2 \rangle} = \frac{[\eta(P_1 - P_0)]^2}{\langle i_{in}^2 \rangle} \quad (5.8)$$

Next, define the optical power amplitude, P , and the "extinction" ratio, M , as:

$$P \equiv P_1 - P_0 \quad (5.9)$$

$$M \equiv \frac{P_0}{P_1} \quad (5.10)$$

allowing the high-level optical power to be written

$$P_1 = \frac{P}{1 - M} \quad (5.11)$$

With these definitions and the noise expression in (5.6), the *SNR*, (5.8), may be written as

$$SNR = \frac{(\eta P)^2}{\frac{2q(\eta P)}{1 - M}B + \frac{4kT}{R_f}B + 4kT\left(\frac{4\pi^2\Gamma}{3}\right)\frac{C_{in}^2}{g_m}B^3} \quad (5.12)$$

This is the key expression in assessing how the various system parameters effect design goal I. Naturally, the SNR , and with it the error probability, improve as the amplitude of the optical signal, P , is increased. P may be limited by a system power dissipation constraint, or by an external factor such as eye safety. The SNR is also improved by reducing the bandwidth B , which may lead to ISI and jitter. P and B play a dominant role in (5.12). Figure 5.10 plots the effect of these two parameters on the error probability. Typical values are assumed for η , C_{in} , R_f and g_m , based on the work in Chapter 4 and Chapter 7 and are indicated in the figure; a best-case extinction ratio of $M=0$ is assumed (the effect of $M>0$ will be examined later).

The main observation from Figure 5.10 is that, for the typical system parameters chosen in this example, an optical power amplitude of around $10 \mu\text{W}$ is required in a 1 GHz system to achieve an error probability of 10^{-20} . If the amplitude is reduced much below this level, the bandwidth must be curtailed sharply and unacceptably. If the bandwidth needs to be increased, it must be accompanied by an increase in amplitude.

Aside from amplitude and bandwidth considerations, SNR is improved by minimizing the extinction ratio, M , and the input capacitance, C_{in} , and by maximizing the feedback resistance, R_f , the transconductance, g_m , and the detector efficiency, η . These variables are generally constrained by practical limitations. For instance, $M>0$ may be needed to control laser turn-on delay (see Figure 6.50), a minimum detector area, and thus input capacitance, may be needed to match a given optical system, the feedback resistor may be limited by the gain achievable in the voltage amplifier used to build the TIA, and the transconductance may be limited by power dissipation.




Figure 5.10: Error probability vs. optical power amplitude and bandwidth

The error probability, (5.3), is determined from the SNR, (5.12), calculated with the system parameters indicated in the table above. $M=0$ is assumed and typical values for the other parameters are taken from Chapter 4 and Chapter 7.

Jitter

Design goal **II** relates to the jitter in the receiver's comparator output. An expression for the jitter variance is derived in Appendix D.

$$\sigma_{jitter}^2 = \sigma_{linear}^2 + \frac{t_{rx}^2}{SNR} + \sigma_{comparator}^2 \quad (5.13)$$

The first term in this expression is the jitter produced within the linear segment of the receiver due, primarily, to ISI. This term will be examined later in this section. The last term allows for timing variations in the comparator itself (other than those related to noise which are included in the second term). As long as a very simple implementation is used for the comparator, such as a basic logic inverter, this term does not contribute significantly to the total jitter. The second and dominant jitter term results from noise at the input of the comparator. An estimate of the linear receiver rise time, t_{rx} , is needed in order to quantify this term.

There are numerous definitions of rise/fall time, as summarized in Table 5.1 [95]. One mathematically useful definition is the standard deviation of the impulse response, T_{σ} . When linear systems are cascaded, their impulse responses are convolved and the impulse responses variances, $(T_{\sigma})^2$, add. As seen in Table 5.1, all of the common definitions of rise/fall time are fairly consistent. As a result, regardless of the specific definition used, the transition time (i.e. step-response rise/fall time) of a cascade of systems may be expressed in terms of its constituents as:

Impulse Response Type	T_{σ}	T_{10-90}	T_{20-80}	T_{slope}	T_{FWHM}	F_{3dB}
One-Pole	1.00	0.877	0.553	0.798	0.28	0.399
Two-Pole Critically Damped	1.00	0.947	0.612	0.900	0.69	0.363
Gaussian	1.00	1.02	0.672	1.00	0.94	0.332

Table 5.1: Comparison of rise time definitions and corresponding bandwidths for various linear systems

T_{σ} is the standard deviation of the system's impulse response, T_{10-90} is the 10%-to-90% transition time, T_{20-80} is the 20%-to-80% transition time, T_{slope} is the transition time based on extrapolating the slope of the signal at the midpoint of its transition, and T_{FWHM} is the full-width-at-half-maximum of the impulse response. In each case, the product of the transition time and the bandwidth, F_{3dB} , is a constant depending on the type of system. Furthermore, this constant does not vary greatly among the various system types [95].

$$(t_{cascade})^2 = (t_1)^2 + \dots + (t_n)^2 \quad (5.14)$$

Furthermore, the product of the transition time of a linear system and its bandwidth is a constant which depends only weakly on the details of the system.

$$(t_{rise})(f_{3dB}) = \gamma \quad (5.15)$$

As seen in Table 5.1, the constant, γ , typically takes values in the range 0.3-0.4.

Using (5.14) and (5.15), the transition time of the linear receiver output may be estimated as:

$$(t_{rx})^2 = (t_{ix})^2 + \left(\frac{\gamma}{B}\right)^2 = \gamma^2 \left(1 + \frac{t_{ix}^2 B^2}{\gamma^2}\right) \frac{1}{B^2} \quad (5.16)$$

Substituting this expression and the SNR , (5.12), into the middle term of (5.13) gives the jitter variance resulting from noise at the comparator input:

$$\sigma_{jitter}^2 = \gamma^2 \left(\frac{2q}{(1-M)(\eta P)B} + \frac{4kT}{R_f} \frac{1}{(\eta P)^2 B} + 4kT \left(\frac{4\pi^2 \Gamma}{3} \right) \frac{C_{in}^2}{g_m (\eta P)^2} \frac{B}{g_m (\eta P)^2} \right) \left(1 + \frac{t_{ix}^2 B^2}{\gamma^2} \right) \quad (5.17)$$

This is the key expression in assessing how the various system parameters effect design goal **II**. As with the SNR , it is desirable to minimize the extinction ratio, M , and the input capacitance, C_{in} , while maximizing the feedback resistance, R_f , the transconductance, g_m , and the detector efficiency, η . Likewise, the optical power amplitude, P , plays a dominant role in reducing jitter. However, unlike the SNR , there is an optimum bandwidth which, for a given set of system parameters, minimizes the jitter.

The jitter, (5.17), is a product of two factors. The first contains the various noise sources and gives the output jitter in the idealized case in which the optical input makes instantaneous transitions. It initially decreases with increasing bandwidth since the resulting drop in t_{rx} is more significant the increase in noise. Once, the FET channel noise term becomes dominant, however, the jitter begins to rise. The optimum bandwidth, in the idealized case of $t_{ix}=0$, is:

$$B_{optimum}|_{t_{ix}=0} = \left[\frac{\frac{2q}{(1-M)(\eta P)} + \frac{4kT}{R_f} \frac{1}{(\eta P)^2}}{4kT \left(\frac{4\pi^2 \Gamma}{3} \right) \frac{C_{in}^2}{g_m (\eta P)^2}} \right]^{\frac{1}{2}} = \left[\frac{\frac{2q(\eta P)}{(1-M)} + \frac{4kT}{R_f}}{4kT \left(\frac{4\pi^2 \Gamma}{3} \right) \frac{C_{in}^2}{g_m}} \right]^{\frac{1}{2}} \quad (5.18)$$

The second factor in (5.17) adjusts the jitter to reflect the finite transition time of the optical signal. Essentially, the second factor accounts for the fact that increasing the bandwidth beyond that of the input signal adds to the noise without reducing t_{rx} , and thus degrades the jitter.

The peak-to-peak jitter is written as

$$(Jitter)_{peak-to-peak} = \alpha \cdot \sigma_{jitter} \quad (5.19)$$

using the relationship pointed out in (5.5) and is plotted as a function of bandwidth in Figure 5.11 with $\alpha=20$. The same typical system parameter of Figure 5.10 are used along with $\gamma=0.35$. When t_{tx} is non-zero it dominates the bandwidth optimization problem, with the optimum bandwidth occurring somewhat below γ/t_{tx} . The jitter is reduced dramatically as the amplitude of the optical signal is increased above the 10 μW level which gave an acceptable static sensitivity result. This points out the important observation that the jitter is a much more sensitive measure of the interconnect performance than is the error probability. This tendency should be accounted for when undertaking laboratory measurements of these systems.

The dependence of jitter on the optical signal amplitude and switching time are shown in more detail in Figure 5.12 and Figure 5.13, respectively. As seen in these figures, the roles of P , B , and t_{tx} are intertwined in determining the jitter. The importance of optical noise relative to thermal noise plays an important role in this interaction, with thermal noise important at low optical amplitudes and optical noise important at higher amplitudes. The thermal noise is determined by the detector and FET capacitances, which are included in C_{in} , the FET transconductance, g_m , and the feedback resistance, R_f .

The dependence of jitter on R_f and g_m are shown in Figure 5.14 and Figure 5.15, respectively. These parameter are also intertwined in establishing the thermal noise level. For instance, decreasing C_{in}^2/g_m reduces the FET thermal noise and allows for further noise reduction by increasing R_f . Conversely, for a given value of C_{in}^2/g_m increasing R_f is beneficial up to a point beyond which further gains are scarce.

Note that in Figure 5.15 the input capacitance, C_{in} , is held constant under the assumption that it is dominated by the detector capacitance. In general, both C_{in} and g_m depend on the width of the input FET. As a result, there is an optimum value for g_m [158,159].

Thus far in this section, the error probability and jitter characteristics of the digital optical interconnect have been examined under the best-case scenario in which the optical signal is fully extinguished ($M=0$) and there is no supply noise or ISI. These non-idealities will be examined next.




Figure 5.11: Jitter vs. bandwidth

The peak-to-peak jitter is given by (5.17) and (5.19). It is plotted here as a function of the linear receiver bandwidth, B , for a typical set of system parameters and various combinations of optical signal amplitude and transition time. A clear minimum in the jitter appears, in this example, near $B=1$ GHz. The jitter drops sharply as the amplitude is increased from $10 \mu\text{W}$ to $50 \mu\text{W}$, and continues to drop, more gradually, as the amplitude is increased further.

α	20
γ	0.35
M	0
η	0.2 A/W
R_f	2 k Ω
C_{in}	150 fF
g_m	4 mS
Γ	1.8




Figure 5.12: Jitter vs. optical power amplitude

(a) Shows the jitter as a function of optical power amplitude, P , for various values of optical transition time, t_{tx} . At each value of P , the bandwidth has been adjusted to minimize jitter. The optimum bandwidth is plotted in (b). The optimum bandwidth curve for $t_{tx}=0$ follows (5.18). The rise in this curve starting at ~ 50 - 100 μ W indicates that optical noise becomes dominant beyond this point. Below this point, thermal noise dominates, jitter variance drops as $1/P^2$, and the peak-to-peak jitter drops as $1/P$. Above this point, jitter variance drops as $1/P$ so the curves in (b) change to a $1/P^{1/2}$ dependence.

α	20
γ	0.35
M	0
η	0.2 A/W
R_f	2 k Ω
C_{in}	150 fF
g_m	4 mS
Γ	1.8




Figure 5.13: Jitter vs. transmitter rise/fall time

(a) Shows the jitter as a function of optical transition time, t_{tx} , for various values of optical signal amplitude, P . At each value of t_{tx} , the bandwidth has been adjusted to minimize jitter. The optimum bandwidth is plotted in (b). At low amplitudes, the optimum bandwidth is relatively small, reducing the importance of t_{tx} . At higher amplitudes, the optimal bandwidth is higher because of the relatively smaller importance of FET channel noise. The higher bandwidth, in turn, gives t_{tx} a greater role in determining the jitter.

α	20
γ	0.35
M	0
η	0.2 A/W
R_f	2 k Ω
C_{in}	150 fF
g_m	4 mS
Γ	1.8




Figure 5.14: Jitter vs. transimpedance amplifier feedback resistance

(a) Shows the jitter as a function of TIA feedback resistance, R_f , for various values of optical signal amplitude, P . At each value of R_f the bandwidth has been adjusted to minimize jitter. The optimum bandwidth is plotted in (b). At low amplitudes, increasing the feedback resistor is effective in reducing jitter up to a value of around 2-3 k Ω . Beyond this, the FET channel noise becomes larger than the resistor thermal noise. At higher powers, optical noise dominates and increasing R_f beyond 1-2 k Ω is not helpful.

α	20
t_{tx}	200 ps
γ	0.35
M	0
η	0.2 A/W
C_{in}	150 fF
g_m	4 mS
Γ	1.8




Figure 5.15: Jitter vs. transimpedance amplifier FET transconductance

(a) Shows the jitter as a function of TIA input FET transconductance, g_m , for various values of optical signal amplitude, P . At each value of g_m , the bandwidth has been adjusted to minimize jitter. The optimum bandwidth is plotted in (b). While this figure focuses on g_m , it should be recalled that it is the ratio C_{in}^2/g_m which determines the impact of FET channel noise. C_{in} includes both the FET gate capacitance and the detector capacitance. In the present example, the latter is assumed to be dominant. At low amplitudes and for the selected value of R_f the drop in jitter flattens out at a g_m of around 4-5 mS. At higher amplitudes, the drop levels off somewhat sooner.

α	20
t_{tx}	200 ps
γ	0.35
M	0
η	0.2 A/W
R_f	2 k Ω
C_{in}	150 fF
Γ	1.8

Non-zero extinction ratio

Maintaining an extinction ratio of $M=0$ may not be possible in a real system because of variations in laser threshold current and the effects of laser turn-on delay. Non-zero extinction ratio leads to a greater level of optical noise in the system, and appears as a $1/(1-M)$ factor in the optical noise term of the SNR , (5.12), and jitter, (5.17). A convenient way to keep track of the effect of a non-zero extinction ratio is to use a “power penalty”.

Consider a system which, with $M=0$, operates with a specified error probability and jitter at a particular optical amplitude, P . If this system is to be used with a non-zero extinction ratio, the optical amplitude may be increased to compensate for the resulting increase error probability and jitter. Since the extinction ratio enters both of these metrics exclusively through the SNR , the new optical amplitude, P' , may be defined by equating the SNR of the original system (optical amplitude P and extinction ratio $M=0$) with that of the modified system (optical amplitude P' and extinction ratio $M>0$):

$$SNR|_{P';M>0} = SNR|_{P;M=0} \quad (5.20)$$

The power penalty may then be defined as the ratio of the new, higher optical amplitude with the original one.

$$K_M \equiv \frac{P'}{P} \quad (5.21)$$

An expression for K_M follows from the use of (5.12) in (5.20), and it depends on the extinction ratio of the new system, M , and the ratio of optical noise to the total noise present in the original system, S_P

$$S_P \equiv 2q \frac{B \cdot SNR}{\eta P} = \frac{\langle i_{in}^2 \rangle_{optical}}{\langle i_{in}^2 \rangle_{total}} \Big|_{P;M=0} \quad (5.22)$$

With this definition, the power penalty for non-zero extinction may be written as

$$K_M = \frac{S_P}{2(1-M)} + \left[\left(1 - \frac{S_P}{2(1-M)} \right)^2 + \frac{MS_P}{1-M} \right]^{\frac{1}{2}} \quad (5.23)$$

Notice that $S_P/(1-M)$ is actually the fraction on optical noise in the new system.




Figure 5.16: Power penalty for non-zero extinction ratio

The degradation in error probability and jitter which results from the use of a non-zero extinction ratio, $M=P_0/P_1$, may be compensated by increasing the optical signal amplitude by a factor K_M . This factor is referred to as the power penalty and is given by (5.23). The fraction of optical noise in the $M=0$ system determines how quickly K_M increases with M .

The non-zero extinction power penalty is plotted in Figure 5.16. The optical noise fraction, S_p , determines how quickly K_M rises with M . When the original optical amplitude is small (small S_p) a relatively small increase in the amplitude is very effective in restoring the SNR so using a non-zero extinction does not result in a significant power penalty. When the original amplitude is large enough that the optical noise is dominant (large S_p) the improvement of the SNR with increasing P is weaker. Thus, if the optical amplitude in the $M=0$ system was large to begin with, the K_M penalty may be significant (~ 3 dB for an extinction ratio as large as 40%).

Supply noise

Since digital optical interconnects are intended for use VLSI environments, supply noise may be a major issue. Although the supply noises is deterministic, it may be assumed for simplicity that, like the other noise sources, the supply noise is adequately modeled as a white, wide-sense-stationary Gaussian random process. Suppose that the transfer function connecting the supply to the linear receiver output is $H=v_{rx}/v_{supply}$, which is constant up to a certain bandwidth¹, and that the variance of the supply noise within this bandwidth is $\langle(V_{supply})^2\rangle$. Then, the noise variance on v_{rx} may then be written as

$$\langle v_{rx}^2 \rangle = \langle i_{in}^2 \rangle (R_f G)^2 + \langle v_{supply}^2 \rangle |H|^2 = \left(\langle i_{in}^2 \rangle + \frac{\langle v_{supply}^2 \rangle}{(PSRR)^2} \right) (R_f G)^2 \quad (5.24)$$

where the power supply rejection ratio, $PSRR$, has been defined as:

$$PSRR \equiv \frac{|R_f G|}{|H|} \quad (5.25)$$

The ratio $\langle(V_{supply})^2\rangle/(PSRR)^2$ is the input referred supply noise. The noise variance, (5.24), may be substituted in (5.8) to determine the SNR with the supply noise included:

$$SNR = \frac{(\eta P)^2}{\frac{2q(\eta P)}{1-M} B + \frac{4kT}{R_f} B + 4kT \left(\frac{4\pi^2 \Gamma}{3} \right) \frac{C_{in}^2}{g_m} B^3 + \frac{\langle v_{supply}^2 \rangle}{(PSRR)^2}} \quad (5.26)$$

As with the extinction ratio, the effect of supply noise on the system may be quantified in terms of a power penalty. Also, since the effect of supply noise on both the error probability and jitter appears exclusively through the SNR , the amplitude, P' , required to undo the effects of a non-zero $\langle(V_{supply})^2\rangle$ may be determined from the expression

$$SNR|_{P'; \langle v_{supply}^2 \rangle > 0} = SNR|_{P'; \langle v_{supply}^2 \rangle = 0} \quad (5.27)$$

using the SNR given by (5.26). The supply-noise power penalty may then be defined as

$$K_{SUP} \equiv \frac{P'}{P} \quad (5.28)$$

1. Whereas this assumption is quite good in the case of the receiver transfer function, it is often not true that the supply-to-output transfer function is flat. The important thing to keep in mind is that it is the input-referred supply noise that appears in the SNR formula.

This power penalty may be written in terms of the fraction of optical noise in the supply-noise-free system

$$S_P \equiv \frac{2q}{1-M} \frac{B \cdot SNR}{\eta P} = \frac{\langle i_{in}^2 \rangle_{optical}}{\langle i_{in}^2 \rangle_{total}} \Bigg|_{P; \langle v_{supply}^2 \rangle = 0} \quad (5.29)$$

and the ratio of the input-referred supply noise to the total noise in the supply-noise-free system given by

$$S_{SUP} \equiv \frac{\left[\frac{\langle v_{supply}^2 \rangle}{(PSRR)^2} \right]}{\left(\frac{1}{S_P} \right) \left(\frac{2q(\eta P)}{1-M} B \right)} = \frac{\left[\frac{\langle v_{supply}^2 \rangle}{(PSRR)^2} \right]}{\langle i_{in}^2 \rangle} \Bigg|_{P; \langle v_{supply}^2 \rangle = 0} \quad (5.30)$$

With these definitions, the supply noise power penalty becomes

$$K_{SUP} = \frac{1 + S_P}{2} + \left[\left(\frac{1 - S_P}{2} \right)^2 + S_{SUP} \right]^{\frac{1}{2}} \quad (5.31)$$

The supply noise power penalty is plotted in Figure 5.17 in terms of the supply noise fraction, S_{SUP} . The plot shows that supply noise begins to be an issue when it accounts for around one tenth of the total noise. The power penalty remains below ~3 dB provided the supply noise is kept to at most one third of the total noise. As is generally true of the power penalties considered in this section, if the optical noise level was large in the original system (large S_P) then the power penalty is higher. In the case of the supply noise power penalty, the dependence on S_P is relatively weak since the presence of a large amount of supply noise automatically means that the optical noise is relatively less important.¹

1. By comparison, the non-zero extinction power penalty, K_M , is approximately equal to $S_P/(1-M)$ as M approaches 1, whereas the supply noise power penalty, K_{SUP} , approximately equals $(S_{SUP})^{1/2}$, independent of the value of S_P , when S_{SUP} is large.




Figure 5.17: Power penalty for supply noise

The degradation in error probability and jitter which results from the presence of supply noise may be compensated by increasing the optical signal amplitude by a factor K_{SUP} . This factor is referred to as the power penalty and is given by (5.31). The power penalty depends on the ratio of supply noise to total noise in the supply-noise-free system, S_{SUP} , and on the fraction of optical noise in the supply-noise-free system, S_P .

Intersymbol interference

The discussions of error probability and jitter above have not dealt with the effects of intersymbol interference (ISI). In a simple model, described in Appendix D and summarized in Figure 5.18, ISI is measured in terms of the fractional eye closure, C_{ISI} . This parameter depends on both the transmitter waveform and the receiver frequency response. Since there is no clear general model for C_{ISI} , the following analysis treats it as a high-level design parameter in much the way that the receiver bandwidth has been.

ISI results in both vertical eye closure, which degrades the error probability, and in jitter. The effect of ISI on the on each of these, as quantified by a power penalty, is somewhat different. To determine the power penalty for maintaining a given error probability when $C_{ISI} > 0$, equate the error probability for the original system (amplitude P and $C_{ISI} = 0$) with that of the modified system (amplitude P' and $C_{ISI} > 0$).

$$Pr\{Error\}_{P;C_{ISI}>0} = Pr\{Error\}_{P;C_{ISI}=0} \quad (5.32)$$

Eye closure enters the error probability by reducing the separation between ISI-free signal mean and the decision level from $(\bar{v}_{rx} - v_d)$ to $(\bar{v}_{rx} - \Delta v_{rx}/2 - v_d)$.

$$Q\left(\frac{\bar{v}_{rx} - \frac{\Delta v_{rx}}{2} - v_d}{\sqrt{\langle v_{rx}^2 \rangle}}\right) \Bigg|_{P;C_{ISI}>0} = Q\left(\frac{\bar{v}_{rx} - v_d}{\sqrt{\langle v_{rx}^2 \rangle}}\right) \Bigg|_{P;C_{ISI}=0} \quad (5.33)$$


The eye closure, Δv_{rx} , is defined by $\Delta v_{rx} = (C_{ISI})(v_{rx,peak-to-peak})$. Referring all quantities back to the receiver input current and equating the arguments of the Q functions then gives,

$$\frac{(\eta P') - \frac{C_{ISI}(\eta P')}{2}}{\sqrt{\langle i_{in}^2 \rangle}} \Bigg|_{P;C_{ISI}>0} = \frac{(\eta P)}{2} \Bigg|_{P;C_{ISI}=0} \quad (5.34)$$

Which can be expressed in terms of the $SNRs$ by squaring each side of the equation

$$(1 - C_{ISI})^2 \cdot SNR \Big|_{P;C_{ISI}>0} = SNR \Big|_{P;C_{ISI}=0} \quad (5.35)$$

Note that the $SNRs$ are not equal although the error probabilities are. The SNR is defined in terms of the amplitude of the signal and the noise variance. It does not reflect the deterministic, pattern-dependent “shifts” in the signal which result from ISI. To compensate for the shift in the signal relative to the decision



$$C_{ISI} = \frac{\Delta v_{rx}}{v_{rx, peak-to-peak}}$$

$$(Jitter)_{peak-to-peak} = C_{ISI} \cdot t_{rx}$$

Figure 5.18: Eye closure and jitter due to intersymbol interference

As discussed in Appendix D, intersymbol interference (ISI) results in vertical eye closure and jitter. (a) Represents the transition region of an ISI-free eye pattern. The underlying signal has the peak-to-peak amplitude and transition times indicated. (b) Shows a simplified model of the effect of ISI on this signal. The fractional eye closure, C_{ISI} , and the transition time determine the jitter.

level, the SNR is increased as shown in (5.35). The ratio of the improved SNR to the original drives the ISI sensitivity power penalty and is defined as $D_{ISI,sens}$.

$$\frac{SNR|_{P';C_{ISI}>0}}{SNR|_{P;C_{ISI}=0}} = \frac{1}{(1 - C_{ISI})^2} \equiv D_{ISI,sens} \quad (5.36)$$

Before stating the explicit formula for the ISI sensitivity power penalty, it is worthwhile examining the ISI jitter power penalty. To find the power penalty for the effect of ISI on jitter, equate the peak-to-peak output jitter of the original system (amplitude P and $C_{ISI}=0$) with that of the modified system (amplitude P' and $C_{ISI}>0$).

$$(Jitter)_{peak-to-peak}|_{P';C_{ISI}>0} = (Jitter)_{peak-to-peak}|_{P;C_{ISI}=0} \quad (5.37)$$

The peak-to-peak ISI jitter is given by the product $(C_{ISI})(t_{rx})$. This peak-to-peak jitter must be added to the peak-to-peak jitter resulting from noise at the comparator input (see Appendix D). The latter term has variance t_{rx}^2/SNR , as in (5.13). Using (5.19) to relate the jitter variance to a peak-to-peak quantity then reduces (5.37) to

$$\left[C_{ISI}t_{rx} + \alpha \sqrt{\frac{t_{rx}^2}{SNR}} \right]_{P';C_{ISI}>0} = \left[\alpha \sqrt{\frac{t_{rx}^2}{SNR}} \right]_{P;C_{ISI}=0} \quad (5.38)$$

This expression may be rearranged to give the ratio of the improved SNR to the original

$$\frac{SNR|_{P';C_{ISI}>0}}{SNR|_{P;C_{ISI}=0}} = \frac{1}{\left(1 - \frac{C_{ISI} \sqrt{SNR|_{P;C_{ISI}=0}}}{\alpha} \right)^2} \equiv D_{ISI,jitter} \quad (5.39)$$

This is the driving term for the ISI jitter power penalty. The quantity $\alpha/(SNR)^{1/2}$ which divides C_{ISI} in this expression is the ratio of the peak-to-peak noise to the peak-to-peak signal in the original system (amplitude P). $D_{ISI,jitter}$ plays the same role in this case as $D_{ISI,sens}$ in (5.36) does for the ISI sensitivity power penalty.

The same power penalty formula applies in both cases using the appropriate definition of D_{ISI} . As with all of the power penalty expressions, the ISI power penalty depends on the fraction of optical noise in the original system:




Figure 5.19: Power penalty for the effect of ISI on sensitivity

The degradation in error probability which results from the presence of intersymbol interference (ISI) may be compensated by increasing the optical signal amplitude by a factor $K_{ISI,sens}$. This factor is referred to as the power penalty and is given by (5.41) with D_{ISI} defined in (5.36) depending on the fractional eye closure, C_{ISI} .

$$S_p \equiv \frac{2q}{1-M} \frac{B \cdot SNR}{\eta P} = \frac{\langle i_{in}^2 \rangle_{optical}}{\langle i_{in}^2 \rangle_{total}} \Big|_p \quad (5.40)$$

The ISI power penalty, determined from the SNR ratio ((5.36) or (5.39)), is given by

$$K_{ISI} = \frac{S_p D_{ISI}}{2} + \left[\left(\frac{S_p D_{ISI}}{2} - 1 \right)^2 + D_{ISI} - 1 \right]^{\frac{1}{2}} \quad (5.41)$$

The ISI sensitivity power penalty is plotted in Figure 5.19 and the ISI jitter power penalty is plotted in Figure 5.20. As long as the eye closure does not exceed ~10-20%, the sensitivity may be compensated by a reasonable (~3 dB) increase in amplitude. Except in very low amplitude conditions, however, compensat-




Figure 5.20: Power penalty for the effect of ISI on jitter

The degradation in jitter which results from the presence of intersymbol interference (ISI) may be compensated, to a point, by increasing the optical signal amplitude by a factor $K_{ISI,jitter}$. This factor is referred to as the power penalty and is given by (5.41) with D_{ISI} defined in (5.39) depending on the ratio of the fractional eye closure, C_{ISI} , and the quantity $C_{noise} = \alpha / (SNR)^{1/2}$ which represents the peak-to-peak noise normalized to the peak-to-peak signal in the original system. When these two quantities are equal, the jitter due to ISI is equal to the noise-related jitter in the original system. At this point $K_{ISI,jitter}$ diverges indicating that no finite increase in power can compensate the jitter.

ing for the jitter effect is not possible. This occurs because $D_{ISI,jitter}$ diverges as the eye closure, Δv_{rx} , becomes comparable to the peak-to-peak noise. At this point, the amount of jitter due to ISI is equal to the original jitter due to noise. To compensate the jitter it would be necessary to reduce the noise-related jitter to zero, and this is not possible with a finite amplitude. The conclusion to be drawn from this is that either ISI must be eliminated, or an acceptable level of ISI jitter must be explicitly included in the jitter budget.

Minimum bandwidth constraint

The proceeding discussion has examined how the various parameter of the digital optical link enter into the error probability and jitter performance metrics. This understanding is useful in guiding the design of the link, and in particular the receiver. However, it must be noted that the design may be constrained by factors which do not explicitly appear in the above formulation. An important instance of this is a constraint on the receiver bandwidth.

The optimum bandwidth discussed above may not be realizable because it is either too high or too low. An upper limit on the bandwidth may be related to the intrinsic capabilities of the underlying electronics technology, or it may be follow, for example, from a power consumption limitation. On the other hand, a lower limit on the bandwidth may originate from the need to control ISI. A minimum bandwidth constraint may also arise in connection with the need to control static timing variations. Unlike jitter, which refers to dynamic timing variation driven by noise and ISI, static variations depend on deviations of electronics process parameters and of the temperature from their nominal values. The major sources of delay variation are related to propagation through the linear receiver and to variations in the decision level.

The propagation delay through the linear portion of the receiver can be quantified in terms of the Elmore delay [199,200]. For a single pole system with time constant τ , the Elmore delay is $\tau=1/(2\pi f_{3dB})$, that is, the delay is equal to the time constant. When cascading systems, the delays add. A cascade of N similar stages, each with a time constant of approximately τ_0 , would have a delay of

$$T_D = N \cdot \tau_0 \quad (5.42)$$

Using (5.14), the bandwidth of the cascaded system would be roughly

$$B = \frac{1}{2\pi\tau_0\sqrt{N}} \quad (5.43)$$

The bandwidth and delay of the cascaded system can then be related as

$$T_D = \frac{\sqrt{N}}{2\pi B} \quad (5.44)$$

The change in the delay can then be related to the change in the bandwidth as

$$\Delta T_D = -\frac{\sqrt{N}}{2\pi B^2} \Delta B \text{ or } \frac{\Delta T_D}{T_D} = -\frac{\Delta B}{B} \quad (5.45)$$

Both the delay and the static delay variation are reduced by increasing the bandwidth. The need to reduce these quantities by increasing B may result in a minimum bandwidth requirement.

Static delay variations may also result from changes in the DC bias-points of the receiver. Specifically, changes in the position of the decision level, v_d , relative to the comparator input, v_{rx} . As with jitter, the static delay variation depends on the rise time of the comparator input

$$|\Delta T_D| = \frac{|\Delta v_d|}{v_{rx, peak-to-peak}} \cdot t_{rx} \quad (5.46)$$

This expression is equivalent to the ISI jitter formula given in Figure 5.18 but with Δv_d playing the role of the eye closure. That is, a displacement of the decision level is equivalent to a displacement of the signal, v_{rx} . Equation (5.46) may be elaborated by using (5.16) for t_{rx} .

$$|\Delta T_D| = \frac{|\Delta v_d|}{v_{rx, peak-to-peak}} \cdot \left(1 + \frac{t_{rx}^2}{\gamma^2} B^2 \right)^{\frac{1}{2}} \frac{\gamma}{B} \quad (5.47)$$

This form of the expression points out the dependence of the rise time on receiver bandwidth and transmitter rise time. Once again, it may be necessary to set a lower bound on the bandwidth in order to control the static delay variation.

If the minimum bandwidth requirement is larger than the optimum discussed above, the optical signal amplitude may be increased in order to reduce the jitter to the specified value. If desired, a power penalty may be formulated to quantify this aspect of the design.

Minimum amplitude constraint

In a practical receiver design, there is likely to be a minimum amplitude constraint on the output of the linear receiver, v_{rx} . For example, the comparator may not function correctly unless its input is switched beyond a certain voltage range (e.g. it may not produce a valid output or it may operate too slowly), or, as it occurs in Chapter 7, it may be desirable to hard-limit v_{rx} .

Let V_{min} be defined as the minimum amplitude requirement on v_{rx} . The corresponding minimum photocurrent amplitude would then be $V_{min}/(R_f G)$. Furthermore, let L be the ratio of the optical amplitude to the minimum required optical amplitude. Using these definitions, the photocurrent amplitude used above can be written as

$$\eta P = \frac{L \cdot V_{min}}{R_f G} \quad (5.48)$$

This formula may be used to recast the SNR , jitter, and power penalty expressions above in terms of the amplitude constraint and the receiver's gain. In particular, the SNR , (5.12), and jitter, (5.17), may be rewritten as

$$SNR = \frac{(LV_{min})^2}{\frac{2q}{1-M} \frac{LV_{min}}{R_f G} B + 4kT R_f G^2 B + 4kT \left(\frac{4\pi^2 \Gamma}{3} \right) \frac{C_{in}^2}{g_m} R_f^2 G^2 B^3} \quad (5.49)$$

$$\sigma_{jitter}^2 = \gamma^2 \left(\frac{2q}{(1-M)(LV_{min})} \frac{R_f G}{B} + 4kT \frac{R_f G^2}{(LV_{min})^2 B} + 4kT \left(\frac{4\pi^2 \Gamma}{3} \right) \frac{C_{in}^2}{g_m} \frac{R_f^2 G^2}{(LV_{min})^2 B} \right) \left(1 + \frac{t_{ix}^2}{\gamma^2} B^2 \right) \quad (5.50)$$

These expressions are somewhat misleading in suggesting that the interconnect performance is improved by increasing V_{min} and decreasing G . This is merely a consequence of the fact that, for a fixed value of L , increasing G or decreasing V_{min} results in a lower optical amplitude which, as seen above, negatively affects interconnect performance. Nonetheless, although the gain did not play a role in the previous expressions, the inclusion of a minimum amplitude requirement makes the receiver gain an important practical design consideration.

The value of (5.49) and (5.50) lies in their potential use, along with a model of the interconnect power dissipation, as part of an optimization procedure. For a given circuit topology, the gain and bandwidth of the receiver may be formulated in terms of physical parameters. Then, for example, the jitter may be mini-

mized while constraints are applied to the power dissipation and *SNR*. Alternately, the power dissipation may be minimized subject to a performance constraint. A similar optimization procedure has been implemented recently [201,202], but in this work the interconnect performance is simply characterized in terms of the receiver output rise/fall time. While this approach is intuitively appealing, it does not recognize the fundamental role of the error probability and jitter in characterizing the interconnect performance.

Three determinants of receiver sensitivity

In an optical communication receiver, the sensitivity is defined as the minimum power (or amplitude) required to meet a given bit-error-rate (or error probability) requirement. The same definition may be applied to the digital optical interconnect receiver, but it must be augmented with two other criteria. The minimum optical amplitude required to meet a given jitter requirement will generally be different from that required to meet the error probability requirement, and is likely to be larger. In addition, as just discussed, a minimum optical amplitude may be required for correct large-signal operation of some portion of the receiver. The latter two issues give rise to a “dynamic” sensitivity while the error-probability related sensitivity may be referred to as the “static” sensitivity since it applies to the ability of the receiver to produce the correct output in the steady state. The actual sensitivity of the receiver is the largest of the three values determined by each of these three constraints.

Summary

The key to understanding the digital optical interconnect is to recognize that the output of the receiver is binary-valued. This means that the performance of the interconnect is quantified by the uncertainty in the output logic value, i.e. the error probability, and by the uncertainty in the time-instance of the logic transitions, that is, the jitter.

A well designed digital optical interconnect may be analyzed in terms of the conceptual model of Figure 5.7. In this model, the transmitter switches an optical signal between two prescribed power levels in a certain time. The “linear receiver”, consisting of a transimpedance amplifier and a post amplifier and having a certain overall bandwidth, converts the optical signal into a voltage waveform which is applied to the input of a comparator. The comparator may be explicitly designed into the receiver, or this function may be performed by the logic gate which follows the receiver.

The designer’s task is to choose the high and low optical power levels, the optical transition time, and the receiver gain and bandwidth in order to meet the desired error probability and jitter specifications.

There are, of course, constraints on the various physical parameters such as the detector capacitance, feedback resistance, etc. Likewise, there are upper and lower bounds on the optical power levels and on the receiver bandwidth. The analytical formalism derived in this section allows the designer to examine the interplay between these parameters in determining the link performance.

As with any model, there is a trade-off between accuracy and simplicity, the latter being favored for use in hand analysis and in building design intuition. This section has opted for simplicity. In reality, the noise processes are not wide-sense-stationary, white, or Gaussian. The noise distributions in the low and high states are not identical. The transmitter includes noise sources in addition to the optical shot noise. The noise variances do not depend only on the receiver bandwidth but on the receiver frequency response, and this dependence is specific to each noise source [194]. The models used to relate ISI and noise to jitter are themselves approximate, as are the estimates of rise/fall time.

Given the simplifications that have been made in deriving the expressions in this section, they should not be expected to give precise numerical predictions. Their primary value is in guiding design judgements. In this vein, a number of observations were made in the proceeding discussion using circuit parameters taken from the receiver design of Chapter 7.

- The optical amplitude needed to produce a vanishingly-small error probability is generally less than the power needed to meet a reasonable jitter requirement in a 1 Gb/s link.
 - An extinction ratio of up to around 40% may be tolerable.
 - The supply noise should be limited to 30% of the total noise, and below 10% would be particularly good.
 - ISI-induced jitter should be explicitly accounted for in the jitter budget. The receiver frequency response and transmitter waveform should be designed to control ISI.
 - There is a bandwidth that minimizes jitter. This bandwidth may not be attainable because of constraints on bandwidth such as the control of ISI or static delay variation.
 - Practical receivers may have a minimum amplitude requirement. This makes receiver gain an important design consideration and introduces a third limitation to receiver sensitivity along with the error probability and jitter.
-

5.5 Transmitter and Receiver Circuit Requirements

This section collects a number of guidelines for the design of the transmitter and receiver interface circuits. The overall objective is to be able to implement a 1 Gb/s digital optical interconnect with a variety of optical system architectures.

Transmitter

The laser driver circuit needs to switch the VCSEL between two current levels corresponding to LO and HI logic inputs. Both the LO and the modulation (HI-LO) currents should be variable and should range up to a few milliamperes. The transition time of the laser driver output should be below 200 ps (it is assumed that the VCSEL bandwidth will not limit the optical transition time). Both the delay through the laser driver and, more importantly, the laser driver jitter should be minimized. The power consumption for the laser driver + VCSEL combination should be around 5 mW.

Receiver

For compatibility with multimode fiber, the receiver should use a relatively large detector (75 μm diameter) corresponding to a capacitance of ~ 100 fF. For the typical g_m achievable with the H-GaAs IV process (at a reasonably low power consumption level), the receiver bandwidth that minimized the jitter appears to be around 1 GHz. The receiver should be designed to maintain this nominal bandwidth with the given input capacitance, and should dissipate around 5 mW. An input current swing of around 10-20 μA should be adequate to produce a valid logic output at 1 Gb/s. Finally, immunity to supply noise should be maximized.

6

A Low Power Laser Driver

The laser driver's purpose is to switch the current level through an integrated laser diode between two values corresponding to logical high (HI) and low (LO) inputs. This is depicted on the generic laser light-versus-current (L-I) plot of Figure 6.1. For experimental purposes, the high and low currents must be externally variable, but in actual systems these values may be hard-wired. A critical circuit parameter is the power dissipation required to operate the laser driver, i.e. in addition to the power dissipated by the laser itself. An optical output power requirement of roughly 1 mW (HI state) and a total budget of 5 mW for the combination of the laser driver and laser were specified in Chapter 5. With VCSEL power efficiencies in the 20-50% range, these specifications translate into approximately 1 mW of power for the driver circuit.




Figure 6.1: L-I characteristics of a generic laser diode

A generic laser light-versus-current input (L-I) curve indicating the desired mode of operation in a digital optical interconnect. The laser current is switched between HI and LO operation points which are both above the lasing threshold.

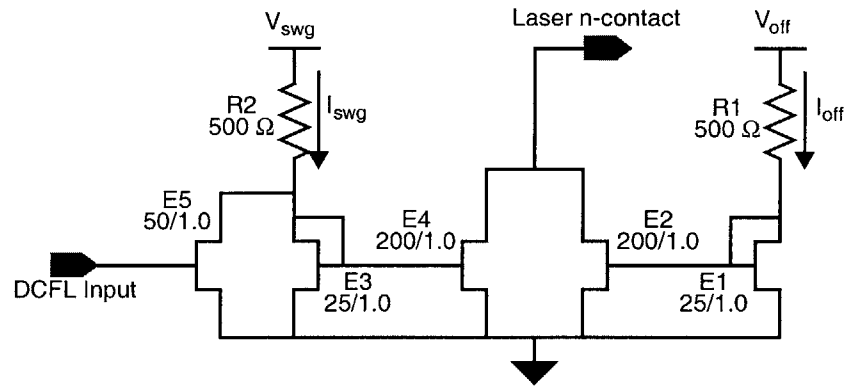


Figure 6.2: Schematic of first-generation switched current mirror laser driver

First-generation laser driver based on a switched current mirror. The current mirror formed by E1 and E2 sets the low-level laser current. A second current mirror, E3 and E4, set the additional current applied for the high-level state. This second mirror is modulated by shunting its reference current through E5.

Numerous laser driver designs have been reported in the literature [203-207]. These are intended for use in lightwave communication systems which operate at high power levels (current levels ~ 50 mA) and are based on current steering using a differential-pair. However, a laser driver design based on a switched current mirror has been pursued for the present application because of its relative simplicity and potentially lower power dissipation. The current mirror output also requires much lower “head-room,” allowing the laser supply voltage to be minimized.

6.1 First-generation Switched Current Mirror Laser Driver

A first-generation implementation of this concept is shown in Figure 6.2. This circuit was fabricated in the H-GaAs III process and was used in the preliminary link demonstration of Chapter 5. The n-type contact of the laser is connected to the driver at the point indicated while its p-type contact is attached to a suitable power supply. In this laser driver, E1 and E2 form a current mirror which sets the output-low laser current level. To reach the output-high level, an additional current is applied by the second current mirror formed by E3 and E4. This “swing” current is modulated by using E5 to shunt the reference current from this second current mirror when the input is high (~ 0.6 V). To achieve a fast rise time, R2 must be made small to reduce the RC time constant on the gate of E4. This leads to greater variability in the swing current in response to the temperature- and process-dependencies of E3, increases the power dissipated due to I_{swg} ,

and requires a larger E5 transistor leading to larger delay and power dissipation in the stage leading up to the laser driver. I_{swg} also rises dramatically when being shunted by E5 leading to a significant increase in power dissipation. A more intractable problem results from an effect known as drain lag. Drain lag will appear numerous times in the discussions to follow, and is explained in Appendix E. In the circuit of Figure 6.2, drain lag in E3 and E5 causes the gate of E4 to exhibit a response tail lasting ~ 100 ns following an initial, rapid rise (the initial rise is ~ 250 ps in the circuit above). This undesirable response is mimicked in the output current. Attempts at eliminating the drain lag in this simple circuit have been unsuccessful.

6.2 Design of an Active Switched Current Mirror Laser Driver

An improved version of a switched-current-mirror laser driver is shown in Figure 6.3. The low-level current is now set by the E210/E211 current mirror. The modulation current is determined by the switched current mirror, E110/E111. As shown symbolically in Figure 6.3(a), this mirror is activated or deactivated through a “switch” controlled by the input. Since the reference side of the current mirror is not switched, drain lag is no longer a problem at this point in the circuit. R100 can now be made larger in order to reduce the variability in the modulation current. The key to the operation of this laser driver is the switch which connects the two sides of the current mirror. The details of this switch are shown in Figure 6.3(b) and a representative transient simulation is shown in Figure 6.4¹.

Turn-on cycle

To turn on the modulation current, the gate of E111 is charged by current which passes primarily through the Schottky gate of the pass transistor, E102. Beginning with the input high², nodes 111 and 102 are discharged so that E111 is off and the two sides of the current mirror are isolated by the off-state pass-transistor. With E103 also initially off, the inverter formed by E101 and D101 produces a high signal in response to the falling edge on the input. Node 102 is low at this point, so E105 begins to conduct current I_{C1} . Since node 111 is also low, the Schottky gate of E102 becomes forward-biased with respect to node 111 and passes a portion of I_{C1} to the right (I_{C2}). Since node 110 is at a higher voltage, as set by the “reference side” of the current mirror, the gate of E102 does not pass significant current to the left (I_{C3}). There is also some amount of capacitive coupling between node 102 and node 111 which helps in the charging pro-

1. This is conceptually similar to the Vitesse-patented “squirt buffer,” but the details of the two circuits differ significantly.

2. The input logic levels are those of direct-coupled FET logic (DCFL): HI= ~ 0.6 V, LO= ~ 0 V.

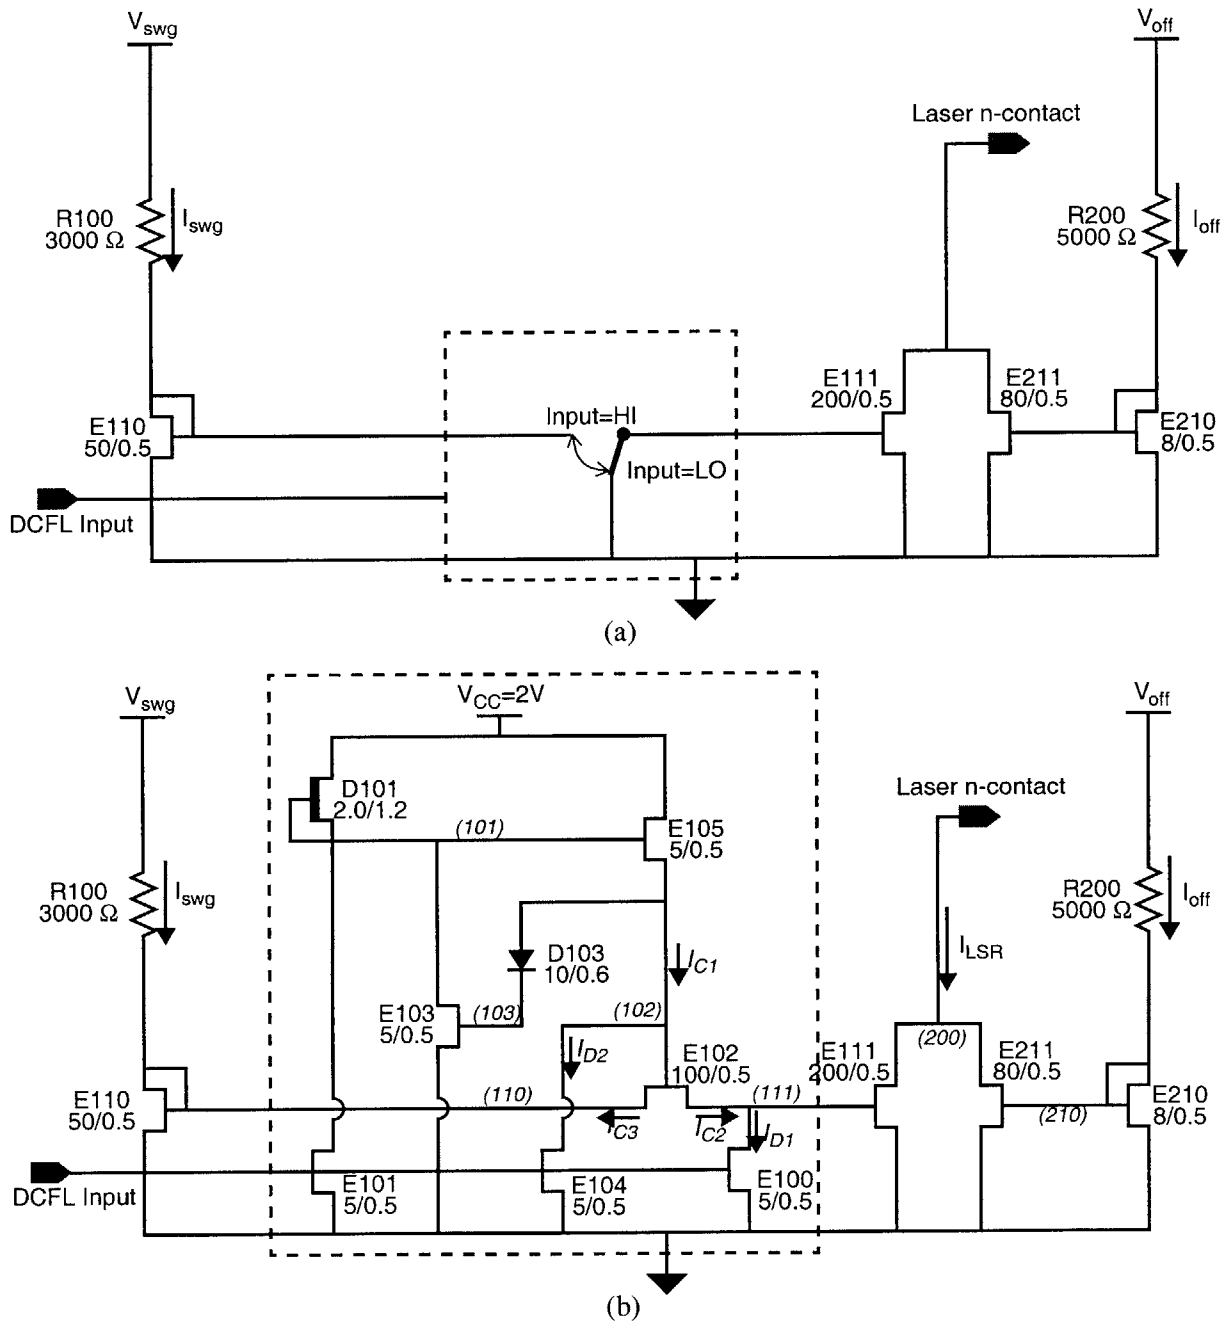


Figure 6.3: Schematic of an active switched current mirror laser driver

Improved laser driver based on a switched current mirror. (a) is a simplified version showing the current mirror which sets the low-state current (E210, E211) and the swing current mirror (E110, E111). The latter is modulated by means of the “switch” controlled by the input. (b) includes the details of the switch which consists of pass-transistor E102 and a circuit to rapidly control its state. The switching operation is shown in Figure 6.4 and described in the text. The values in parenthesis are node identification numbers.

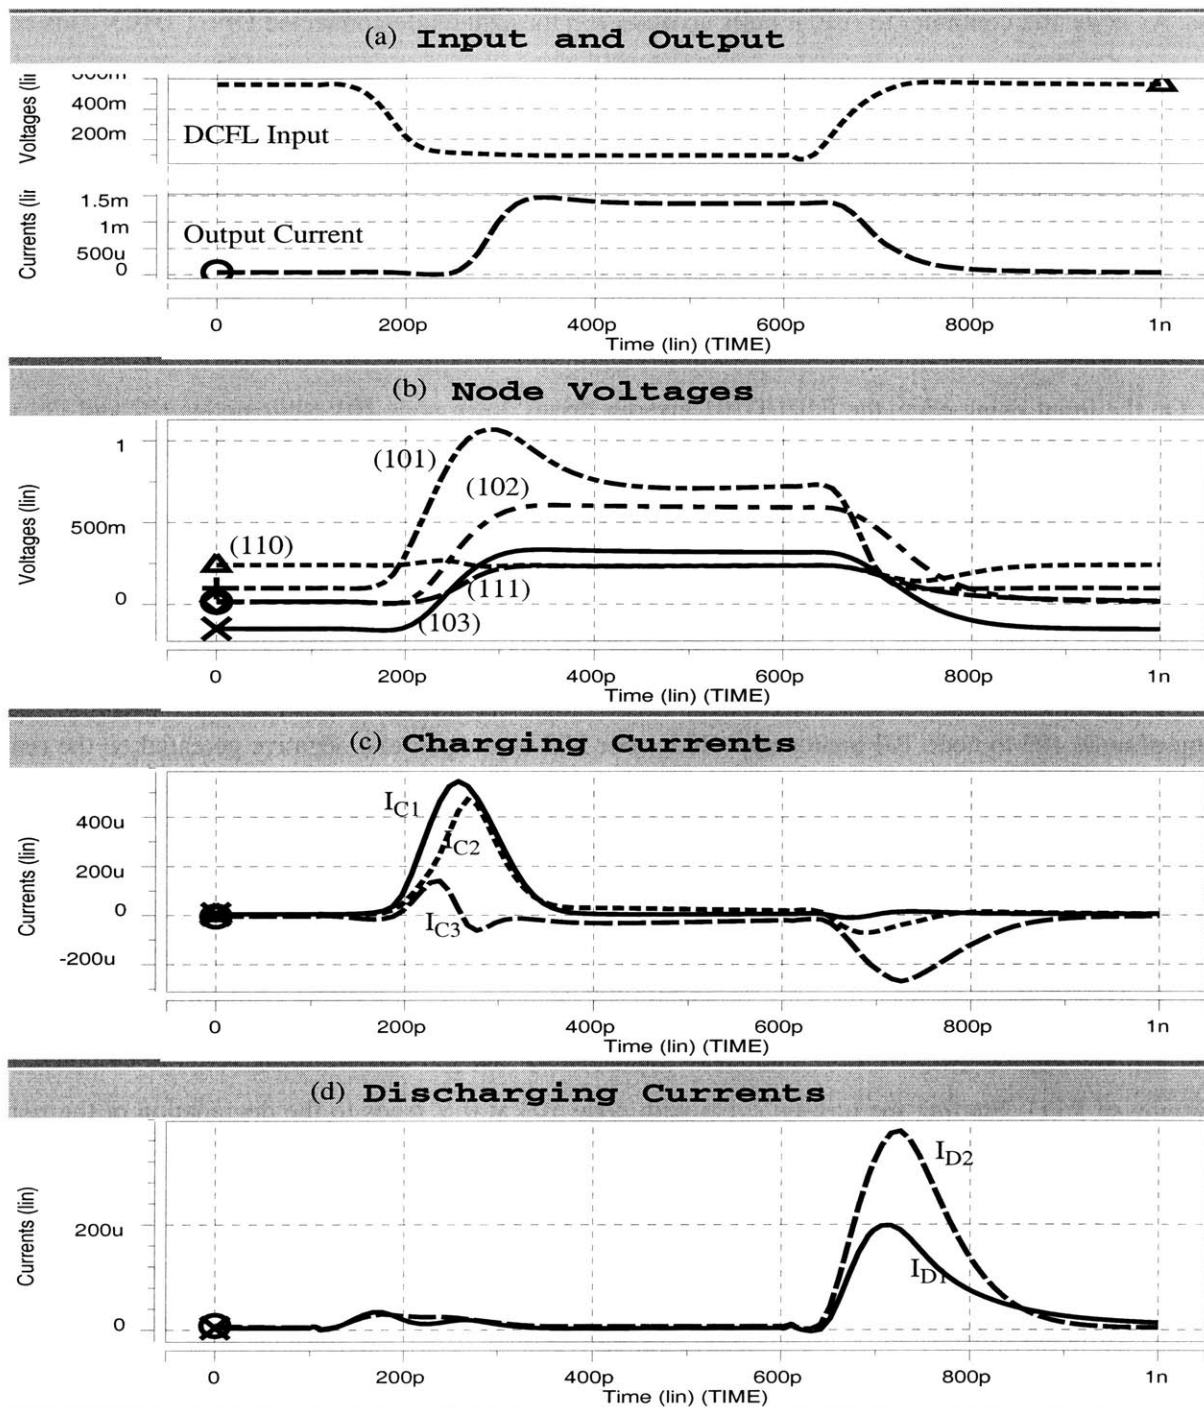


Figure 6.4: Simulation of active switched current mirror laser driver

Transient simulation of the improved laser driver showing the current-mirror switching operation. (a) The DCFL input and the laser current output (only the modulated current is shown). (b,c,d) Voltages and currents as indicated. See Figure 6.3 for node and current identifications.

cess. As node 102 continues to rise, it pulls up node 103 through diode-connected DFET D103. This continues until E103 turns on and pulls down node 101, which in turn shuts off E105. At this point node 102 is high enough that the pass transistor is on and the current mirror is “activated.” If the reference current is set too high then node 111 will not be fully charged by the gate current from the pass transistor. In this case, node 111 will be further charged by R100; this process is slow because of R100’s large resistance.

Turn-off cycle

On the input rising edge, the E101/D101 inverter drives down node 101 while nodes 102 and 111 are discharged by E104 and E100, respectively. This again isolates the two sides of the current mirror (by turning off E102) and turns off the output current (by turning off E111).

Node 103 floats

Notice in Figure 6.4 that the voltage on node 103 begins and ends at a negative value. The details of this behavior were not observed prior to submission of the design for fabrication. Due to the capacitive coupling of node 103 to node 101 and/or node 102, node 103 is pumped to a negative potential on the rising edge of the input (turn-off cycle). With the input high, node 103 is isolated by the gate of E103 and the diode-connected DFET D103. In the most extreme case, node 103 is pushed to -150 mV following the rising input. This condition was, in fact, used throughout the design. If the turn-on cycle begins with node 102 precharged to a less negative value (such as 0 V), node 103 reaches the threshold of E103 sooner than if node 103 had been precharged negative. This results in the premature termination of the charging current, I_{C1} , leaving the gates of E102 and E111 inadequately charged and requiring R100 to complete the charging of E111. Starting the turn-on cycle with node 103 at 0 V leads to the degradation of the rising edge of the output current.

Normal operation of the circuit drives node 103 to successively lower values on each turn-off cycle. It takes around 1000 1 ns cycles to reach a steady state starting with node 103 at 0 V. On the other hand, the charge on node 103 leaks out through the gate of E103, forward conduction of D103, and the source/drain parasitic junctions of D103. Starting at -150 mV, node 103 discharges to -50 mV in 50 ns and continues to discharge very slowly thereafter, remaining below -40 mV for more than 1 μ s. Simulations have shown that starting the turn-on cycle with node 103 at this point results in acceptable output current rising edges for current levels up to 1 mA. Thus, provided the circuit is allowed to run for a short time to precharge, and the

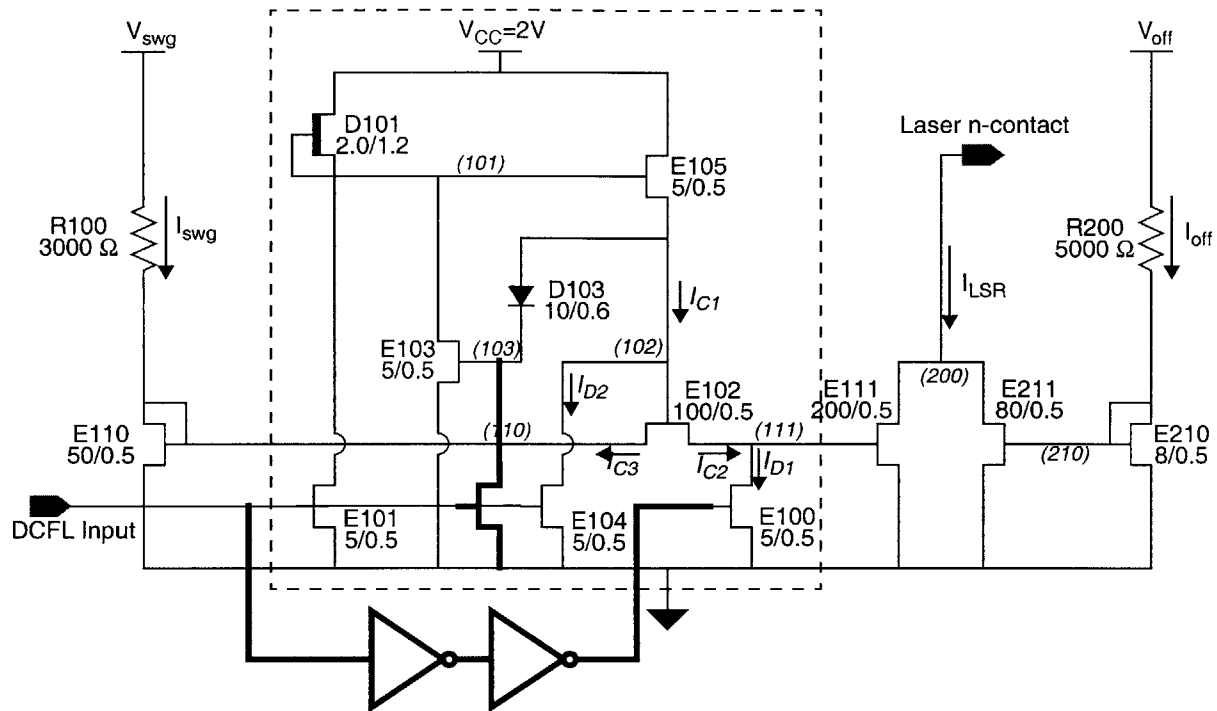


Figure 6.5: Modifications to laser driver of Figure 6.3

Modifications are shown in bold. An EFET is added to discharge node 103 on each turn-off cycle to prevent it from floating, and a pair of inverters are used to form a delay leading up to E100 in order to reduce the delay asymmetry between the rising and falling transitions

maximum input high duration is limited to around 1 μ s, it should be possible to use this circuit for experimental purposes. To move beyond this experimental use, the circuit must be improved by adding one transistor to discharge node 103 as shown in Figure 6.5, and resizing E103 and D103 to maintain the duration of the charging cycle.

Delay asymmetry

The input to output delay is roughly 100 ps shorter for a falling output than for a rising output. This occurs because the additional circuitry in the signal path used in the turn-on cycle. This asymmetry may be improved in a future version of the circuit by inserting a delay between the input and the gate of E100, as in Figure 6.5.

Performance

In simulations, using nominal device parameters at 85°C, this laser driver switches the modulation current in roughly 100 ps with 100 ps and 50 ps delays to turn on and turn off the output, respectively. Over a temperature range of 25°C to 125°C and process variations of two standard deviations, the delay parameter vary between 20ps and 200 ps, and the rise/fall times vary between 50 ps and 400 ps.

Under nominal process conditions at 85°C, the circuit draws 0.6 mW of average power (50% duty cycle) in addition to the reference currents and the laser current. The design goal of 5 mW total laser-and-driver power dissipation for a 1 mW high-level output power may be achieved using this circuit in conjunction with a VCSEL having a power efficiency of around 35%, threshold current of around 0.5 mA and an operating voltage of around 1.6 V. These VCSEL specifications are consistent with reported VCSEL performance [178].

Physical design

The laser driver of Figure 6.3 was designed in Vitesse's four metal layer H-GaAs IV process. A layout plot of the circuit is shown in Figure 6.6. The cell occupies a 148 μm x 54 μm area, and requires an external connection to the n^+ isolation ring which surrounds it.

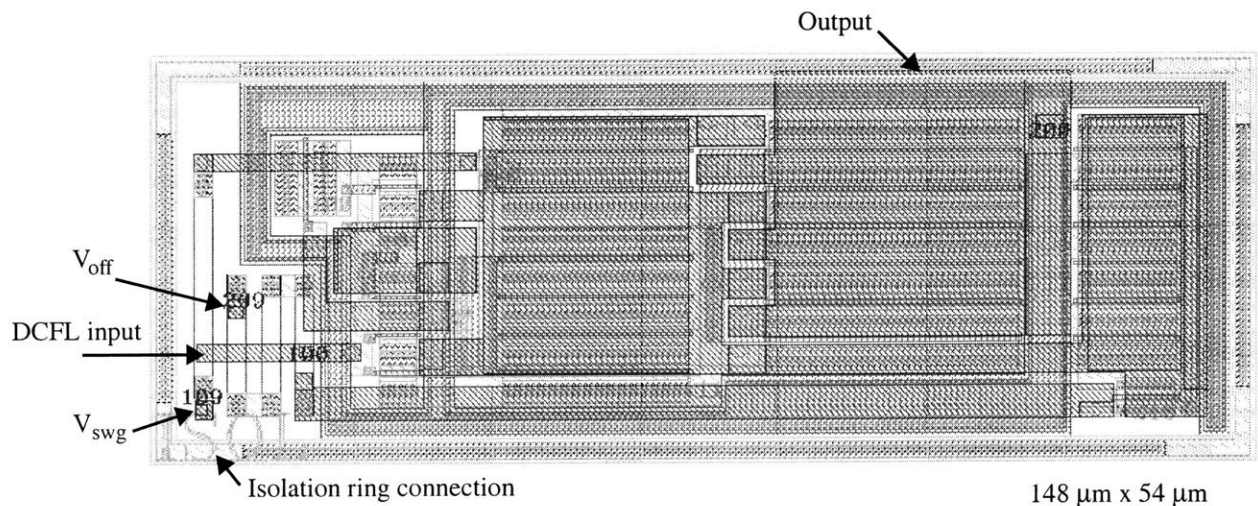


Figure 6.6: Layout of the active switched current mirror laser driver

Designed in Vitesse's H-GaAs IV process, the layout of the laser driver in Figure 6.3 is shown here. For clarity, this plot does not include the sheets of Metal-3 and Metal-4 used for $V_{CC}=2\text{V}$ and ground, respectively. The isolation ring should be connected to the most positive supply available.

The source and drain regions have been lengthened to 3 μm in order to minimize contact resistance. According to the H-GaAs IV Technology Reference Manual, the nominal contact resistance drops from 750 $\Omega\cdot\mu\text{m}$ at a contact length of 1.2 μm down to 204 $\Omega\cdot\mu\text{m}$ for contact lengths exceeding 2.8 μm , and contact lengths of at least 3 μm are recommended when contacts are shared [208]. However, the junction capacitance increases as the source/drain regions are lengthened. Subsequent to the completion of the above design, it was learned that a contact length of around 2 μm is believed to be optimal [209].

Resistors were implemented as source/drain implants with a nominal sheet resistance of 300 Ω/\square . The temperature coefficient of variation of these resistors is reported to be zero [208], and the resistors are typically matched to within 2% across a die. The width of 2.5 μm was chosen for the resistor implants in order to minimize across-die variability. However, there is around a 20% run-to-run variation in resistance. The ohmic contact length to the resistors was set at 3 μm . Whereas in the source/drain sizing the contact area is a performance issue, here it is one of predictability (i.e. by the extraction tool) and uniformity.

Liberal use was made of p-contacts in order to minimize the resistance to the backgate and to control the generation of substrate noise by the laser driver. The active portions of the laser driver were surrounded by a p-contact ring placed within the n^+ isolation ring. The p-contacts were connected to ground within the cell.

The interconnect metallization was also sized liberally in order to minimize the likelihood on electromigration failure during circuit testing. Most likely, the interconnect sizing should be reduced in the future to minimize stray capacitance.

6.3 Measurement Setup

The MIT-OEIC-7 test chip shown in Appendix C (Figure C.1) may be cut to isolate the die shown in Figure 6.7. This die contains the laser driver shown in Figure 6.3 and Figure 6.6. This circuit is referred to as the VCSEL driver because the current levels it is designed to drive are suitable for a typical VCSEL. The die in Figure 6.7 also includes a scaled-up version of this laser driver design which is suitable for driving the higher currents required by in-plane lasers. This circuit is referred to as the IPSEL driver and will not be discussed further. This section describes the experimental setup used to characterize the VCSEL driver.

Clock and data

The primary laser driver test block is at the bottom of the die in Figure 6.7. In addition to the laser driver, it includes digital circuitry for on-chip clock and data generation. Details of this circuit are provided in Appendix C which also describes the construction of the chip-on-board packaging into which the die was incorporated. The clock source on the die uses a number of ring oscillators and D-flip-flops to generate 16 different frequencies between 25 MHz and 1 GHz. Four input bits select one out of the sixteen signals. Alternately, an external clock may be connected to the circuit in place of the on-chip source. Close examination of the internal clock showed it to exhibit a large amount of jitter (see Appendix C). For this reason, the external clock option was utilized in testing the laser driver.

The desired clock (either the external clock or one of the sixteen internal clocks) is output from the die as an ECL signal and is used to drive the internal pseudorandom bit sequence (PRBS) generator. The PRBS circuit is a state machine consisting of a 23-bit shift generator and an XNOR feedback. It generates a sequence of $2^{23}-1$ bits before repeating. Taken 23 bits at a time, the sequence will enumerate all possible 23-bit binary numbers with the exception of the all-1 state (i.e. the number $2^{23}-1$). The all-1 state locks-up the sequence because of the use of XNOR feedback (as opposed to XOR feedback which would result in the all-0 state being the lock-up state). An input is provided to flush the shift register with 0's in this event.

The PRBS output is available from the circuit as an ECL signal, and is routed internally to the laser driver. However, operating the ECL outputs results in a large amount of supply bounce which corrupts the output of the laser driver. To circumvent this problem while allowing the data pattern to be viewed externally, one module was used solely for clock/data generation while the laser driver was operated on a separate module. The experimental setup used to generate an oscilloscope trigger and the signals needed to test

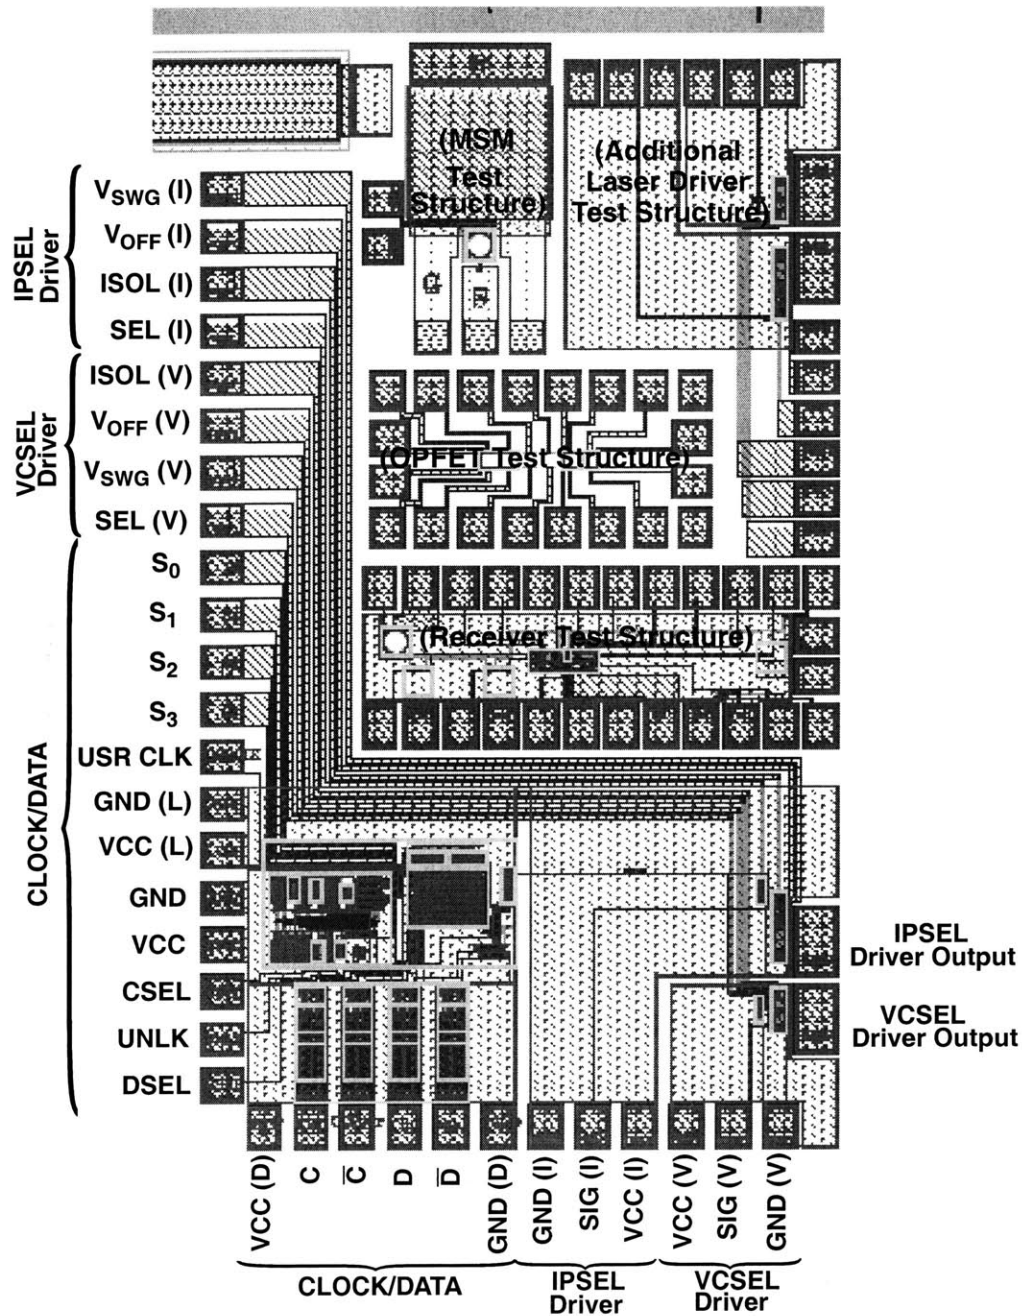


Figure 6.7: Layout of die cut from MIT-OEIC-7 for laser driver characterization

The MIT-OEIC-7 test chip shown in Appendix C (Figure C.1) may be cut to isolate the die shown here. This die contains the laser driver shown in Figure 6.3 and Figure 6.6, referred to here as the VCSEL driver, as well as similar laser driver scaled to support the higher current requirements of in-plane laser. The primary laser driver test block at the bottom of the die includes on-chip clock and data generation and is detailed in Appendix C.

the laser driver is shown in Figure 6.8. The clock output from the Clock/Data module is used to elicit the switching response of the laser driver while either of the data outputs may be used to produce an eye diagram. The clock output of the module is used rather than directly using the external clock in order to produce consistent rise/fall time and jitter characteristics in the switching response and eye pattern measurements.

A key observation regarding the data pattern generated by the Clock/Data module is made in Figure 6.9. A 185 ps delay is observed between the rising and falling edges. This delay will appear in the laser driver eye patterns. Although it reduces the eye width, it allows the jitter in the rising and falling edges of the laser driver output to be measured separately.

Electrical test module

Two circuit boards were constructed to support electrical and optical testing of the laser driver (either of the boards can be used to for the Clock/Data module). The construction of these modules is detailed in Appendix C. A diagram of the signal path implemented for electrical testing of the laser driver is shown in Figure 6.10. The three-resistor network was used to improve the termination of the transmission line leading to the input. Since the input to the chip is a DCFL logic gate, a high-level input signal is clamped at ~ 0.6 V. When the input is low, the input impedance is high since the Schottky gate of the input MESFET is off (the input is modeled as a capacitor in this case.) But, when the input is high, the input impedance becomes small because the Schottky diode turns on. If the transmission line is terminated directly into the MESFET gate, both extremes of input impedance are encountered during each transition. The resistive network limits the variation in the load impedance to $33\ \Omega$ when the input is high to and $50\ \Omega$ when the input is low. A 6 dB reduction in the input amplitude is also produced; the input signal being used is large enough so that this is not an issue. Overall, the matching network provides an incremental improvement in signal fidelity, but it is not crucial to the experiment.

The signal input to the chip passes through a multiplexer before arriving at the laser driver. The purpose of the multiplexer is to select between the data output of the on-chip PRBS source and the external signal source. In these experiments, the input to the module was set to select the external signal source. The on-chip PRBS generator and the clock and data ECL outputs were held in a static state in order to reduce supply bounce (the clock source was set to the user clock option, and no external clock was applied).

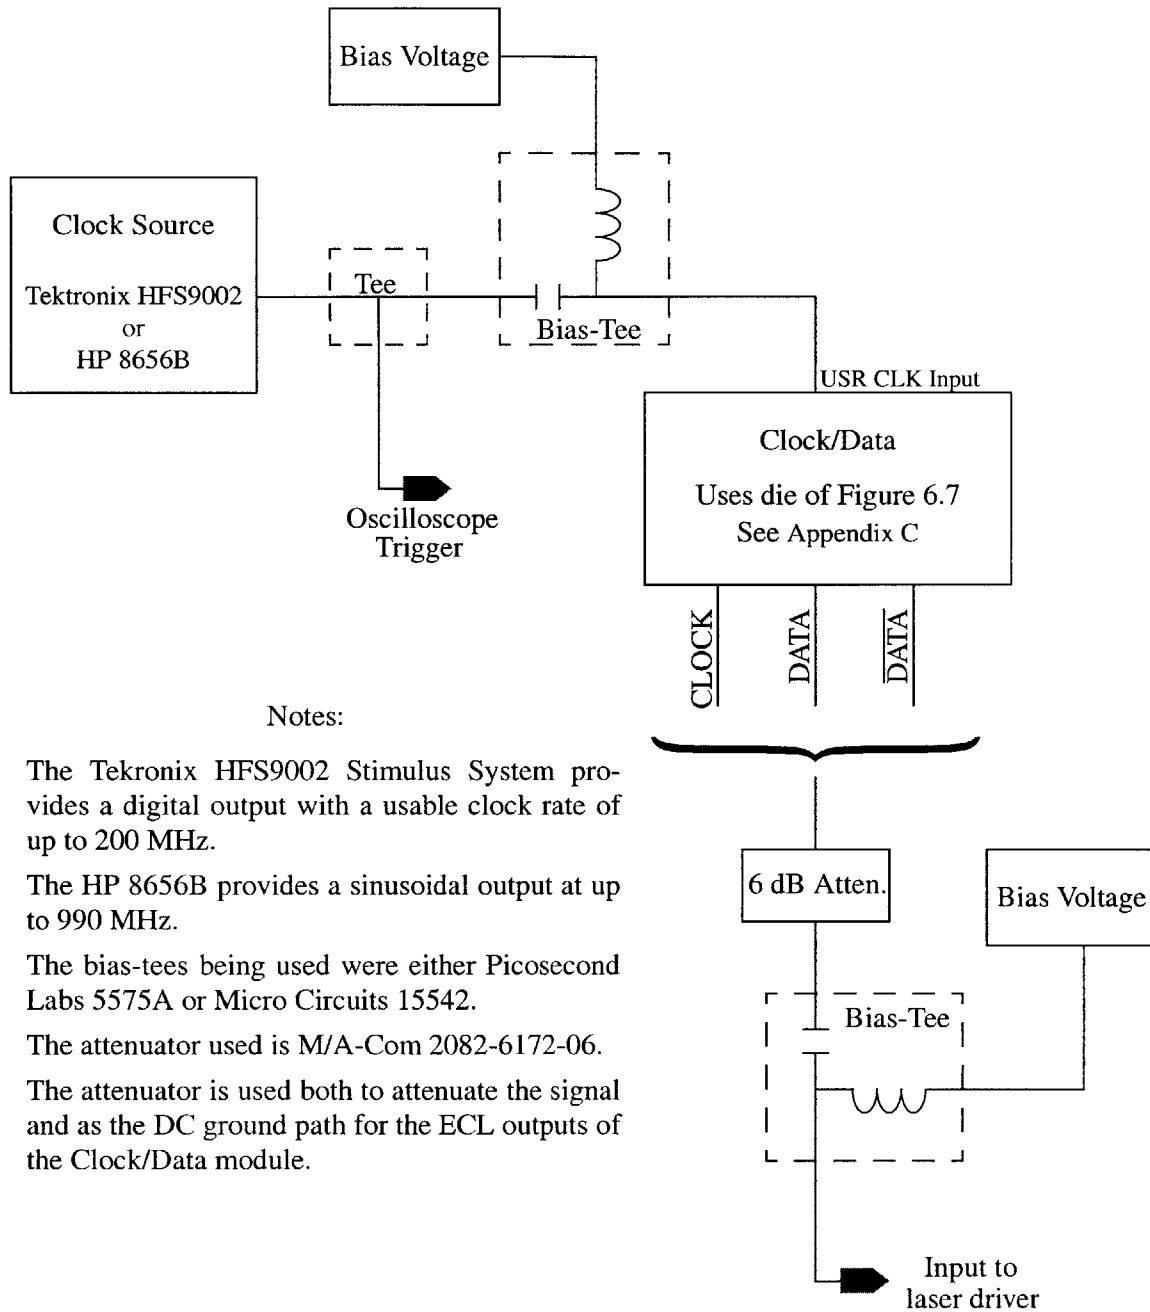


Figure 6.8: Signal source for characterization of laser driver

A $250 \text{ mV}_{\text{p-p}}$ clock output is biased to match the trip-point of the DCFL “USR CLK” input on the Clock/Data circuit of Figure 6.7. The construction of this module is detailed in Appendix C. A clock or data output from this module may be selected as the input to the laser driver being tested. The bias voltage is again adjusted to match the trip point of the laser driver test circuit’s input.

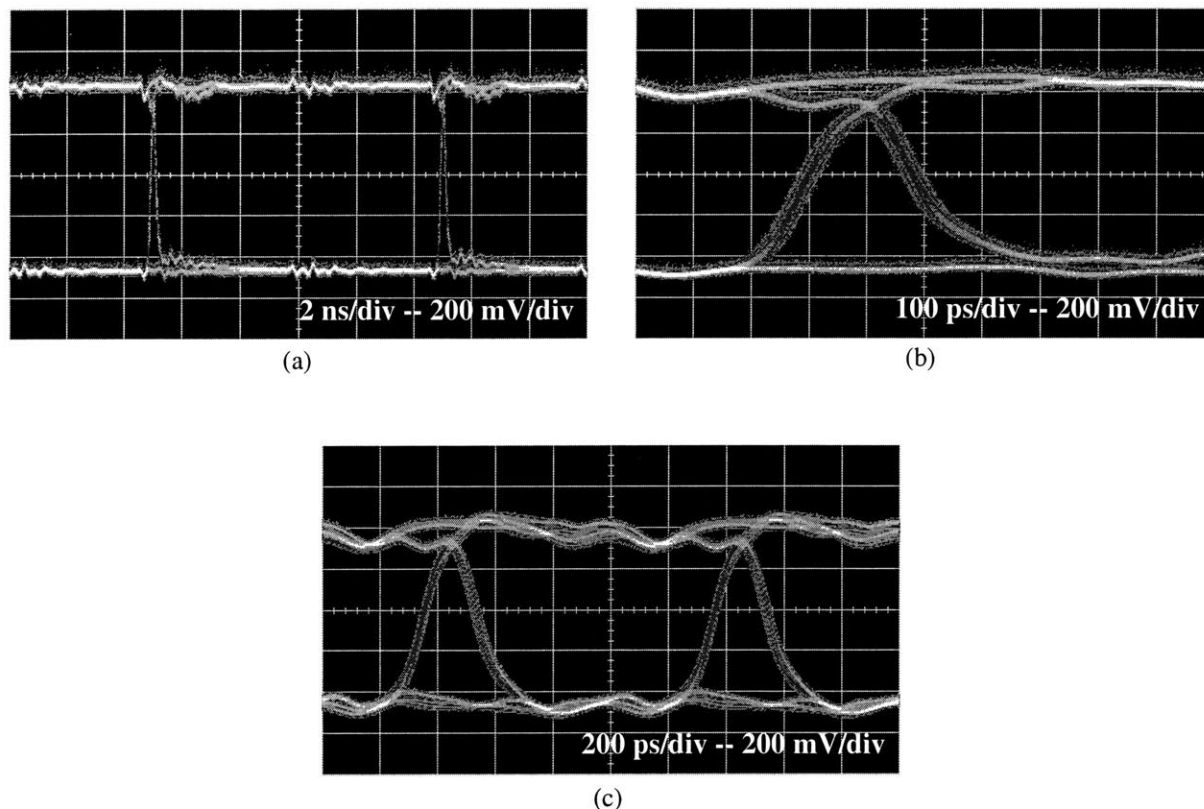


Figure 6.9: Data signal used to characterize laser driver

Eye diagrams of the data output generated by the setup shown in Figure 6.8. (a) and (b) were measured at 100 MHz and (c) was measured at 990 MHz. A 185 ps delay is observed between the rising and falling edges. The rise and fall times are roughly 200 ps.

The output of the laser driver is wire bonded to a 50 Ω transmission line. A 50 Ω resistor at the source end of the transmission line pulls-up the laser driver output to the required bias voltage. For static characterization of the laser driver, this bias port is left open and direct connection is made to the laser driver output through the SMA connector (the bias tee is removed). For dynamic measurements, the bias port is connected, and the bias tee is used to block the dc path.

In observing the output waveforms from the electrical test module, it is important to distinguish between the intrinsic response of the laser driver and parasitic effects created by the board. Figure 6.11 shows the output of the module in response to rising and falling edges at its input with the laser driver powered off. In these plots, the measured output voltage has been converted to a current by dividing by the 25

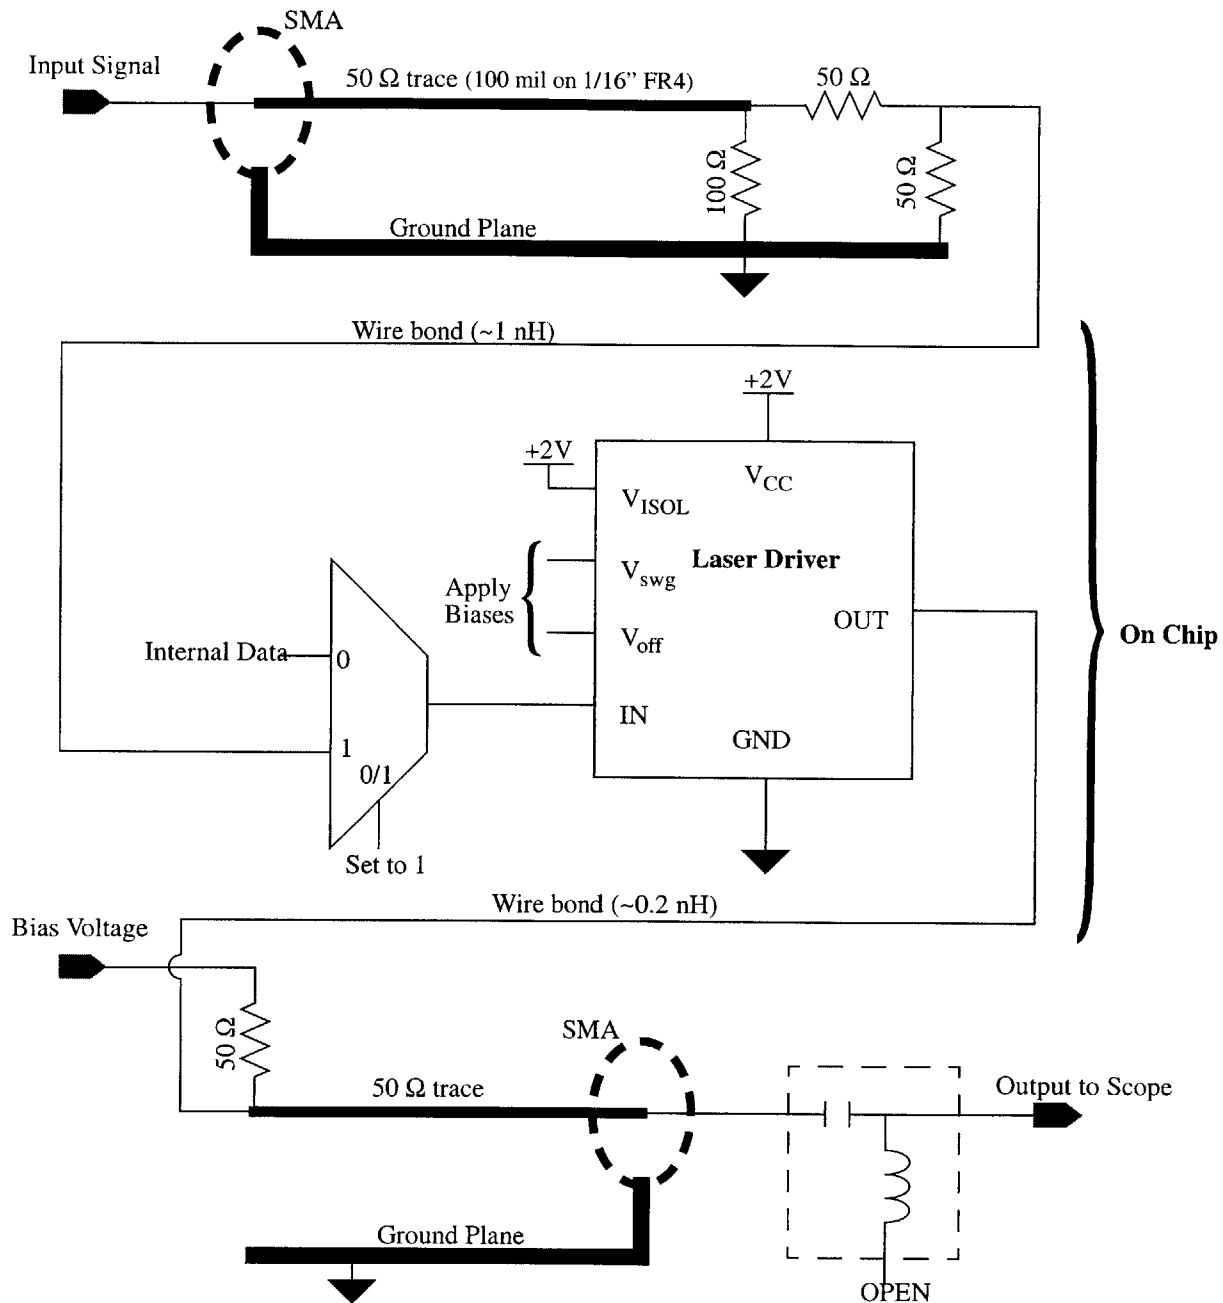


Figure 6.10: Signal path for electrical characterization of the laser driver

The output from Figure 6.8 is applied as the "Input Signal" above. The resistor network is designed to improve the termination into the MESFET gate input of the multiplexer on the chip. The multiplexer is configured to pass the external signal, rather than an on-chip signal, onto the laser driver. The output of the laser driver is pulled up to the V_D bias voltage, also known as V_{200} , through a $50\ \Omega$ resistor. Finally, a bias-tee is used to ac-couple the signal to the HP 5475A oscilloscope.

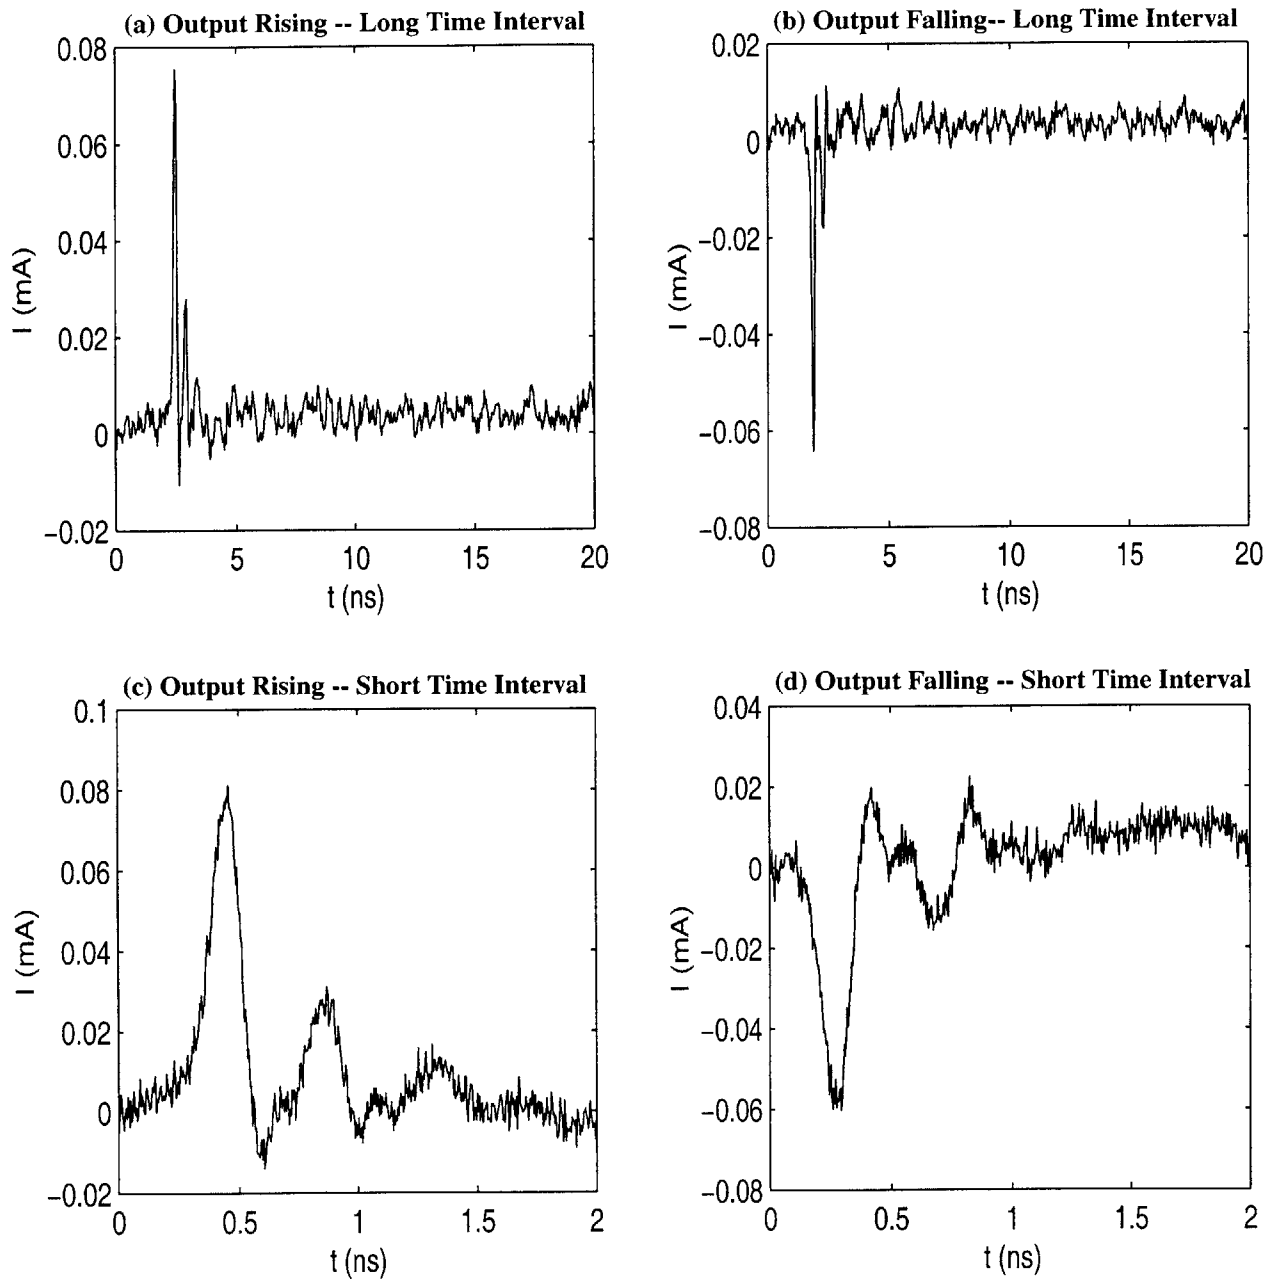


Figure 6.11: Parasitic coupling of input to output in laser driver electrical test module

With the laser driver circuit turned off, the clock output from Figure 6.8 is applied to the input to the electrical test module in Figure 6.10. (a) and (c) show the response of the output to a rising input edge while (b) and (d) show the falling edge response. The measured output voltage is converted to a current by dividing by 25Ω . This is done for consistency with the laser driver output measurements. The parasitic coupling seen in the plots here must be kept in mind when studying the laser driver output data.

Ω equivalent resistance appearing at the laser driver output. A similar procedure will be followed when measuring the output step response of the laser driver. The parasitic coupling seen in Figure 6.11 must be kept in mind when studying the laser driver output data.

Optical test module

A similar module is used for testing the laser driver-plus-VCSEL combination. The signal path in this case is shown in Figure 6.12. No matching network is used in this case (this module was designed prior to the electronic module). Instead, a 50 Ω resistor is placed at the SMA connector end of the 3 cm long trace to improve the termination. The output of the laser driver is wire-bonded to the adjacent laser die. This die contains a 2x10 array of VCSELs made on a p-type substrate. Each VCSEL has an individual, top-side n-type contact which may be driven by the laser driver. The VCSEL was made at Sandia National Laboratory, emits at 850 nm, and uses oxide-apertures to achieve a low threshold current and high power efficiency [85,178-187]. Its characteristics will be examined shortly.

The laser is butt-coupled to a 62.5 μm graded index core multimode fiber. Index matching gel is applied between the surface of the VCSEL and the fiber in order to reduce reflections from the fiber facet into the laser. Such reflections vary the laser's operating characteristics and are a significant source of noise. The fiber is connected to a 13 GHz bandwidth photodetector from Discovery Semiconductor (model DSC50).

Oscilloscope

The output from both the electrical and optical dynamic test was measured using an HP 5475A digitizing oscilloscope with an HP54751A 20 GHz input module. The measurements were bandwidth limited to 12.4 GHz through an internal oscilloscope setting in order to reduce noise contribution from the scope. This was done after verifying that the bandwidth limitation did not limit the measurements, that is, that the signals being observed occupied a smaller bandwidth.

Laser characteristics

Static characteristics from a VCSEL of the type used in this work are shown in Figure 6.13. The lasing threshold for this device is found to be 1 mA. Threshold currents of similar devices have been found to vary from 0.5 mA to 1.6 mA. The fiber-coupled power is seen to rise linearly with current above threshold. The diode turns on at around 1.2 V, and the current rises exponentially up to the lasing threshold of 1 mA

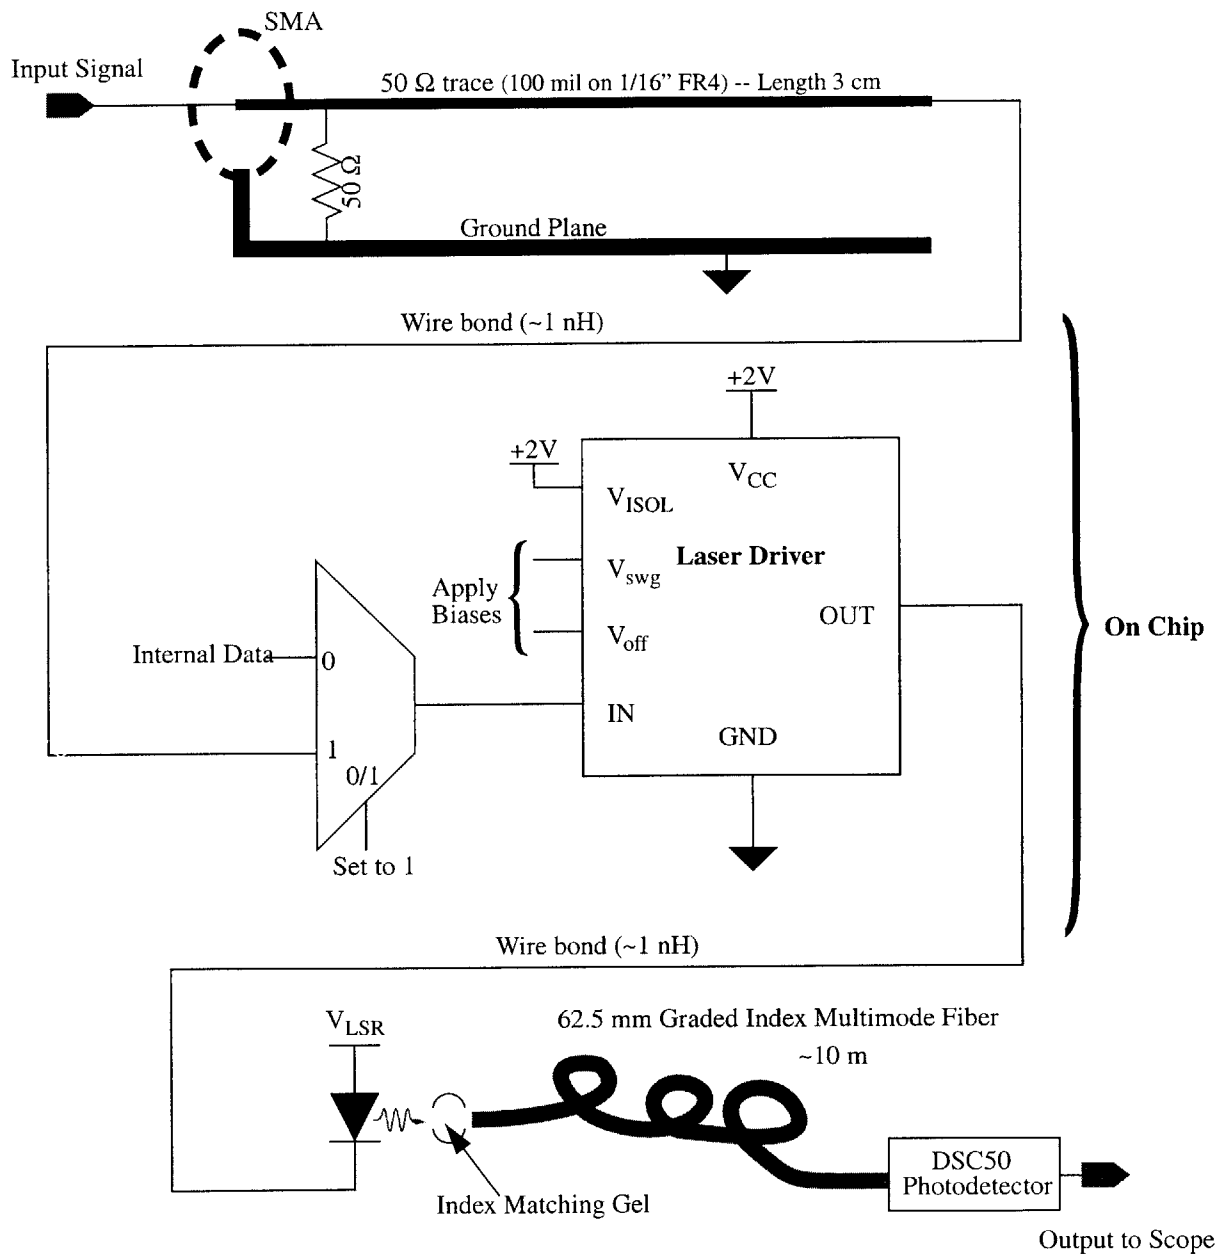


Figure 6.12: Signal path for optical characterization of the laser driver

The output from Figure 6.8 is applied as the "Input Signal" above. The multiplexer is configured to pass the external signal, rather than an on-chip signal, onto the laser driver. The output of the laser driver is wire-bonded to the adjacent laser die. The laser is butt-coupled to a multimode fiber. Index matching gel is used to reduce reflections from the fiber facet into the laser. The fiber connects to a Discovery Semiconductor DSC50 13 GHz photodetector, the output of which is viewed on the HP 5475A oscilloscope.

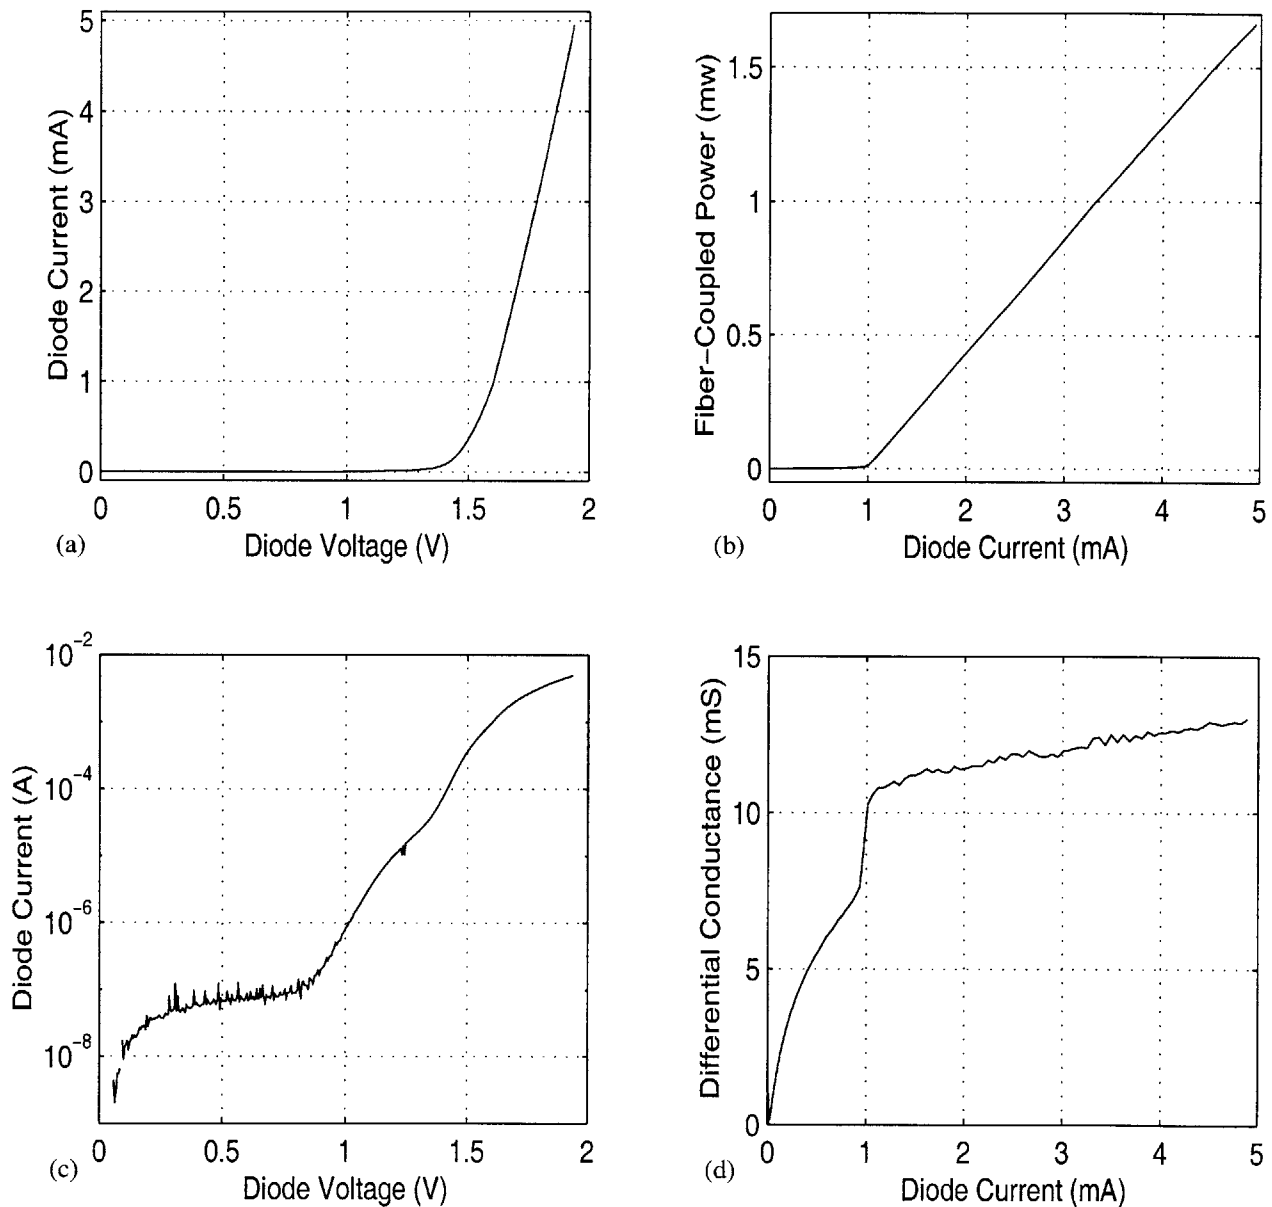


Figure 6.13: Static characteristics of a typical VCSEL.

These characteristics are typical of the oxide-apertured VCSELs fabricated at Sandia National Laboratory and used in this work. (a) Shows the I-V curve. The diode turns on at around 1.2 V, and the current rises exponentially up to 1 mA at which point the current rise becomes linear with voltage. This pattern of behavior is seen more vividly in the log(I)-V plot in (c). There appears to be $\sim 0.1 \mu\text{A}$ of leakage current either in the device or, more likely, in the measurement setup. (b) Shows the optical power coupled into a multimode fiber ($62.5 \mu\text{m}$ graded index core) as a function of the diode current. The threshold current of this device is found to be 1 mA. Finally (d) is a plot of differential conductance, dI/dV , versus the diode current. Above threshold, this quantity is around 12 mS corresponding to a resistance of 83Ω .

at which point the current rise becomes linear with voltage. This occurs at around 1.6 V, slightly above the Fermi-level separation of 1.46 eV needed to support lasing at 850 nm. The sharp change from an exponential diode response to a linear I-V occurs because of the clamping of the carrier concentration in the active region at the onset of lasing. The differential conductance, dI/dV is seen to rise gradually from 11 mS to 13 mS above threshold. The typical conductance of 12 mS corresponds to a resistance of 83 Ω . The resistance is due primarily to the distributed Bragg reflectors used to form the VCSEL's optical cavity.

The static characteristics suggest that, with the exception of packaging parasitics, the above laser can be modeled as a series combination of the 83 Ω mirror resistance with a 1.6 V drop representing the clamped junction. Unfortunately, this is far from true when modulating the laser. Figure 6.14 shows the S_{11} characteristics of a similar VCSEL measured by direct co-planar ground-signal-ground probing over a frequency range of 45 MHz to 12 GHz. Measurements were made above threshold at bias currents of 1, 2, and 3 mA. The input admittance corresponding to these measurements is shown in Figure 6.15.

The equivalent circuit shown in Figure 6.16 was used to fit the data at 2 mA (only a slight bias dependence was found at the other current levels). Conceptually, C_p and R_p represent the parasitics associated with bonding pads and interconnects (a capacitance to the substrate, and the substrate resistance), R_m represents the mirror resistance, and C_a and R_a account for the active device. Unfortunately, the active device is not well modeled by a simple parallel R-C combination. As a result, all of the components have taken on non-physical values. Notwithstanding this lack of physicality, however, this equivalent circuit fits the measurements very well, and allows the input characteristics of the VCSEL to be accounted for in SPICE simulations of the laser driver.

The AC input impedance, around 500 Ω at 45 MHz, differs greatly from the small-signal DC impedance of 83 Ω found from the static measurements. This has been attributed to the dynamics of lateral carrier diffusion in the active region [181]. As a historical matter, only the latter information was available during the design of the laser driver. The difference in loading characteristics will be seen to have a significant effect in Section 6.7.

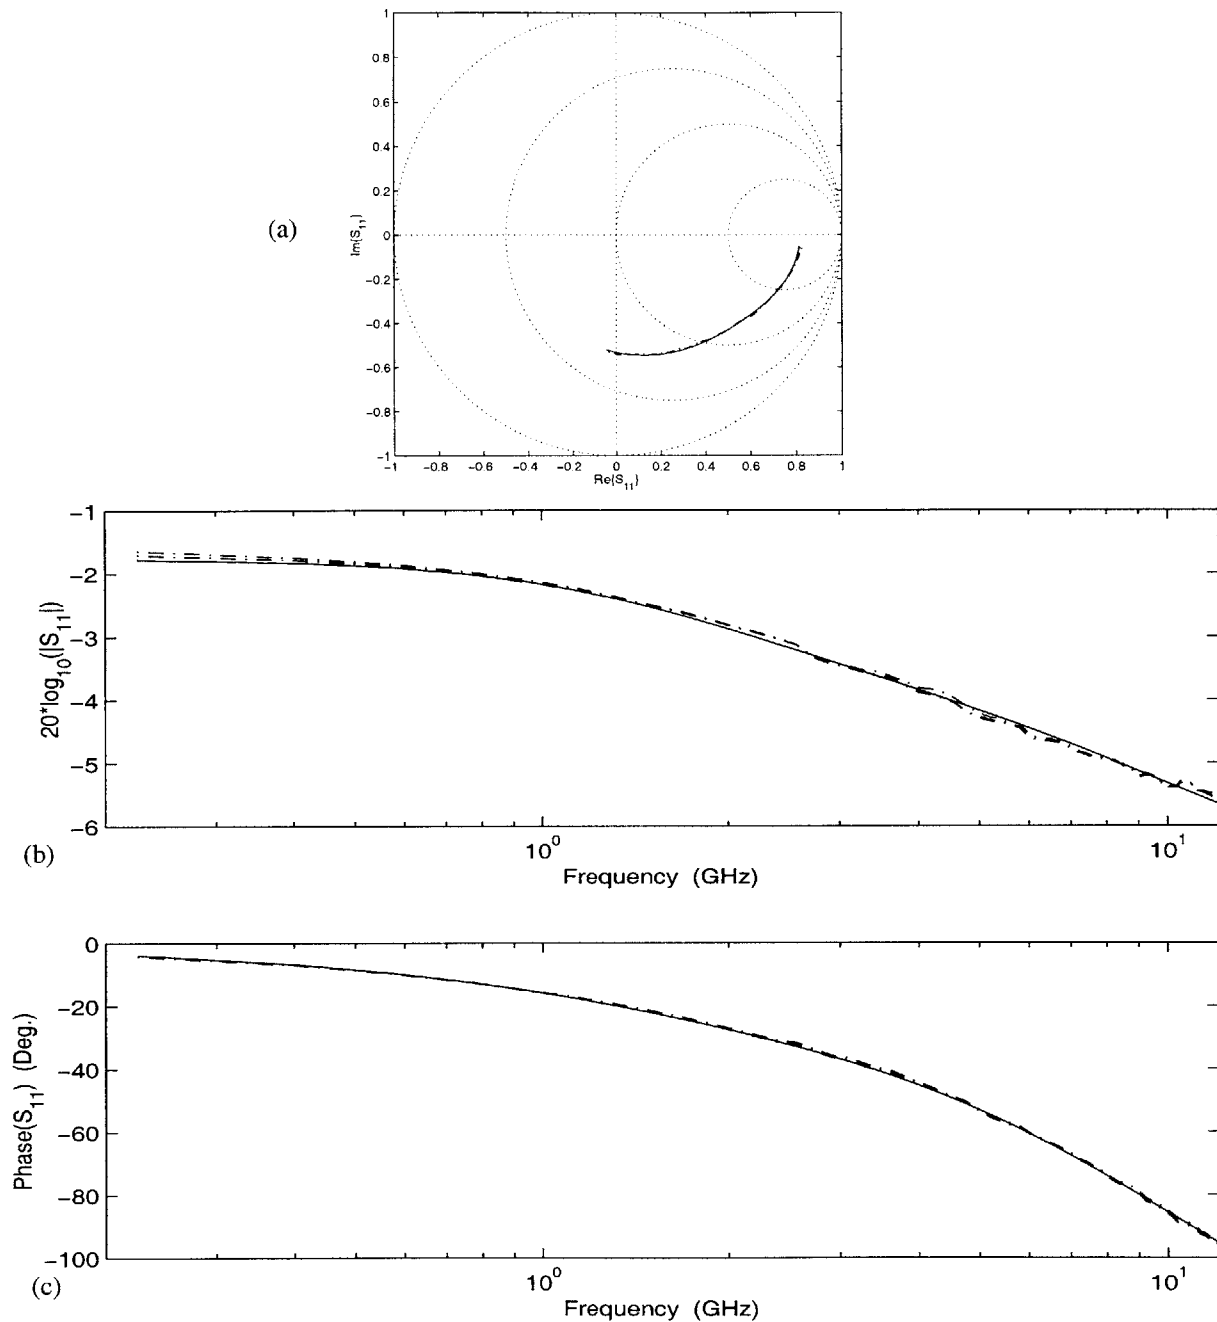


Figure 6.14: S_{11} characteristics of a typical VCSEL

These characteristics are typical of the oxide-apertured VCSELs fabricated at Sandia National Laboratory and used in this work. (a) Shows S_{11} measured between 45 MHz and 12 GHz plotted on a Smith chart, while (b) and (c) show the log-magnitude and phase of the same S_{11} data vs. frequency [210]. The dashed-dotted lines are measured values at laser bias currents of 1, 2, and 3 mA; the threshold current for this device is below 1 mA. The solid line is simulated on the basis of the equivalent circuit shown in Figure 6.16.




Figure 6.15: Input admittance characteristics of a typical VCSEL

The input admittance was calculated from the S_{11} data shown in Figure 6.14. The dashed-dotted lines again are the measured values at laser bias currents of 1, 2, and 3 mA; the threshold current for this device is below 1 mA. The solid line is simulated on the basis of the equivalent circuit shown in Figure 6.16.




Figure 6.16: Equivalent small signal input circuit for a typical VCSEL

Equivalent input circuit optimized to match the S_{11} data shown in Figure 6.14 for the data taken at 2 mA laser bias current. Conceptually, C_p and R_p represent the parasitics associated with bonding pads and interconnects, R_m represents the mirror resistance, and C_a and R_a account for the active device. Unfortunately, the active device is not well modelled by this circuit, and as a result all of the components take on non-physical values. Notwithstanding its lack of physicality, this model fits the measurements very well, and allows the input characteristics of the VCSEL to be accounted for in SPICE simulations of the laser driver.

6.4 Output Current Static Characteristics

This section will analyze the static operation of the laser driver. This consists of examining the characteristics of the “off-state” current mirror formed by E210 and E211 and those of the “on-state” current mirror which controls the modulation of the laser driver output.

Simulation of the off-state current mirror

Simulation results of the off-state current mirror are shown in Figure 6.17. Here, the off-state reference voltage, V_{off} , is swept while the reference current, I_{off} , and the output current I_{LSR} are monitored in addition to the gate voltage, V_{210} , of E210 and E211. To disable the modulation portion of the laser driver, a high-level input has been applied. The output node, 200, has been set to 1 V.

In Figure 6.17, I_{off} and I_{LSR} are seen to be flat until V_{210} rises to ~ 150 mV, the EFET threshold voltage. I_{off} and I_{LSR} then rise linearly with, maintaining roughly a 10:1 ratio. This is in agreement with the 10:1 ratio between the widths of E211 and E210. As V_{210} reaches ~ 500 mV, the rise in I_{LSR} becomes sub-linear. This is due to the onset of gate conduction (primarily in E211 due to its larger size).

Simulation of the modulation current mirror

Simulation results of the modulation current mirror are shown Figure 6.18. In this case, the modulation reference voltage, V_{swg} , is swept. The reference current, I_{swg} , and output current, I_{LSR} , are monitored along with the gate voltages on the current mirror transistors, E110 and E111, and the pass transistor, E102. The off-state current mirror has been disabled by setting $V_{\text{off}}=0$ V, and as before the output node, 200, is maintained at 1 V. A high input applied to the laser driver in order to enable the modulation output. The switching circuit thus raises up V_{102} in order to turn on the pass transistor.

With the pass transistor on, V_{110} and V_{111} are able to track each other. Unlike the off-state current mirror, however, V_{110} and V_{111} do go to zero when V_{off} is at zero. This is due to current passing through the gate of E102 and pulling up nodes 110 and 111. This effect is dominant when V_{110} & V_{111} drop below ~ 250 mV. Since V_{110} does not go to zero, I_{swg} is actually becomes negative when V_{swg} is near zero. More importantly, I_{LSR} does not go to zero. In the simulation, I_{LSR} has a minimum value of $260 \mu\text{A}$.




Figure 6.17: Simulated static characteristics of off-state current mirror

The off-state reference voltage, V_{off} , is swept and the reference current and resulting laser current are monitored along with the gate of the current mirror transistors, V_{210} . The modulation portion of the laser driver has been disabled by applying a high input, and the output node, 200, is maintained at 1 V. I_{off} and I_{LSR} begin to ramp up once V_{210} reaches 150 mV, the threshold voltage of the EFETs. The current mirror is designed with a 10:1 ratio between E211 and E210. I_{LSR} and I_{off} maintain a ratio near this design value until V_{210} reaches ~ 0.5 V at which point gate current begins to flow into the MESFETs and the mirroring action gradually becomes distorted.

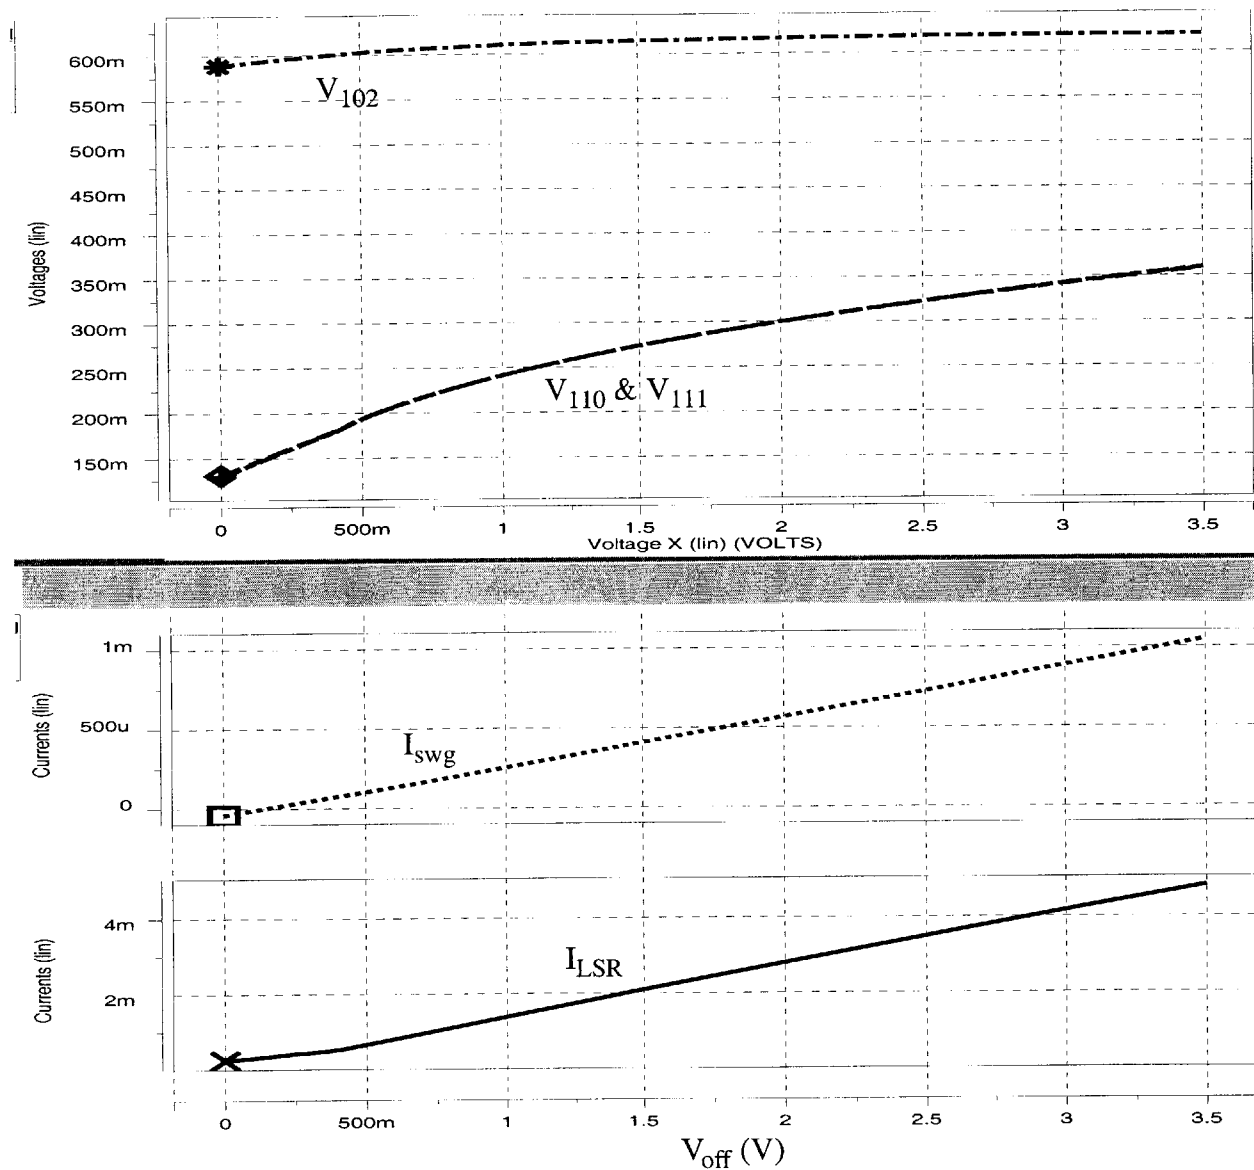


Figure 6.18: Simulated static characteristics of on-state current mirror

The modulation reference voltage, V_{swg} , is swept and the reference current and resulting laser current are monitored along with the gate voltages of the current mirror and pass transistor EFETs. The off-state current mirror has been disabled by setting $V_{off}=0$ V, and the output node, 200, is maintained at 1 V. With a high input applied to the laser driver, V_{102} is set by the switching circuit to maintain the pass transistor, E102, in the on state. This allows V_{110} and V_{111} to track each other. Note that V_{110} and V_{111} do go to zero when V_{off} is at zero. This is due to gate current from E102 which flows as V_{110} & V_{111} drop below around 250 mV. As a result, I_{swg} is actually negative when V_{swg} is near zero. Similarly, I_{LSR} does not go to zero. For V_{off} above 0.5 V, I_{LSR} and I_{swg} track with roughly the designed 4:1 ratio.

For V_{off} above 0.5 V, V_{110} and V_{111} are high enough to keep the Schottky gate of E102 turned off. This allows the current mirror to function in the same manner as the off-state circuit. I_{LSR} and I_{swg} now rise linearly with V_{swg} and maintain a ratio of roughly 4:1, as set by the width ration of E111 and E110.

Effect of output node bias voltage

In each of the above simulations, the ratio of the output current to the reference current matched the design value only approximately. The design value is set by the ratio of the gate widths of the current mirror transistors under the assumption that the pair of transistors are essentially identical in all other respects. This assumption does hold precisely because the drain voltage of the transistor on the output side of the current mirror (E111 or E211) depends on what is applied at the output node, 200, and does not generally match the drain voltage of the transistor on the reference side of the current mirror (E110 or E210).

In the above simulations, V_{200} was set to 1 V to match the value used in the dynamic characterization results of the next section. This value was determined on the basis of the measurements reported in Figure 6.19. These plots give the value of V_{200} (V_D in the figure) needed achieve the targeted mirror ratio in each of the current mirrors. The results indicate that $V_{200}=1.25$ V is needed for the off-state current mirror while $V_{200}=0.4$ V is more appropriate for the modulation current mirror.

Deviating from the values found in Figure 6.19 will result in different current mirror gains. However, as long as the output FETs of the current mirrors (E111 and E211) are maintained in saturation, no other negative effects on the laser driver performance are observed. Thus, it is possible to select a value for V_{200} that gives a reasonable match to the designed current gain values and allows for correct switching operation of the laser driver. The value of 1 V was selected in this way.

Measured static characteristics

The static characteristics of the off-state current mirror are shown in Figure 6.20. In this figure, the off-state reference voltage, V_{off} , was swept while measuring the reference current, I_{off} , and the laser output current, I_{LSR} (I_D in the figure). The output bias voltage, V_{200} (V_D in the figure), was stepped from 0.25 V to 2.0 V in 0.25 V increments, and the modulation current mirror was turned off.

From the plot of V_{off} vs. I_{off} , the value of internal resistor R200 as determined to be 6.3 k Ω . This is 26% larger than the design value, but is in line with run-to-run resistor variations for the Vitesse process.




Figure 6.19: Dependence of laser driver static characteristics on output bias voltage

V_D above refers to the voltage at the output node of the laser driver, node 200, and I_D is the laser driver output current, I_{LSR} . (a) Shows the characteristics of the off-state current mirror (modulation current mirror is off) by plotting I_D along with 10 times the reference current I_{off} for various V_{off} settings. The factor of 10 is the designed current gain of the off-state current mirror. Thus, the intersection of the I_D and I_{off} curves, marked by dots in the plot, indicate the value of V_D needed to achieve the targeted mirror ratio. (b) Repeats this measurement for the modulation current mirror (the off-state mirror is off). In this case the mirror gain is 4.




Figure 6.20: Static characteristics of the off-state current mirror

The reference voltage, V_{off} , was swept while the reference current, I_{off} , and the laser output current, labeled I_D above, were measured. The output bias voltage, labeled V_D above, was stepped from 0.25 V to 2.0 V in 0.25 V increments. I_D and V_D are referred to as I_{LSR} and V_{200} in the text. (a) Shows the $V_{\text{off}}-I_{\text{off}}$ relationship, which does not depend on V_D . The slope of this curve (for $V_{\text{off}} > 0.5 \text{ V}$) gives the value of R_{200} as $6.3 \text{ k}\Omega$. (b) Shows the output current as a function of the reference current. The dashed line has a slope of 10 representing the designed current mirror gain. The modulation current mirror is kept off during this measurement.

The current mirroring functionality of the circuit are shown in the plot of I_{LSR} vs. I_{off} . For $V_{200}=1.25$ V, the I_{off} -to- I_{LSR} current transfer characteristics are nearly ideal while I_{LSR} is below ~ 3 mA. As in the simulations, the transfer characteristics are distorted beyond this point due to the onset of gate conduction. For the selected value of $V_{200}=1$ V, the transfer characteristic is only slightly modified.

Next, the characteristics of the modulation current mirror are shown in Figure 6.21. The modulation reference voltage, V_{swg} , is swept while the off-state current mirror is kept off and the reference current, I_{swg} , and resulting laser current are monitored. The output bias voltage was again stepped from 0.25 V to 2.0 V in 0.25 V increments.

The V_{swg} - I_{swg} curve gives the value of R_{100} as 3.3 k Ω . In this case the resistance is 10% larger than the target value.

The I_{swg} -to- I_{LSR} transfer characteristics are also seen to be consistent with the simulations. I_{LSR} achieves a minimum value of around 200 μ A and I_{swg} is -60 μ A at $V_{\text{off}}=0$ V. The transfer curve becomes linear for I_{LSR} above ~ 1 mA, and $V_{200}=0.5$ V gives the best agreement with the designed current mirror gain, though operation at $V_{200}=1$ V poses no problem.

Finally, Figure 6.22 shows the laser driver output with contributions from both the off-state and modulation current mirrors. In this case V_{200} is maintained at 1 V while V_{swg} is swept and V_{off} is stepped. The range of control over the low and high laser current levels is seen to be very well matched to the task of transmitting binary data over a variety of optical systems using a VCSEL.




Figure 6.21: Static characteristics of the modulation current mirror

The reference voltage, V_{swg} , was swept while the reference current, I_{swg} , and the laser output current, labeled I_D above, were measured. The output bias voltage, labeled V_D above, was stepped from 0.25 V to 2.0 V in 0.25 V increments. I_D and V_D are referred to as I_{LSR} and V_{200} in the text. (a) Shows the $V_{\text{swg}}-I_{\text{swg}}$ relationship, which does not depend on V_D . The slope of this curve gives the value of R_{100} as 3.3 k Ω . (b) Shows the output current as a function of the reference current. The dashed line has a slope of 4 representing the designed current mirror gain. Notice that I_D does not go to zero, and that the range of I_{swg} begins at a negative value. The off-state current mirror is kept off during this measurement. The glitch occurring at $I_{\text{swg}} = 0.8 \text{ mA}$ is due to an instrumentation fault.




Figure 6.22: Combined output of off-state and modulation current mirrors

The output current, referred to as I_D above, was measured with both the off-state and modulation current mirrors active. V_{swg} was swept and V_{off} was stepped while V_{200} was maintained at 1 V. Correct addition of the off-state and modulation current is seen.

6.5 Switching Characteristics

The setup shown in Figure 6.10 was used along with the clock output from Figure 6.8 was used to characterize the turn-on and turn-off waveforms of the laser driver output. The pull-up bias voltage in Figure 6.10 was set to 1 V; the laser driver output, V_{200} , is thus biased around 0.1 V below this value, depending on the current level of the laser driver output. V_{off} was set to 1 V throughout these measurements, resulting in $I_{\text{off}} = 123 \mu\text{A}$ and contributing 1.2 mA of off-state current to the laser driver output. The switching behavior, however, was found to be independent of the off-state current. Figure 6.23 shows the measured turn-on switching characteristics on a 20 ns time scale, and Figure 6.24 gives a closer look at the first 2 ns of the step. Similarly, Figure 6.25 and Figure 6.26 show the measured turn-off characteristics. This section will analyze the qualitative features of the step responses before examining the rise and fall times qualitatively.

General observations

The amplitude of the step responses are consistent with the static measurements of the previous section. The turn on pulse shapes are seen to have a significant bias dependence while the turn-off behavior is largely independent of bias. This is consistent with the operation of the switching mechanism described in Section 6.2--the turn-off operation simply requires the discharge of node 111 through MESFET E100, while the turn-on cycle proceeds through a more elaborate switching process.

Parasitic input-output coupling

The parasitic contributions to the output measurement which were presented in Figure 6.11 should be borne in mind when viewing the current results. In particular, variations in the output which occur just prior to the output rising edge are accounted for by the parasitic coupling. Additional variation in the pulse shapes, which are significant at low modulation bias levels, also arise from parasitic coupling.

Slow turn-on at low bias

In Figure 6.23, a slow (~ 10 ns) turn-on response is observed for bias levels below $V_{\text{swg}}=0.75$ V. As discussed in Section 6.4, this corresponds to the range of bias levels where the gate current through the pass transistor rather than the reference voltage primarily determines the output current. The simulation results shown in Figure 6.27 shed light on the cause of this slow-rise characteristic. The effect is seen to originate

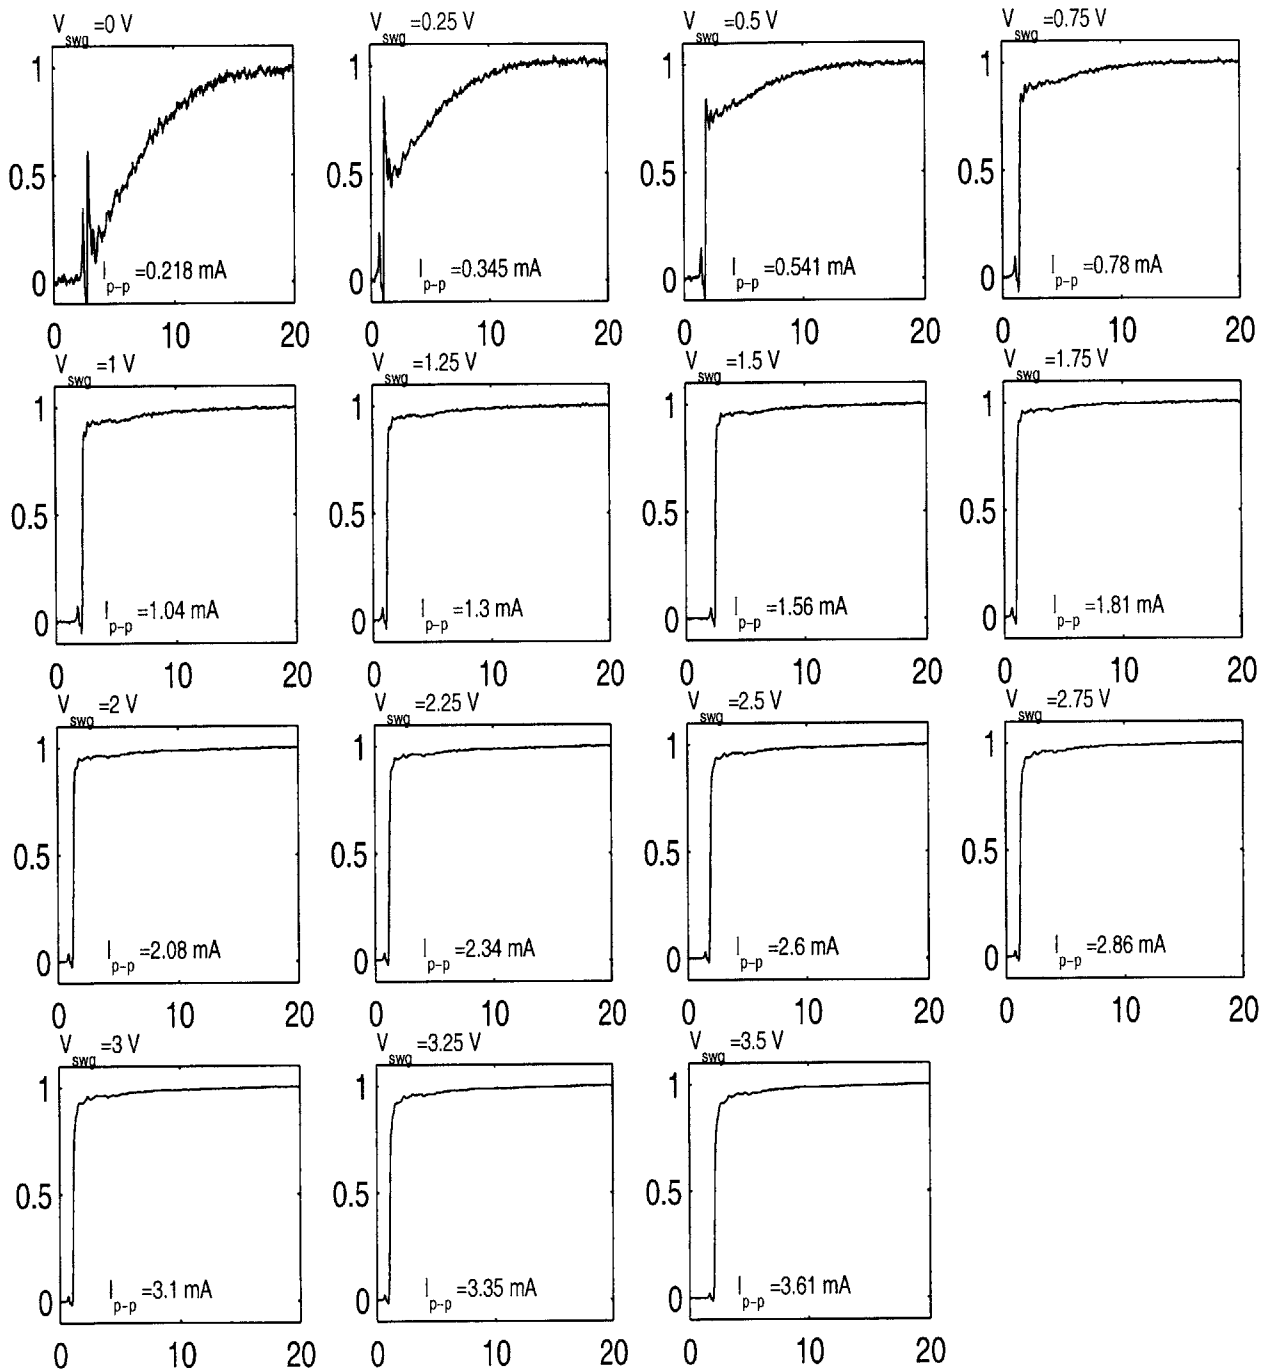


Figure 6.23: Laser driver turn-on characteristics at various modulation levels -- 20 ns time scale

Measurements were made using the setup of Figure 6.10 with the laser output pull-up bias voltage set to 1 V. The step responses have been normalized and the height has been noted. Figure 6.24 shows the initial ramp-up characteristics.




Figure 6.24: Laser driver turn-on characteristics at various modulation levels -- 2 ns time scale
 Initial ramp characteristics of step responses shown in Figure 6.23. The dotted lines correspond to the values of 0.1, 0.2, 0.8, and 0.9 on the vertical axis.

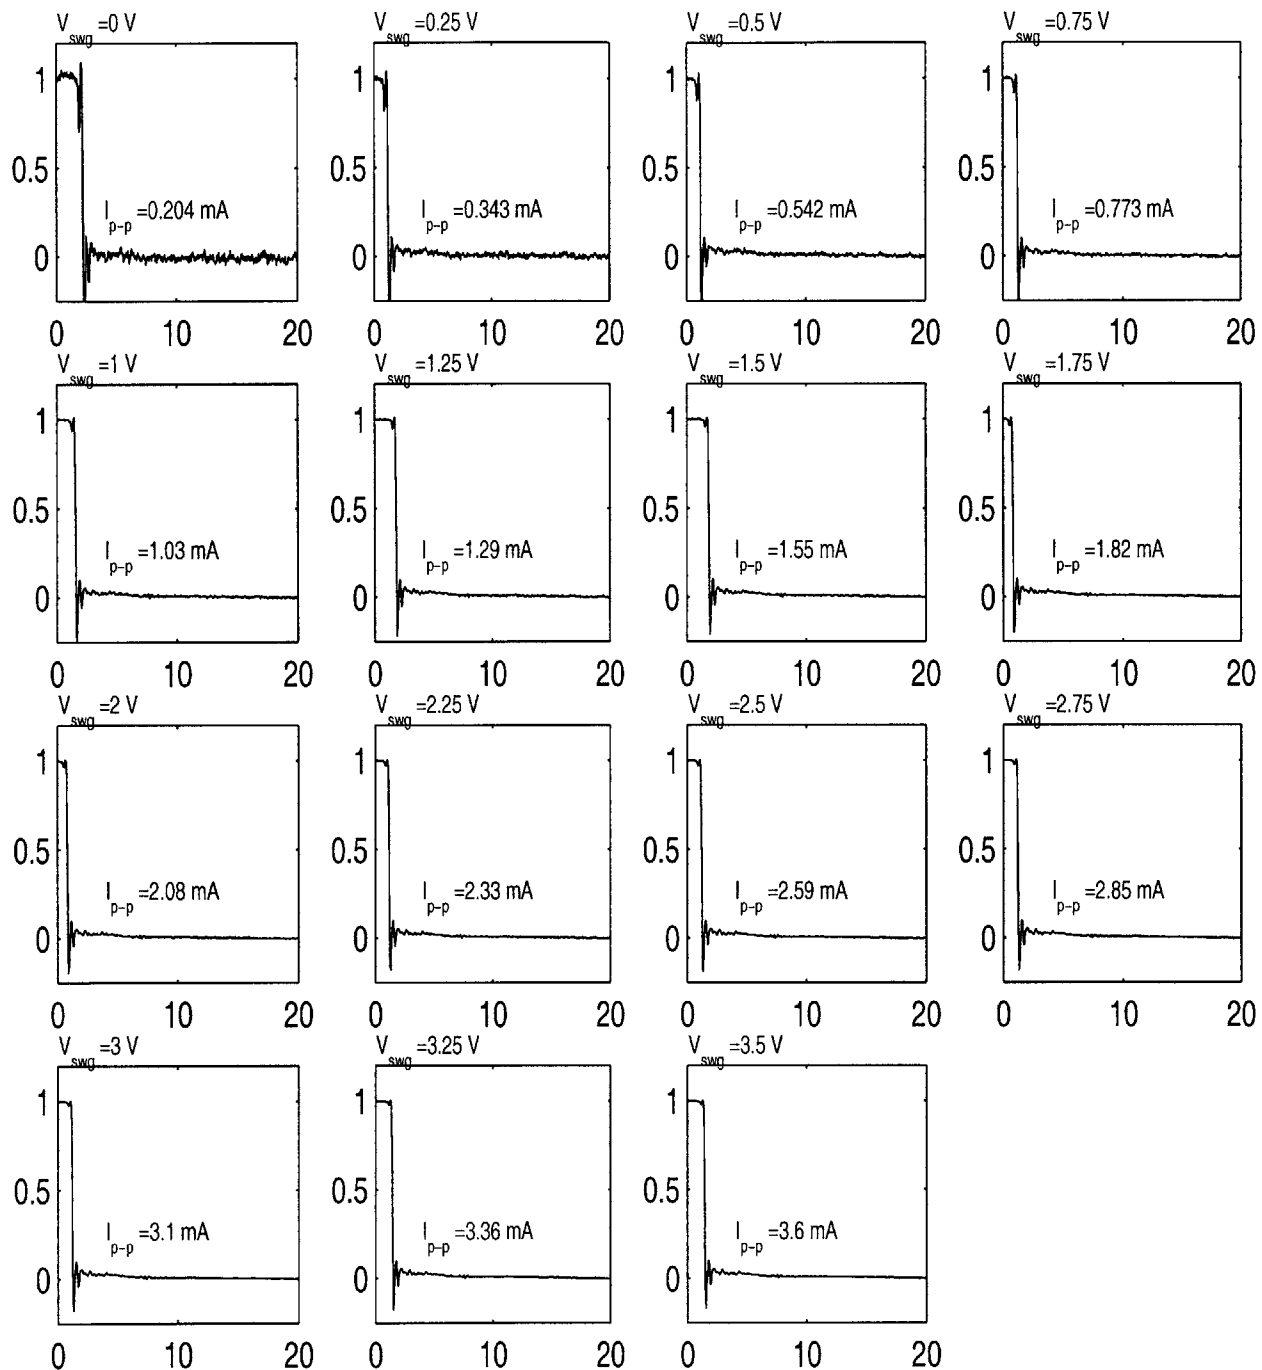


Figure 6.25: Laser driver turn-off characteristics at various modulation levels -- 20 ns time scale

Measurements were made using the setup of Figure 6.10 with the laser output pull-up bias voltage set to 1 V. The step responses have been normalized and the height has been noted. Figure 6.26 shows the initial ramp characteristics.

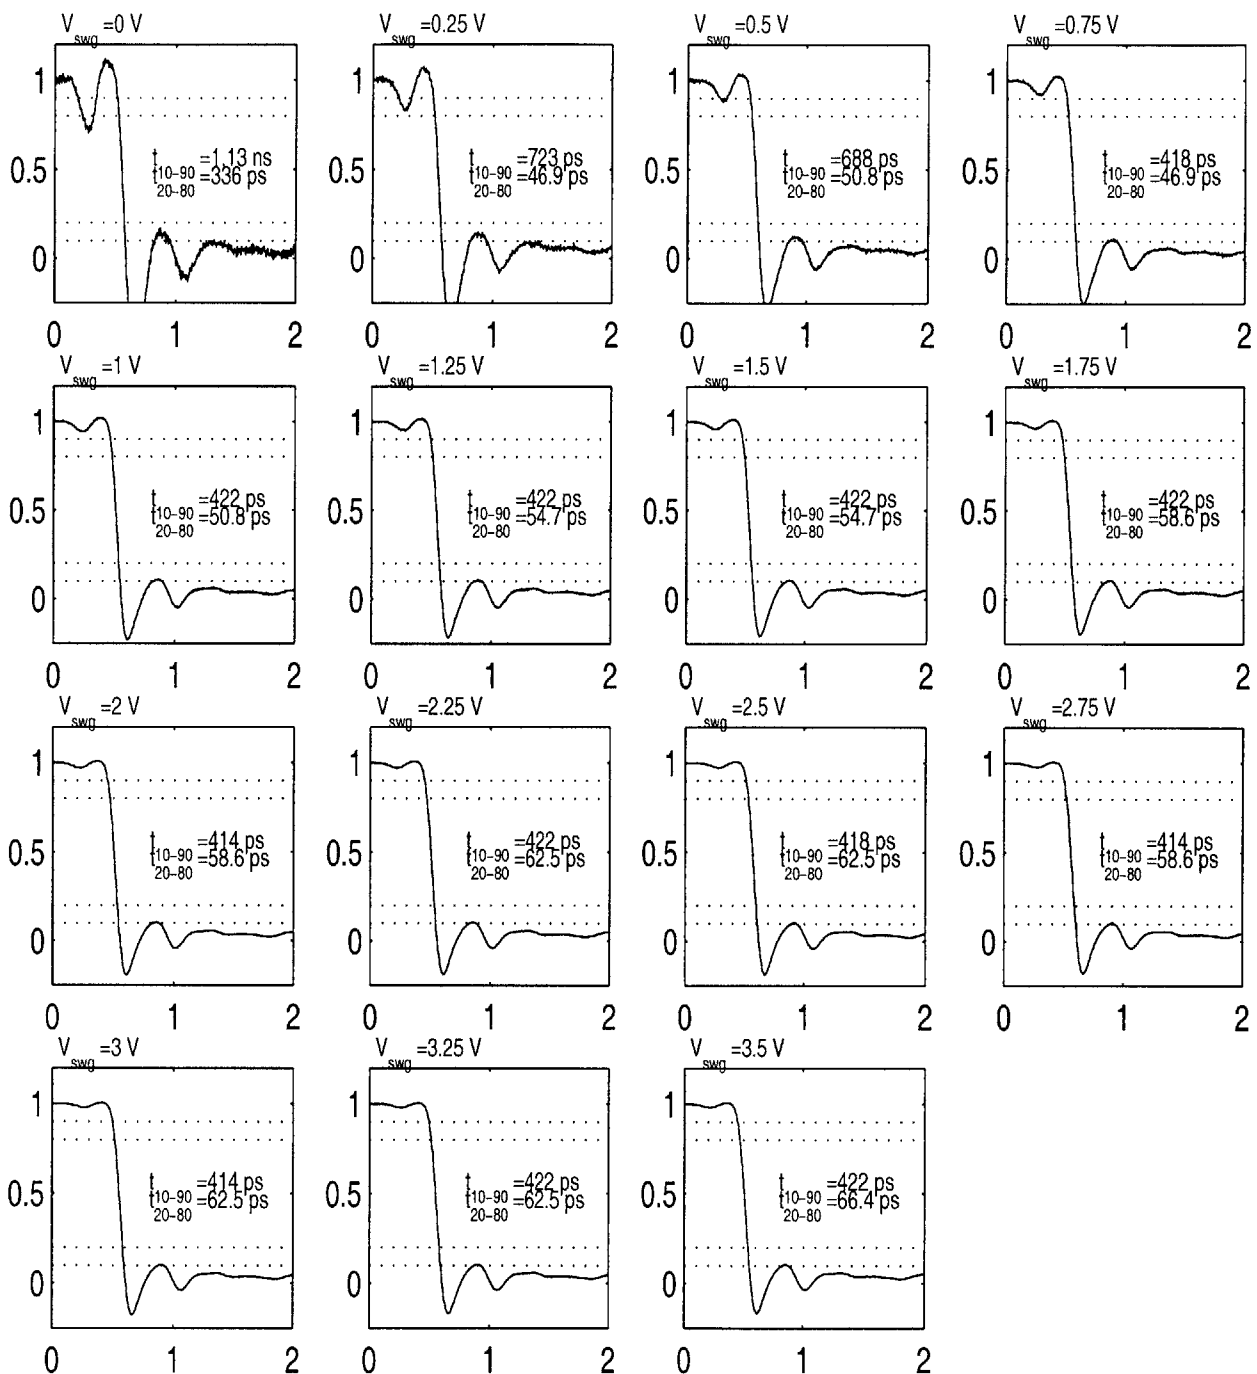


Figure 6.26: Laser driver turn-off characteristics at various modulation levels -- 2 ns time scale

Initial ramp characteristics of step responses shown in Figure 6.24. The dotted lines correspond to the values of 0.1, 0.2, 0.8, and 0.9 on the vertical axis.




Figure 6.27: Simulation of laser driver slow turn-on effect at low modulation level

In this simulation, $V_{\text{off}} = 0$ V and the off-state current mirror is turned off. A falling edge input is applied at 1 ns, and the initial conditions of the simulation have been set to reflect a degree of precharging on the floating node 103. The slow turn-on characteristics observed in Figure 6.23 are reproduced here. The origin of the effect is seen to be the slow rise of V_{101} . This is a consequence of the drain lag effect.

with response of node 101. V_{101} rises rapidly immediately after the input falling edge, but continues to rise gradually for around 50 ns. This behavior is a result of an effect known as drain lag which is described in Appendix E. Drain lag is a consequence of coupling between the drain and backgate of a MESFET through the (reverse biased) drain junction. The drain-to-backgate transfer function includes a pole at approximately 2 MHz and a zero at a somewhat higher frequency. This results in the backgate being driven by the drain at frequencies above 2 MHz. As observed from the drain, this appears as an increase in drain conductance at frequencies above 2 MHz. The inverter formed by E101 and D101 thus has a higher gain at low frequencies than at high frequencies. Its step response rises quickly to a level set by the high frequency gain, then continues to rise to the value determined by the low frequency gain on a time scale set by the low frequency pole.

The slow rise of V_{101} is seen to propagate onto nodes 102, 110, and 111. V_{111} ultimately drives the output current. The time scale for the slow rise seen in the simulation, ~50 ns, differs from the measured slow rise of ~10 ns for a couple of reasons. One reason is that the drain lag effect is over represented in Vitesses's HSPICE models [211]. An additional reason relates to the floating node, 103. The details of the switching dynamics are effected by the value of V_{103} used to initialize the simulation. The value of $V_{103} = -45$ mV reflect a typical value found on node 103 after allowing it settle for ~0.5 μ s. Deviation of this value from that of the actual circuit contributes to the discrepancy found in the simulation.

For bias levels at or above $V_{\text{swg}} = 1$ V, Figure 6.23 suggests that the slow turn-on effect is absent. In fact, close observation of the data reveals that a similar slow rise is present at all biases, but that the magnitude of the additional signal is inconsequential in comparison with the primary, fast output step.

Ringling on falling edge

In Figure 6.26, the falling edge of the laser driver output is seen to be followed by significant ringing which is not present in the rising edge data. The latter fact may imply that the ringing is intrinsic in the dynamics of the falling edge response, however, it is more likely due to packaging parasitics. This is motivated by a number of observations. First, the falling edge of the output is significantly faster than the rising edge. It thus includes more power at higher frequencies and can excite parasitic responses at these frequencies which the rising edge can not. The frequency of the ringing is seen to match that of the purely-parasitic response in Figure 6.11. The amplitude of the ringing scales with the amplitude of the output step. This is not strictly true at lower biases due to the proportionally more significant contribution from para-

sitic input coupling (Figure 6.11). Furthermore, the laser driver + laser data to be shown later in this chapter do not exhibit the ringing seen in Figure 6.26. Thus, as a practical matter, the effect in question is important only in that it obscures the measurement of the fall time of the laser driver output.

Rise and fall times

The amplitude and transition times extracted from the above data are summarized in Table 6.1. The transition times are plotted with respect to bias level in Figure 6.26. Transition time values were determined between the 10% and 90% points of the step (relative to the peak values) and between the 20% and 80% points. The utility of the two definitions will be discussed shortly.

The rise time values corresponding to bias points below $V_{\text{swg}}=0.75$ V are long due to the slow turn-on effect discussed above. The laser driver is not intended, nor is it useful, for operation in this regime. The rise time data at these biases will not be examined further.

Also, Figure 6.26 reveals that the 10%-to-90% fall time data is a consequence of the parasitic ringing at the laser driver output and does not reflect the performance of the laser driver itself. Thus, this data will similarly not be examined further.

Furthermore, the initial drop in the 10%-to-90% rise time is also seen in Figure 6.24 to be an artifact of the constant parasitic ringing superimposed on the growing laser driver output. It is evident that the edges in Figure 6.24 are consistently becoming slower with increasing bias, not initially decreasing as the extracted rise time data indicates.

Rise time vs. modulation level

The non-excluded rise time data may be summarized as follows: The 10%-to-90% rise time increase from 400 ps to 550 ps over the range of modulation levels corresponding to V_{swg} greater than 2 V; the 20%-to-80% rise time grows from 100 ps to 200 ps for V_{swg} at or above 1 V.

The significant difference between the 10%-to-90% and 20%-to-80% results from a number of factors. The first is the slow turn-on effect caused by the drain lag on node 101. As pointed out above, the slow rise of V_{101} makes some contribution to the output at all bias levels. Drain lag also effects the output FET, E111, directly. However, the drain lag initiated on the laser output, node 200, is small because of the small gain at the output resulting from the small, 25 Ω , load resistance.

V_{SWG} (V)	I_{SWG} (μA)	$\langle I_D \rangle$ (mA)	Amplitude (mA)	Rise Time 10%-90% (ps)	Rise Time 20%-80% (ps)	Fall Time 90%-10% (ps)	Fall Time 80%-20% (ps)
0.00	-34	1.26	0.21	9770	6170	1130	336
0.25	14	1.34	0.35	7190	4960	723	47
0.50	82	1.46	0.54	5040	2190	688	51
0.75	156	1.60	0.78	1360	105	418	47
1.00	232	1.75	1.04	469	117	422	51
1.25	309	1.91	1.30	441	121	422	55
1.50	386	2.06	1.56	387	125	422	55
1.75	464	2.22	1.82	223	129	422	59
2.00	542	2.37	2.08	301	133	414	59
2.25	619	2.51	2.34	402	145	422	63
2.50	696	2.66	2.60	426	152	418	63
2.75	774	2.80	2.86	445	160	414	59
3.00	852	2.94	3.10	473	176	414	63
3.25	929	3.08	3.35	496	195	422	63
3.50	1001	3.21	3.61	539	230	422	66

Table 6.1: Summary of laser driver output step characteristics

Summary of parameters extracted from step measurements in Figure 6.23-Figure 6.26. $\langle I_D \rangle$ is the average current sunk by the laser driver output from the pull-up bias supply in Figure 6.10. I_{SWG} is the modulation reference current. In these measurements, V_{off} was set to 1 V resulting in $I_{off} = 123 \mu A$ and contributing 1.2 mA of off-state current. However, the switching behavior was found to be independent of the off-state current. The amplitudes measure here are consistent with the static data of the previous section. Transition time data were measured both between the 10% and 90% points of the transition (relative to the peak-to-peak swing) and between the 20% and 80% points.

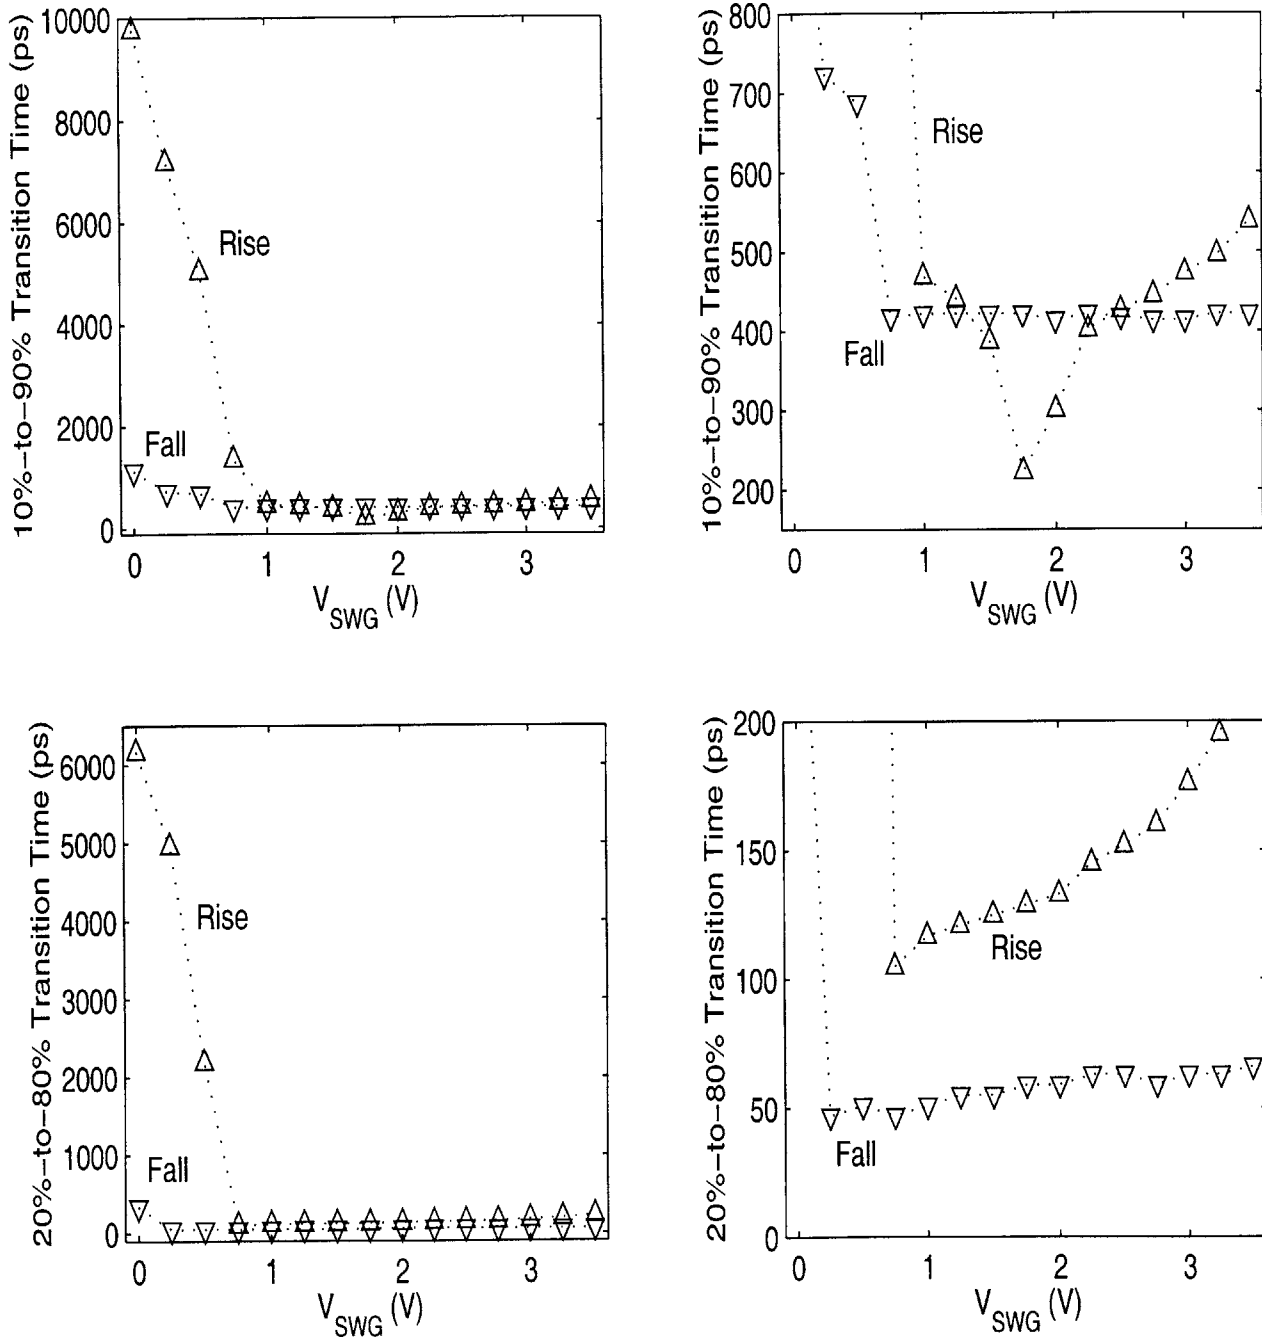


Figure 6.28: Plots of rise and fall time vs. modulation bias level

Plots are based on the data summarized in Table 6.1. (a) and (c) plot all of the 10%-to-90% and 20%-to-80% data, respectively. (b) and (d) exclude the low bias level data points which have dramatically longer transition times.

A more significant cause of the difference between the 10%-to-90% and 20%-to-80% values is the incomplete charging of the node 111 by the pass transistor gate current, I_{C1} . The amount of charge transferred via I_{C1} is determined by the dynamics of the D103/E103 feedback path and does not change to reflect the modulation bias level. Over the range of bias levels where the above 10%-to-90% data is valid, the amount of charge transferred is not adequate to bring V_{111} up to the level set at V_{110} . Node 111 must then be additionally charged through the pass transistor by current originating from R100. Like the first generation laser driver shown in Figure 6.2, this process is RC limited. In the simulation, several hundred picoseconds are required to complete the charging; this is consistent with shape of the transitions in Figure 6.24 and with the extracted rise time data.

Fall time vs. modulation level

The 20%-to-80% fall time data increases from 50 ps to 65 ps for V_{swg} above 0.25 V. The shape of the transitions in Figure 6.26 suggest that the 10%-to-90% fall times would not deviate dramatically from the 20%-to-80% values. That is, there is no extended response tail as in the rising case. The turn-off process depends on the discharging of node 111 by E100 from the on state value, which increases with bias level, down to the threshold voltage of E111. For smaller initial values of V_{111} (lower bias level) the discharge process reaches the threshold voltage sooner.

10%-to-90% and 20%-to-80% transition times

There is no unique definition for the rise and fall times. Commonly used definitions include the 10%-to-90% and 20%-to-80% transition times as well as a parameter based on the slope of the transition at its midpoint. Selecting among these conventions requires an understanding of the application for which the design is intended. The purpose of the laser driver considered here is to implement a digital interconnect in which binary data is transmitted to a receiver which generates its output by continuously comparing its input with a decision level. More accurately, the receiver may be conceptualized as an ideal comparator preceded by a linear filter representing the pulse-shaping effects of the laser, channel, photodetector, and initial receiver gain stages.

In a digital interconnect application, as opposed to a communication application, the link is required to operate with essentially no errors. This means that the transmitted signal must be designed to maintain a

high signal-to-noise ratio at the ideal comparator which is consistent with a vanishingly small bit error rate¹. Under the assumption that this critical design goal is met, the additional key link specifications are delay and delay variation, or jitter. The jitter characteristics of the laser driver itself will be studied in the next section. The receiver will also add jitter to the signal, and the jitter characteristics of the receiver will depend on the filter characteristics of the various components as well as the transition characteristics of the transmitted signal. The definition of the transition time used to characterize the laser driver output must be selected in the context of meeting the bit error rate and jitter requirements.

As examined in Appendix D, jitter is created at the input of the of the comparator by two mechanisms: additive noise and intersymbol interference (ISI). Additive noise results in jitter with a peak-to-peak value proportional to the rise/fall time of the signal as it appears near the decision level. A rise/fall time definition based on the slope of the transition at its midpoint is thus quite appropriate in addressing this type of jitter. Likewise the 20%-to-80% definition, which ignores the extended response tail of the transition, is useful in this context and is more readily extracted from the waveform measurement than the derivative of the signal. Since the 10%-to-90% definition looks at nearly the entire transition, rather than focusing on the midpoint region, it gives an overly pessimistic appraisal of the laser driver's rising edges.

ISI refers to the distortion of a waveform during a given bit period by the sum of the response tails of previous data symbols. Irrespective of its settling time, the laser driver output is ISI-free by construction². However, it is possible to produce jitter as a result of the combined effect of a long settling time at the laser driver output and the linear filtering effect of subsequent link components. Suppose, for example, that the 20%-to-80% fall time of the laser driver exceeds the response time of the linear channel. This scenario is examined in Figure 6.29. In this case, the time at which a falling edge reaches the decision level depends on the final value of the rising edge before the new falling transition. If the rising edge settling time

1. The meaning of vanishingly small depends on the bit rate and the required mean-time-to-error. For example, if the link is included in a 1 GHz computer which is to operate for at least 1 year between hardware errors, the bit error rate on the link must be significantly less than 3×10^{-17} , that is, one divided by the number of bits transmitted in one year.

2. The laser driver attempts to reset its internal state at each transition so that it is inherently free of ISI. It fails to fully achieve this goal because of drain lag on node 101 and charge storage on the floating node 103. Although this failure does result in pattern dependent jitter in the laser driver output and, to a lesser extent, in the details of the output waveform, there is no pattern dependence in the transmitted signal as a result of its rise/fall time. The same is not necessarily true of the received signal.

exceeds the bit period, the timing of the falling edge will be pattern dependent. In this context, the 10%-to-90% transition time definition is an appropriate figure of merit.

Once the receiver has recognized a transition at its input, meeting the zero bit error requirement amounts to maintaining the signal within the appropriate noise margins. A transmitted signal which fully settles immediately after transitioning into the noise margin region allows the amplitude of the transmitted signal to be minimized. By contrast, if a signal has a sharp initial edge followed by an extended response tail, its amplitude must be scaled so that the initial response completes the transition into the noise margin. The signal then continues to rise, but the additional signal is essentially “wasted.” These two cases are shown graphically in Figure 6.30. Here again, the 10%-to-90% transition time definition is useful.

Summary

It is now apparent why both the 10%-to-90% and the 20%-to-80% transition times have been determined for the laser driver, and it is appropriate at this point to appraise the laser driver’s performance. Since the 10%-to-90% transition time is below 540 ps, the laser driver is not expected to cause intersymbol




Figure 6.29: Jitter caused by combined effect of slow laser driver settling and linear channel response

In this example, the rising edge is limited by the laser drivers slow settling time, while the falling edge is limited by the response time of linear channel components. (a) shows the waveform for pattern ‘0010’ while (b) represents the pattern ‘0110’. The time at which the falling edge crosses the decision level differs in the two cases because the final value of the rising edge differs.

interference at the design bit rate of 1 Gb/s. The useful range of the laser driver in this regard is somewhere below 2 GHz.

The significance of the 20%-to-80% transition times, which remain below 200 ps, must be considered in comparison with the noise present at the receiver due to both the laser driver, the laser, and the receiver. Furthermore, this transition time represents only a lower bound on the rise/fall time of the signal at the receiver input since the transition may be degraded by the intervening components, most notable the laser and the photodetector.

Finally, while it is certainly desirable to have a short 10%-to-90% transition time, the relative system-level cost of attaining this goal must be compared to the option of increasing the amplitude on a signal having the same 20%-to-80% rise/fall time but a longer 10%-to-90% settling time.




Figure 6.30: Comparison of transition times specifications in meeting noise margin requirements

The signals in both (a) and (b) satisfy the same noise margin requirements. They also meet the same jitter requirement by having the same initial edge rate; the two signals have similar 20%-to-80% rise times. However, the amplitude of signal (b) is larger than that of signal (a). The relative cost of maintaining a short 10%-to-90% rise time, as in (a), versus allowing a long response tail on a larger signal must be considered at the system level.

6.6 Eye Patterns, Delay Asymmetry, and Jitter

This section will quantify the delay asymmetry and pattern dependent degradation of the laser driver output. Pattern dependencies may be observed by forming an eye pattern of the output signal. An eye pattern is constructed by graphically superimposing the signal from multiple bit sequences. This is carried out by making use of an oscilloscope's infinite persistence mode¹. The pseudorandom bit sequence (PRBS) generator incorporated into the module shown in Figure 6.8 allow the effects of all 23-bit patterns to be examined.

There are two primary causes for pattern dependencies in the laser driver: one is the drain lag effect on node 101, and the other is charge storage on node 103. The coupling of the output to node 210 of the off-state current mirror through the drain-to-gate capacitance of E211 also allows a linear ISI contribution from node 210. However, this effect was addressed by sizing E211, E210, and R200 to ensure that the settling time on this node is well below 1 ns. Likewise, the linear response on node 110 has been designed to avoid ISI related problems.

Note that there is no significant drain lag at the laser driver output because of the low gain afforded by the 25 Ω effective load resistance. This issue will be revisited later in this chapter. The packaging at the output of the laser driver, however, does contribute to the measurements made here through the ringing which was observed in the previous section.

Eye patterns

Figure 6.31 shows eye patterns formed at 100 Mb/s with various modulation levels. Figure 6.32 focuses in on the transitions on the eye patterns in Figure 6.31, and Figure 6.33 repeats the measurement at a bit rate of 990 Mb/s. The eye patterns are formed from the voltage output from the electrical test module shown in Figure 6.10. Since the polarity of the voltage waveform is the opposite of that of the current sunk by the laser driver, the eye patterns shown here have been inverted--the rising edge of the eye patterns seen in these figures corresponds to the turn-on cycle of the laser driver. The bias conditions used for the eye pattern measurements are identical to those used in the previous section for the switching characteristics.

1. The HP 5475A digitizing oscilloscope goes a step further by using a color code to indicate the concentration of samples at a given point in the pattern.

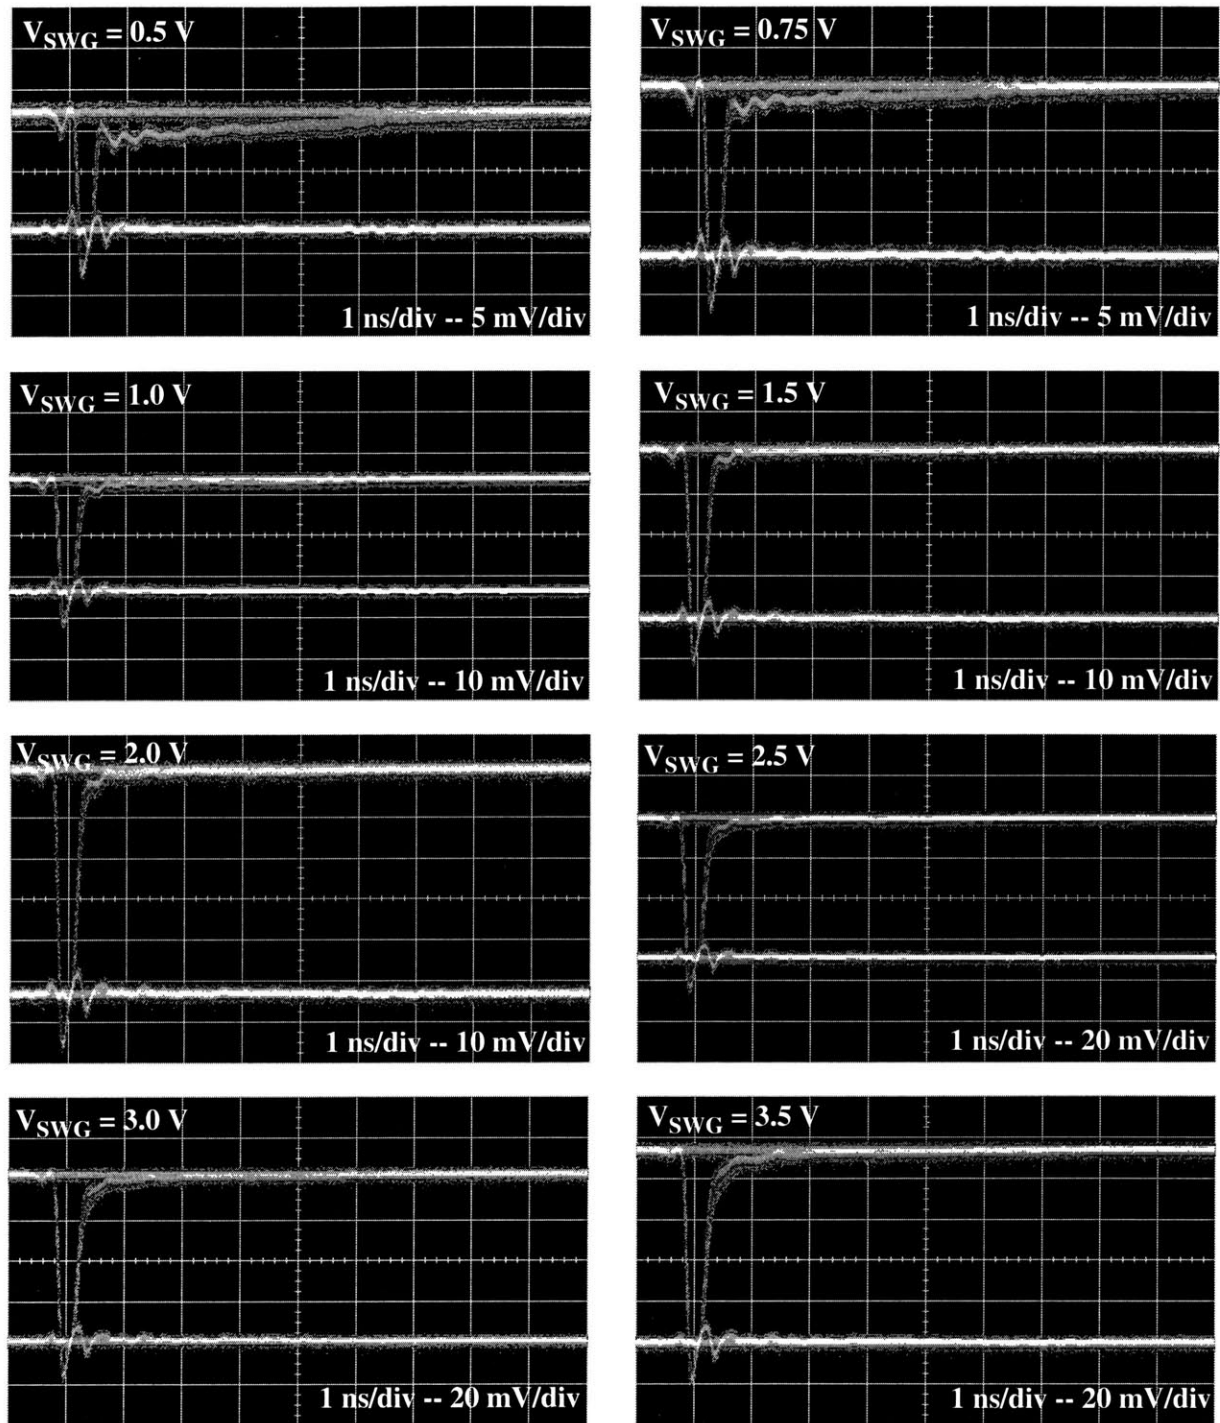


Figure 6.31: 100 MHz eye pattern of laser driver output at various modulation levels

Bias conditions are same as those given in Table 6.1. Output current is related to the above voltages by a 25Ω load. The voltage waveform has been inverted to match to polarity of the laser driver output current.




Figure 6.32: Close-up of transitions in 100 MHz eye patterns given in Figure 6.31
Eye closure is apparent at the $V_{swg}=0.5 \text{ V}$ and 0.75 V bias levels.

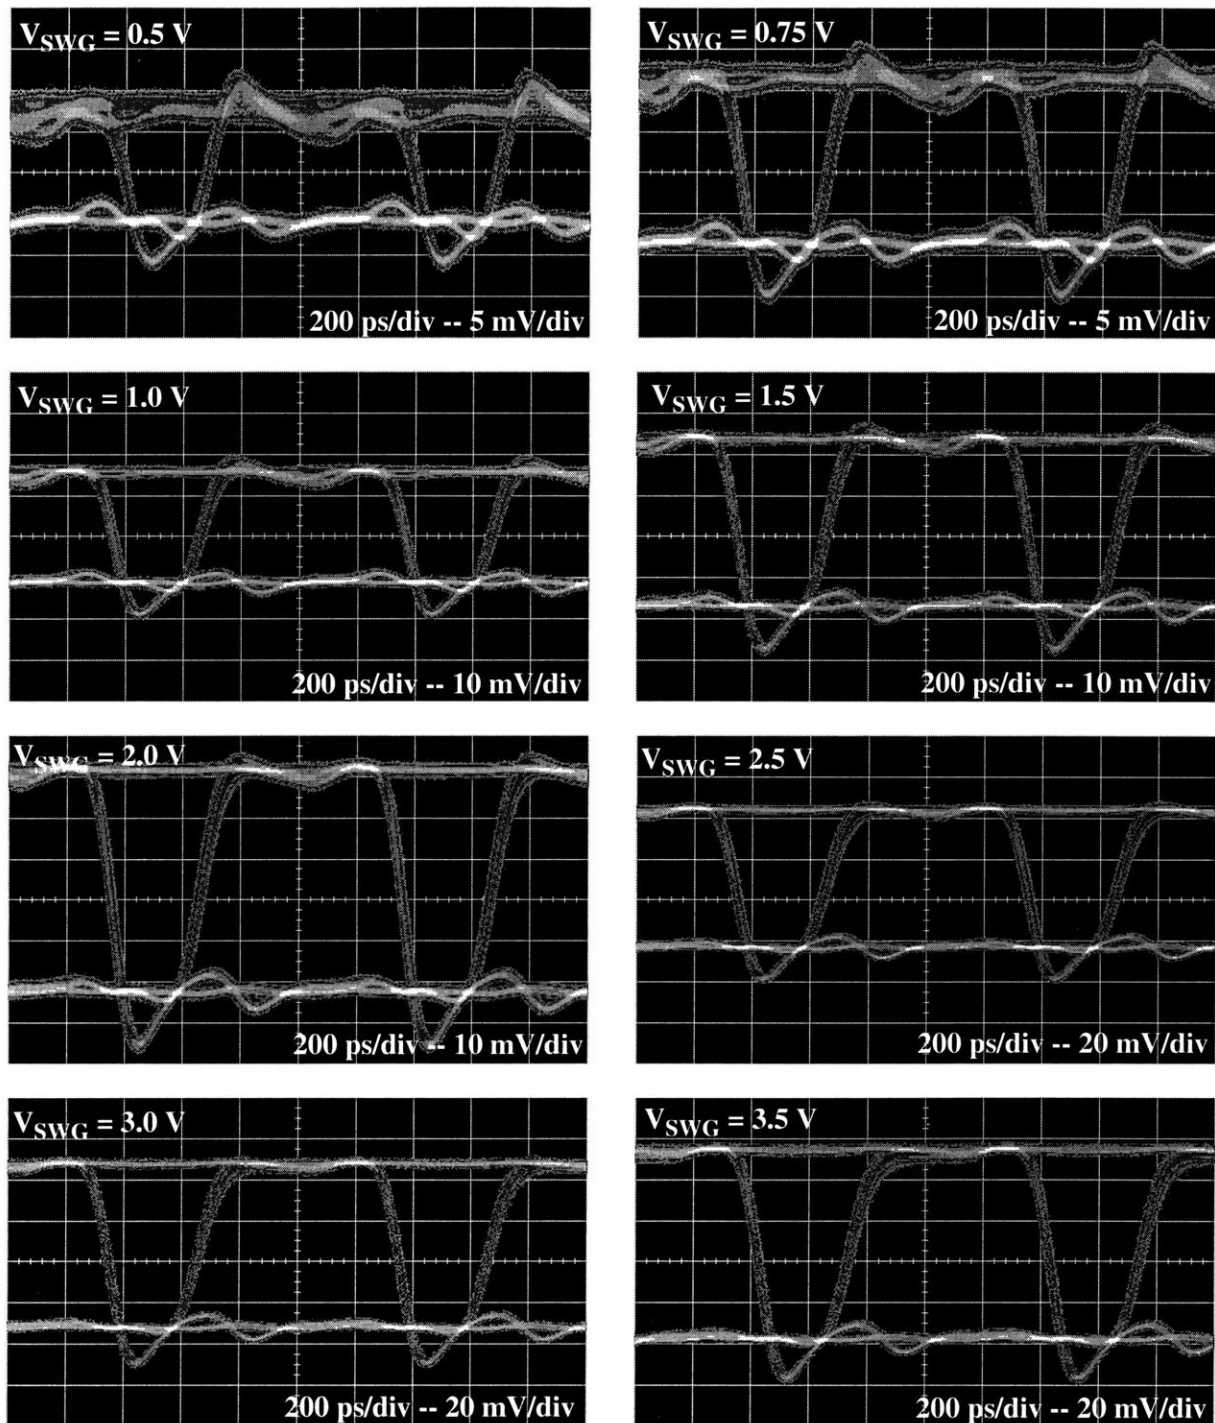


Figure 6.33: 990 MHz eye pattern of laser driver output at various modulation levels

Bias conditions are same as those given in Table 6.1. Output current is related to the above voltages by a 25Ω load. The voltage waveform has been inverted to match to polarity of the laser driver output current.

The rising and falling edges of the eye patterns are not aligned in time. The timing asymmetry varies between roughly 300 ps and 350 ps and includes contributions from both the input signal and the laser driver. As seen in Figure 6.9, an asymmetry of 185 ps exists between the transitions of the data signal applied to the laser driver. The remainder of the timing asymmetry is caused by the laser driver. This occurs because of the different signal paths involved in the turn-on and turn-off cycles. Turn-off requires only the discharge of node 111 by E100; this can be thought of as the equivalent of one inverter delay between the input and node 111. The turn-on cycle, on the other hand, must propagate through inverter E101/D101, source follower E105, and pass transistor E102 in order to reach node 111. This amounts to 2 to 3 inverter delays. In a subsequent design, this path delay difference may be equalized by adding additional delay elements between the input and E100. However, the dependence of the delay asymmetry on modulation level and bit-rate, which will be examined shortly, will mean that the delay can not be equalized under all operation conditions.

Vertical eye closure

As explained above, vertical closure of the eye pattern immediately preceding the bit transitions can result in jitter. At 100 Mb/s, the laser driver output has fully settled within one bit period. This is true even of the $V_{\text{swg}}=0.5$ V and 0.75 V modulation levels. Some eye closure may be observed as a result of ringing caused by the packaging. Some of the packaging effects originate at the module's input and do not scale with the modulation level. As a result, the eye closure is relatively more significant at lower modulation levels. The eye opening is also reduced by the presence of additive noise. This too is more significant at lower modulation levels. Note that most of the additive noise is due to the oscilloscope. With its input left open and the bandwidth limited to 12.4 GHz, the scope exhibits 2.3 mV of peak-to-peak noise. The noise distribution is well approximated as Gaussian with a standard deviation of 220 μ V. Also note that it is the additive noise on the transition edge that effects jitter; the noise should not be included in the eye closure since this would double-count its effect.

For modulation levels at or above $V_{\text{swg}}=1$ V, the eye closure characteristics at 990 Mb/s are essentially identical to those at 100 Mb/s. This confirms the assessment made in the previous section that the 10%-to-90% settling times, at these modulation levels, are short enough to support 1 Gb/s operation. In fact, if the delay asymmetry were to be eliminated, the eye patterns indicate that the laser driver could be operated at

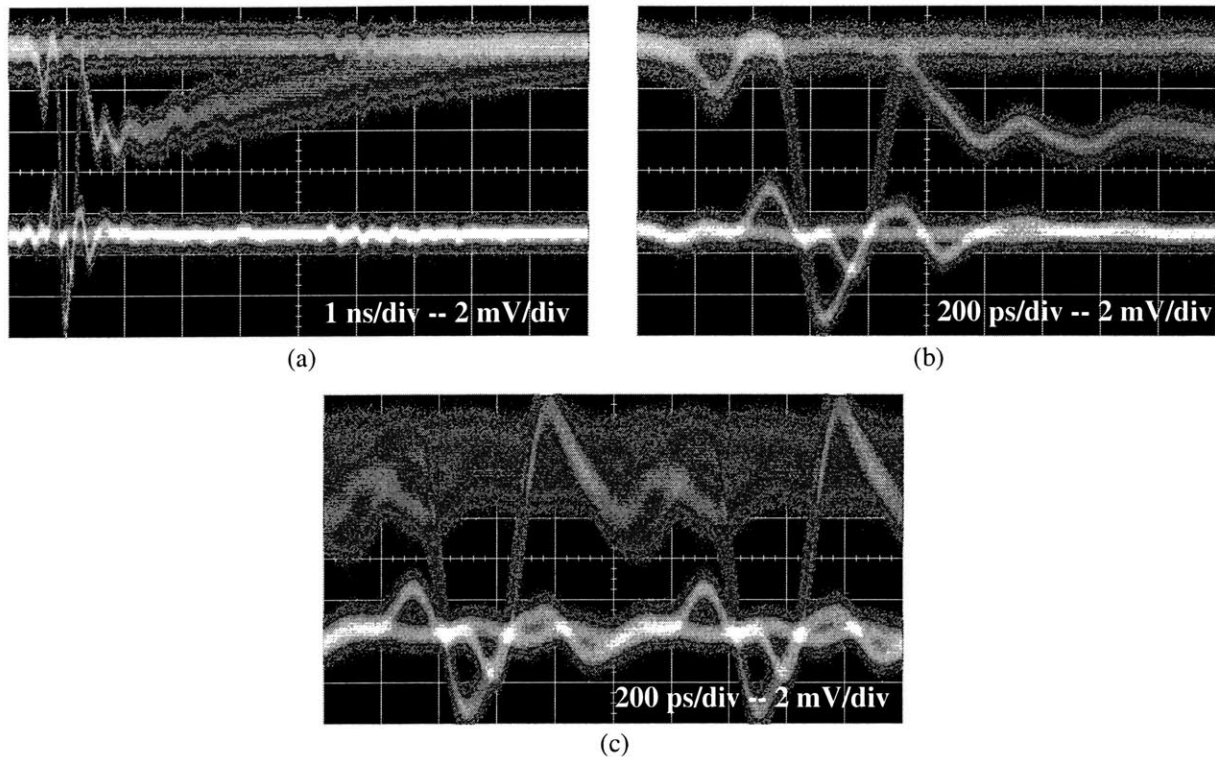


Figure 6.34: Laser output eye patterns for $V_{\text{swg}} = 0.25 \text{ V}$

(a) is the 100 MHz eye pattern. (b) shows the transitions from this pattern more closely. (c) is the 990 MHz eye pattern. Significant eye closure is observed due to the slow turn-on effect resulting from drain lag on node 101.

up to 1.7 GHz (600 ps bit period) before the settling time of the rising edge begins to reduce the eye opening.

For the $V_{\text{swg}} = 0.5 \text{ V}$ and 0.75 V modulation levels, the 990 Mb/s eye patterns include significant eye closure in the high state. This is the expected consequence of the slow turn on effect. An even more extreme case, that of $V_{\text{swg}} = 0.25 \text{ V}$, is shown in Figure 6.34

Sharper edges at higher bit-rates

Careful examination of Figure 6.32 and Figure 6.33 reveal that the rising edges are actually sharper in the 990 Mb/s eye patterns than in the 100 Mb/s patterns. The 100 Mb/s eye patterns are consistent with the turn-on waveforms of the previous section. This is to be expected since the step responses were measured at a low frequency. In particular, the waveforms rise rapidly at first, then continue to rise more slowly for

around 400 ps. The rising edges in the 990 Mb/s eye patterns, however, exhibit this slow feature to a much lesser extent, and then only at the highest modulation level.

The 400 ps feature was explained in the previous section as being caused by incomplete charging of node 111 by I_{C1} and the subsequent completion of the charging of node 111 through the pass transistor and reference resistor, R100. As discussed in Section 6.2, the amount of charge transferred via I_{C1} in the turn-on cycle is effected by the charge storage on node 103. The more negatively charged node 103 becomes, the greater the charge transfer by I_{C1} becomes, and the more complete the initial, rapid output transition becomes. Now, node 103 is driven successively lower on each turn off cycle and decays toward 0 V while in an idle state. Thus, operating the laser driver at a higher bit rate will, on average, maintain V_{103} at a more negative voltage since this allows less time between turn-off cycles for the charge to decay.

Delay asymmetry and jitter vs. modulation level

Figure 6.35 and Figure 6.36 take a more detailed look at the transitions of selected eye patterns at 100 and 990 Mb/s, respectively. The overlaid histograms represent the density of sample points within 1mV of the midpoint of the transition. The mean and width of the histograms give a measure of the delay asymmetry and jitter characteristics of the laser driver. These parameters are summarized in Table 6.2 and Table 6.3. Both the standard deviation (Std. Dev.) of the timing distribution and the difference between its extreme values (Pk.-to-Pk.) are given. These parameters were determined by the HP 5475A oscilloscope based on the histogram data. The edge-to-edge delay is the difference between the histogram means of the turn-off and turn-on edges. Corresponding data is reported for the input signal from Figure 6.8 which is used to form the eye patterns.

The falling edge of the data signal in Figure 6.9 appears later in time than the rising edge. This falling data edge triggers the laser driver's turn-on cycle. Since the propagation delay of the laser driver is greater for the turn on cycle, the edge-to-edge delay is seen to grow in the laser driver output waveforms. The delay asymmetry of the laser driver can be recovered by subtracting the edge-to-edge delay of the data signal from that of the laser driver output. The resulting delay asymmetry is plotted in Figure 6.37. The bit rate is seen to have an important effect on the delay asymmetry; this will be examined below. The delay asymmetry is also seen to rise with modulation level. It is not clear, however, if this is simply a consequence of the increase in rise time with modulation level, as seen in Figure 6.28, or if there is also a modulation level dependence internal to the laser driver.




Figure 6.35: Edge timing vs. modulation level at 100 Mb/s

Detailed look at the transitions of selected eye patterns from Figure 6.31. The overlaid histograms represent the density of sample points within 1 mV of the midpoint of the transition. 'INPUT' is the data signal used to form the eye patterns. The laser driver output waveforms shown here are inverted relative to the eye patterns of Figure 6.31--the rising edge here corresponds to the turn-off cycle of the laser driver. The mean and width of the histograms give a measure of the delay asymmetry and jitter characteristics of the laser driver. These parameters are summarized in Table 6.2.

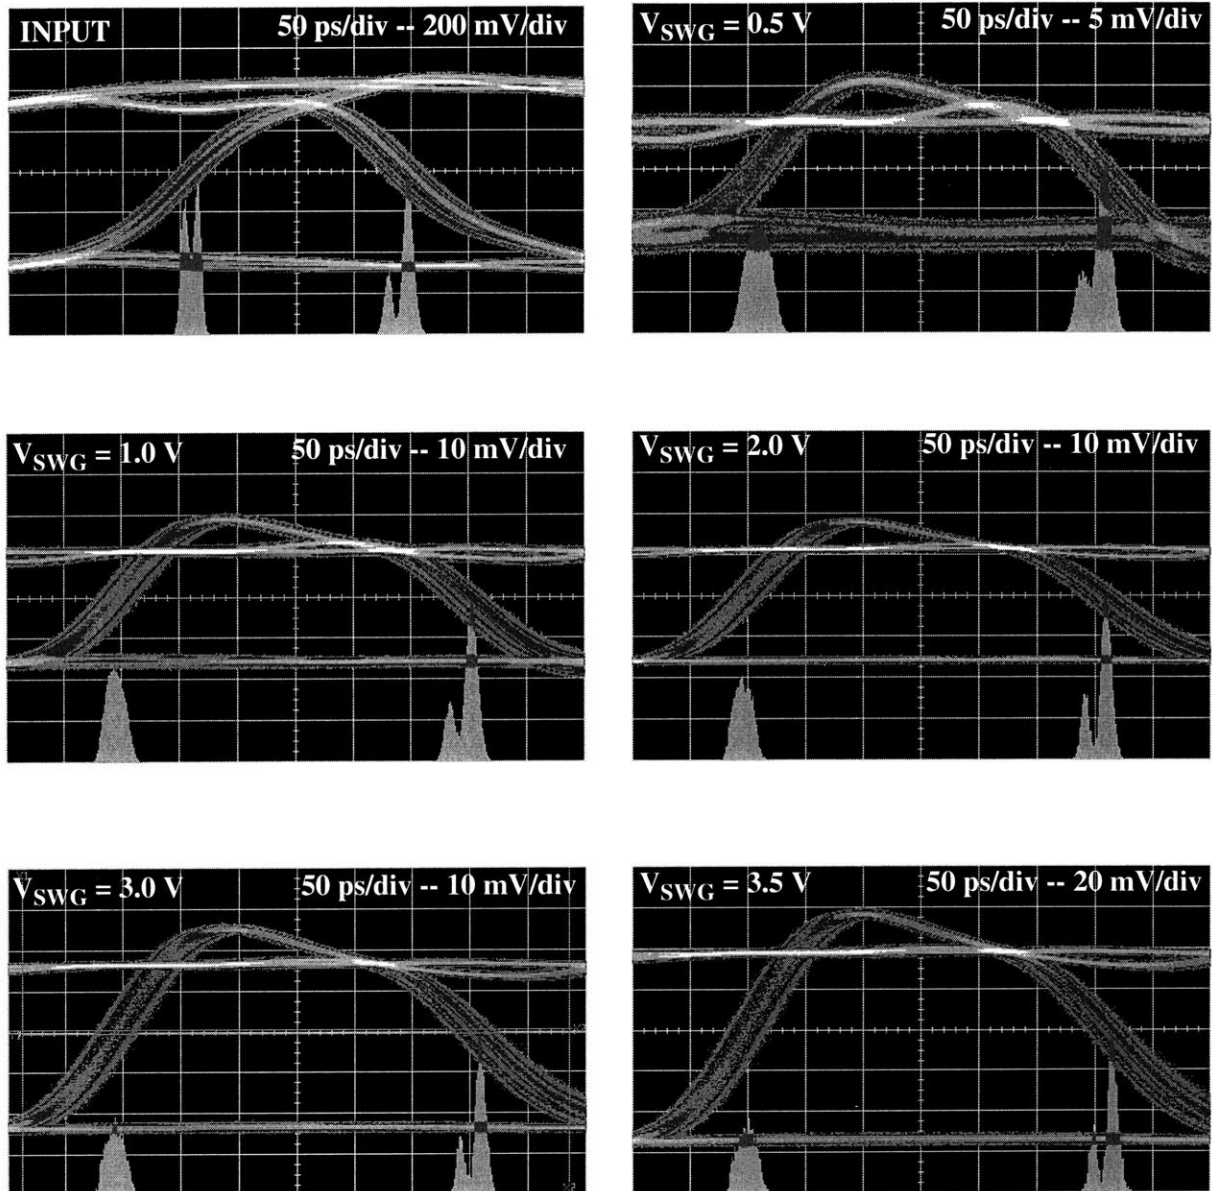


Figure 6.36: Edge timing vs. modulation level at 990 Mb/s

Detailed look at the transitions of selected eye patterns from Figure 6.32. The overlaid histograms represent the density of sample points within 1mV of the midpoint of the transition. 'INPUT' is the data signal used to form the eye patterns. The laser driver output waveforms shown here are inverted relative to the eye patterns of Figure 6.32--the rising edge here corresponds to the turn-off cycle of the laser driver. The mean and width of the histograms give a measure of the delay asymmetry and jitter characteristics of the laser driver. These parameters are summarized in Table 6.3. The bimodal distributions in the INPUT waveform are a result of ISI. This jitter pattern is propagated by the laser driver, but is obscured in some of the turn-off edges.

V_{SWG} (V)	Rising Jitter Std. Dev (ps)	Rising Jitter Pk.-to-Pk. (ps)	Falling Jitter Std. Dev. (ps)	Falling Jitter Pk.-to-Pk. (ps)	Delay Edge-to-Edge (ps)
0.5	11.3	93.3	8.9	65.6	328
1.0	11.2	83.3	8.5	63.3	346
2.0	11.2	81.1	8.5	64.4	355
3.0	12.3	92.2	8.5	60.0	364
3.5	13.4	107.8	8.4	64.4	368
INPUT	7.7	56.7	7.1	51.1	185

Table 6.2: Summary of jitter and delay asymmetry characteristics at 100 Mb/s

Values extracted from Figure 6.35. The rising and falling edge refer to the turn-on and turn-off phases of the laser driver, respectively. The jitter parameters were determined by the HP 5475A oscilloscope based on the histogram data. Both the standard deviation (Std. Dev.) of the distribution and the difference between its extreme values (Pk.-to-Pk.) are given. The edge-to-edge delay is the difference between the histogram means of the turn-off and turn-on edges. 'INPUT' refers to the data signal from Figure 6.8 which was used to form the eye patterns.

Appendix D examines the propagation of jitter from the input signal to the laser driver output. A simple model gives the standard deviation and peak-to-peak jitters as

$$\sigma_O^2 = \sigma_I^2 + \frac{\tau^2}{SNR} + \sigma_D^2 \quad (6.1)$$

$$\Delta T_O = \Delta T_I + \frac{\Delta V}{A} \tau + \Delta T_D \quad (6.2)$$

where σ_O is the standard deviation of the output jitter, σ_I is the standard deviation of the input jitter, τ is the rise/fall time, SNR is the signal-to-noise ratio at the output, σ_D is the standard deviation of the delay through the laser driver, ΔT_O is the output peak-to-peak jitter, ΔT_I is the input peak-to-peak jitter, $\Delta V/A$ is the peak-to-peak output noise normalized to the peak-to-peak amplitude of the output, and ΔT_D is peak-to-peak variation in the delay through the laser driver.

Since the jitter variances and the peak-to-peak jitter add, the jitter contribution from the input signal may be removed from the data. However, due to the bimodal timing distributions in Figure 6.36, which

V_{SWG} (V)	Rising Jitter Std. Dev (ps)	Rising Jitter Pk.-to-Pk. (ps)	Falling Jitter Std. Dev. (ps)	Falling Jitter Pk.-to-Pk. (ps)	Delay Edge-to-Edge (ps)
0.5	10.1	62.2	9.8	58.9	295
1.0	9.5	51.1	7.6	45.6	303
2.0	8.9	50.8	7.5	45.6	308
3.0	8.6	50.0	7.7	44.4	313
3.5	8.5	52.2	7.9	44.4	312
INPUT	8.5	43.3	6.2	31.1	179

Table 6.3: Summary of jitter and delay asymmetry characteristics at 990Mb/s

Values extracted from Figure 6.36. The rising and falling edge refer to the turn-on and turn-off phases of the laser driver, respectively. 'INPUT' refers to the data signal from Figure 6.8 which is used to form the eye patterns. The jitter parameters were determined by the HP 5475A oscilloscope based on the histogram data. Both the standard deviation (Std. Dev.) of the distribution and the difference between its extreme values (Pk.-to-Pk.) are given. Because of the presence of bimodal distributions in the INPUT timing, the standard deviation is not a useful measure of the distribution widths. The edge-to-edge delay is the difference between the histogram means of the turn-off and turn-on edges.

result from ISI in the input signal path, the 990 Mb/s standard deviation jitter data is not useful; jitter variances no longer add since the timing distribution is not well approximated as Gaussian. The peak-to-peak value of the jitter gives consistent results independent of the timing distribution. The second terms in (6.1) and (6.2) are due to additive noise at the output. It would be desirable to remove the effects of the noise added by the oscilloscope, but this is not readily achieved since the relative contribution of the laser driver to the additive noise is not known.

Figure 6.38 plots the output jitter data after removing the contribution from the input signal. The falling edge jitter is seen to be independent of bias level and frequency with the exception of $V_{swg}=0.5$ V data points. At this bias point, the drain lag effect on node 101 causes pattern dependent jitter. This effect also applies to the rising edge data. At higher modulation levels, the 100 Mb/s rising edge jitter data increases with modulation level while the 990 Mb/s data remains constant at a lower level (again, the 990 Mb/s standard deviation jitter data is not admissible). This difference is due to the dependence of the rising edge




Figure 6.37: Delay asymmetry vs. modulation level

The laser driver's delay asymmetry is determined from the data of Table 6.2 and Table 6.3 by subtracting the edge-to-edge delay of the input signal from that of the laser driver output. The bit rate is seen to have an important effect on the delay asymmetry; this will be examined in Figure 6.41. Here, the delay asymmetry is seen to rise with modulation level. It is not clear if this is simply a consequence of the increase in rise time with modulation level, or if there is also a modulation level dependence internal to the laser driver.

waveforms on bit rate. As discussed above, the 990 Mb/s eye patterns have sharper rising edges, and maintain this pattern up to the highest modulation level. The jitter data is consistent in this regard, rising slightly at $V_{\text{SWG}}=3.5$ V. The 100 Mb/s rising edges, on the other hand, become slower with increasing modulation level, and this leads to a larger measured jitter through the second (additive noise) terms of (6.1) and (6.2).

With the exception of the 100 Mb/s rising edge data, which is obscured by an additive noise contribution, the jitter data indicate that the propagation delay of the laser driver is independent of modulation level. The falling edge peak-to-peak jitter is around 15 ps, and the rising edge jitter is between 10 and 50 ps.

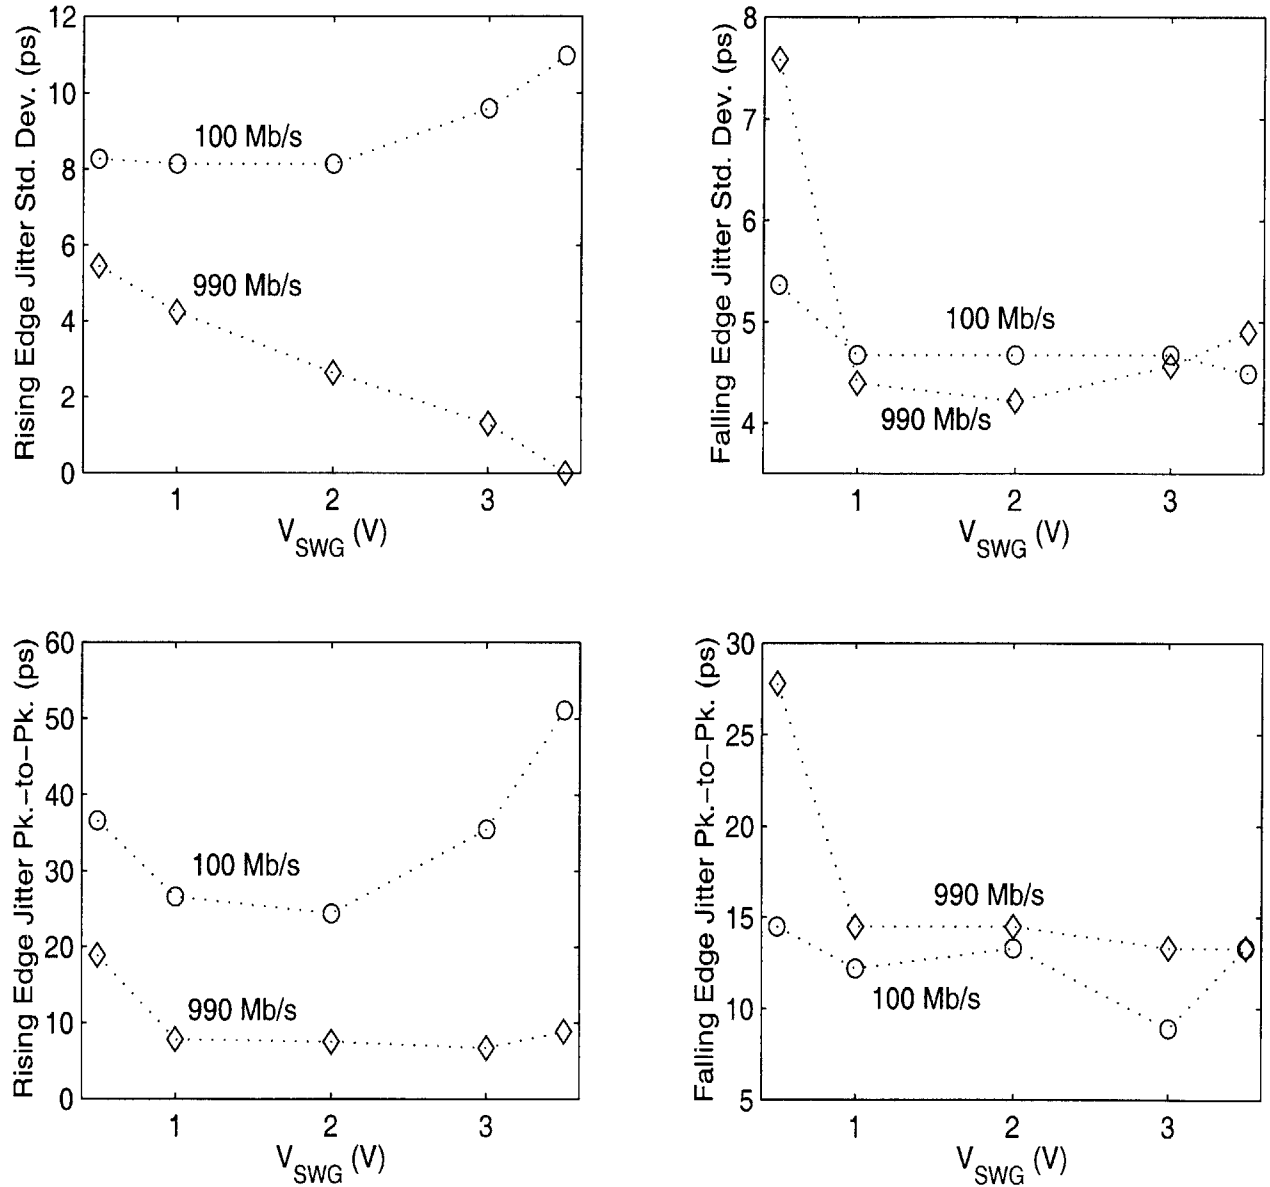


Figure 6.38: Jitter vs. modulation level

Based on the data of Table 6.2 and Table 6.3, the standard deviation jitter is determined as the root-difference-square of the output and input standard deviation values (this assumes that the jitter variances add, as developed in Appendix D). The peak-to-peak jitter is the difference between the input and output peak-to-peak values. Due to the bimodal timing distributions in Figure 6.36, the 990 Mb/s standard deviation jitter data is not admissible. The falling edge jitter is seen to be independent of bias level and frequency with the exception of $V_{SWG}=0.5$ V data points. At this bias, the drain lag induced slow turn-on effect causes pattern dependent jitter. This effect also applies to the rising edge data. At higher modulation levels, the difference in the 100 and 990 Mb/s jitter data is attributable to the dependence of the rising edge waveforms on bit rate.

Jitter vs. bit-rate

Figure 6.35 and Figure 6.36 revealed a frequency dependence to the laser driver's delay asymmetry and rising edge jitter. The following data looks at these trends, focusing on a single modulation level corresponding to $V_{\text{swg}}=1.5$ V. Figure 6.39 looks at the input signal while Figure 6.40 details the transitions in the laser driver output. Corresponding measurements are summarized in Table 6.4 and Table 6.5. As before, the delay asymmetry is determined as the difference between the output and input signals' edge-to-edge delays and is plotted in Figure 6.41. Likewise, Figure 6.42 plots the jitter characteristics after removing the input signal's jitter contribution.

The delay asymmetry is found as the difference between output and input edge-to-edge delays. As previously seen in Figure 6.37, the delay asymmetry in Figure 6.41 decreases with bit rate. The best explanation for this, once again, is the bit-rate dependence of the charge storage on the floating node 103. The fact node that 103 is more negatively charged at higher bit rates not only improves the characteristics of the rising edge, it also leads to more rapid signal propagation through the turn-on cycle switching circuit.

The jitter characteristics shown in Figure 6.42 are a further result of the bit-rate dependence of the charge on node 103. As expected from the circuit and seen previously in Figure 6.38, the falling edge jitters do not have a strong frequency dependence. (The falling edge peak-to-peak data point at 50 Mb/s is inconsistent with the standard deviation figure and appears to be a measurement error.) Like the delay asymmetry, the rising edge jitter becomes smaller as the bit rate is increased. The jitter is reduced by both the reduction in the internal jitter, which is expected to accompany the reduced turn-on propagation delay, and the increased edge rate, which reduces the influence of additive noise. As already discussed, these effects are related to the floating node issue.

Summary

The static, transient, and timing measurements support the assessment that the laser driver is well suited to operation in a 1 Gb/s digital interconnect. The primary target for improvement is the elimination of the floating node. This design flaw has led to a variation in delay asymmetry, making it difficult to correct for the asymmetry by the inclusion of delay elements in the turn-off path. The floating node has also resulted in propagation delay variations, although the magnitude of the resulting jitter is probably acceptable. More critical, however, is the accompanying degradation of the rising edge waveform which reduces the design margins allowed in the other link components.

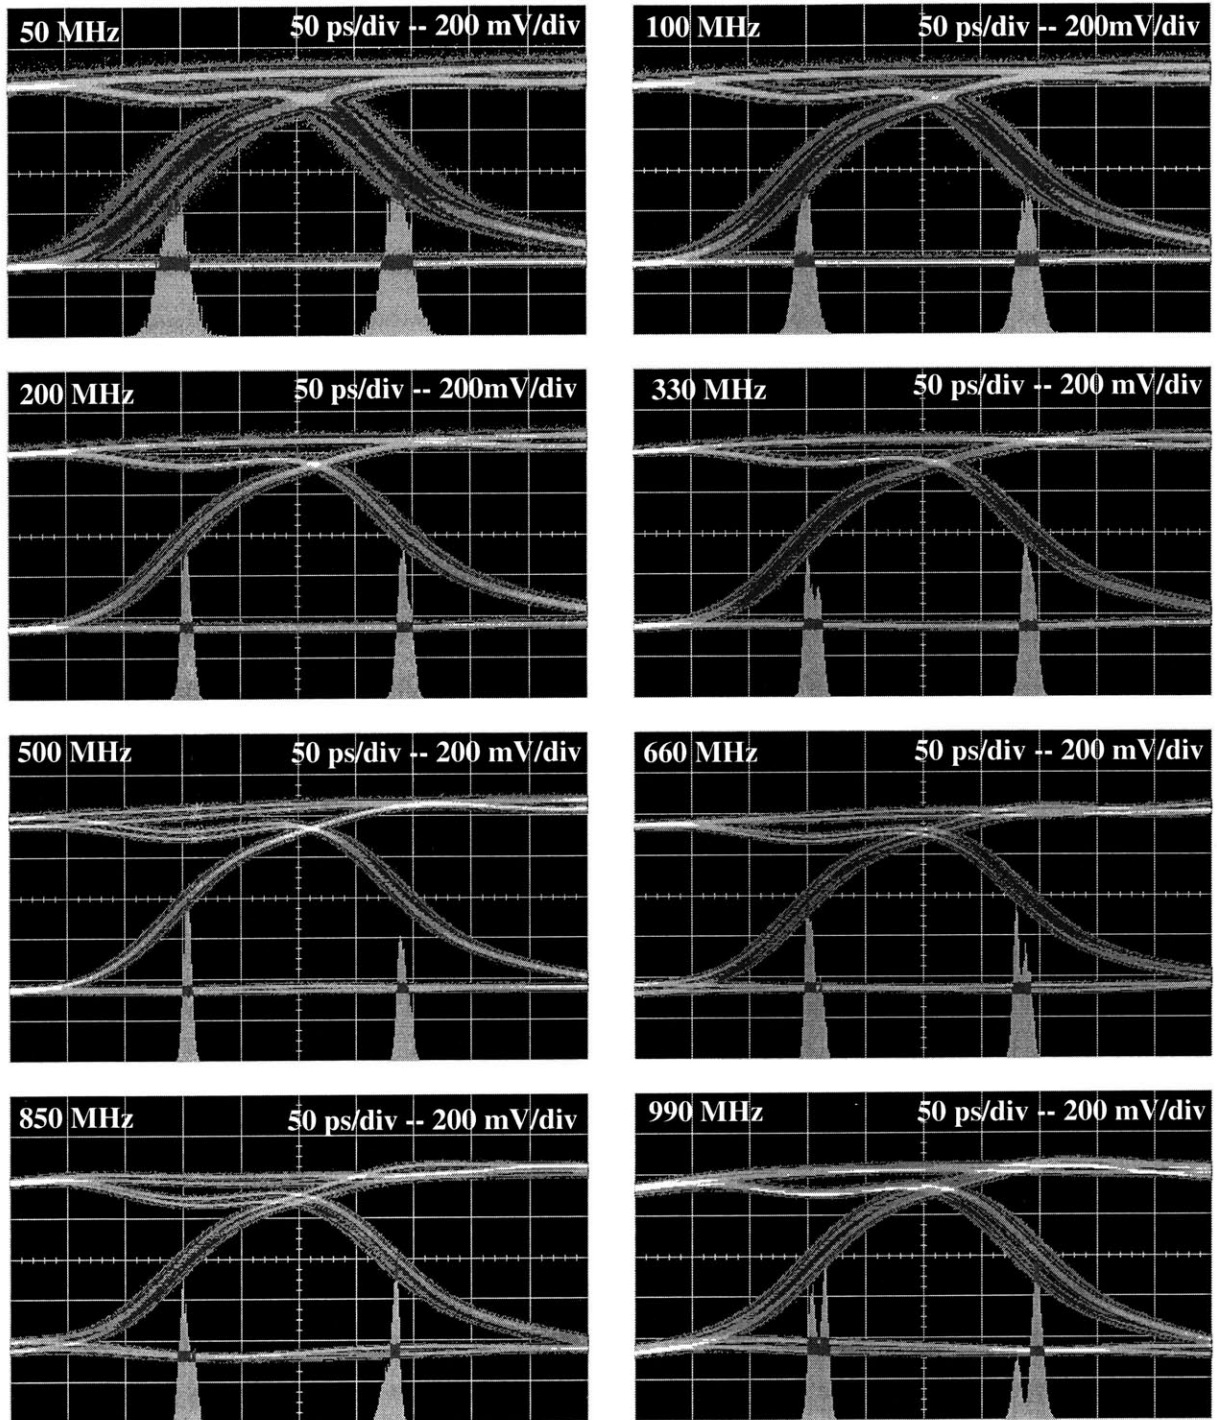


Figure 6.39: Input signal edge timing vs. bit-rate

The mean and width of the overlaid histograms are used to determine the edge-to-edge delay and jitter characteristics. The histograms indicate the crossing time of the waveform through a 1mV window at the transition midpoint.

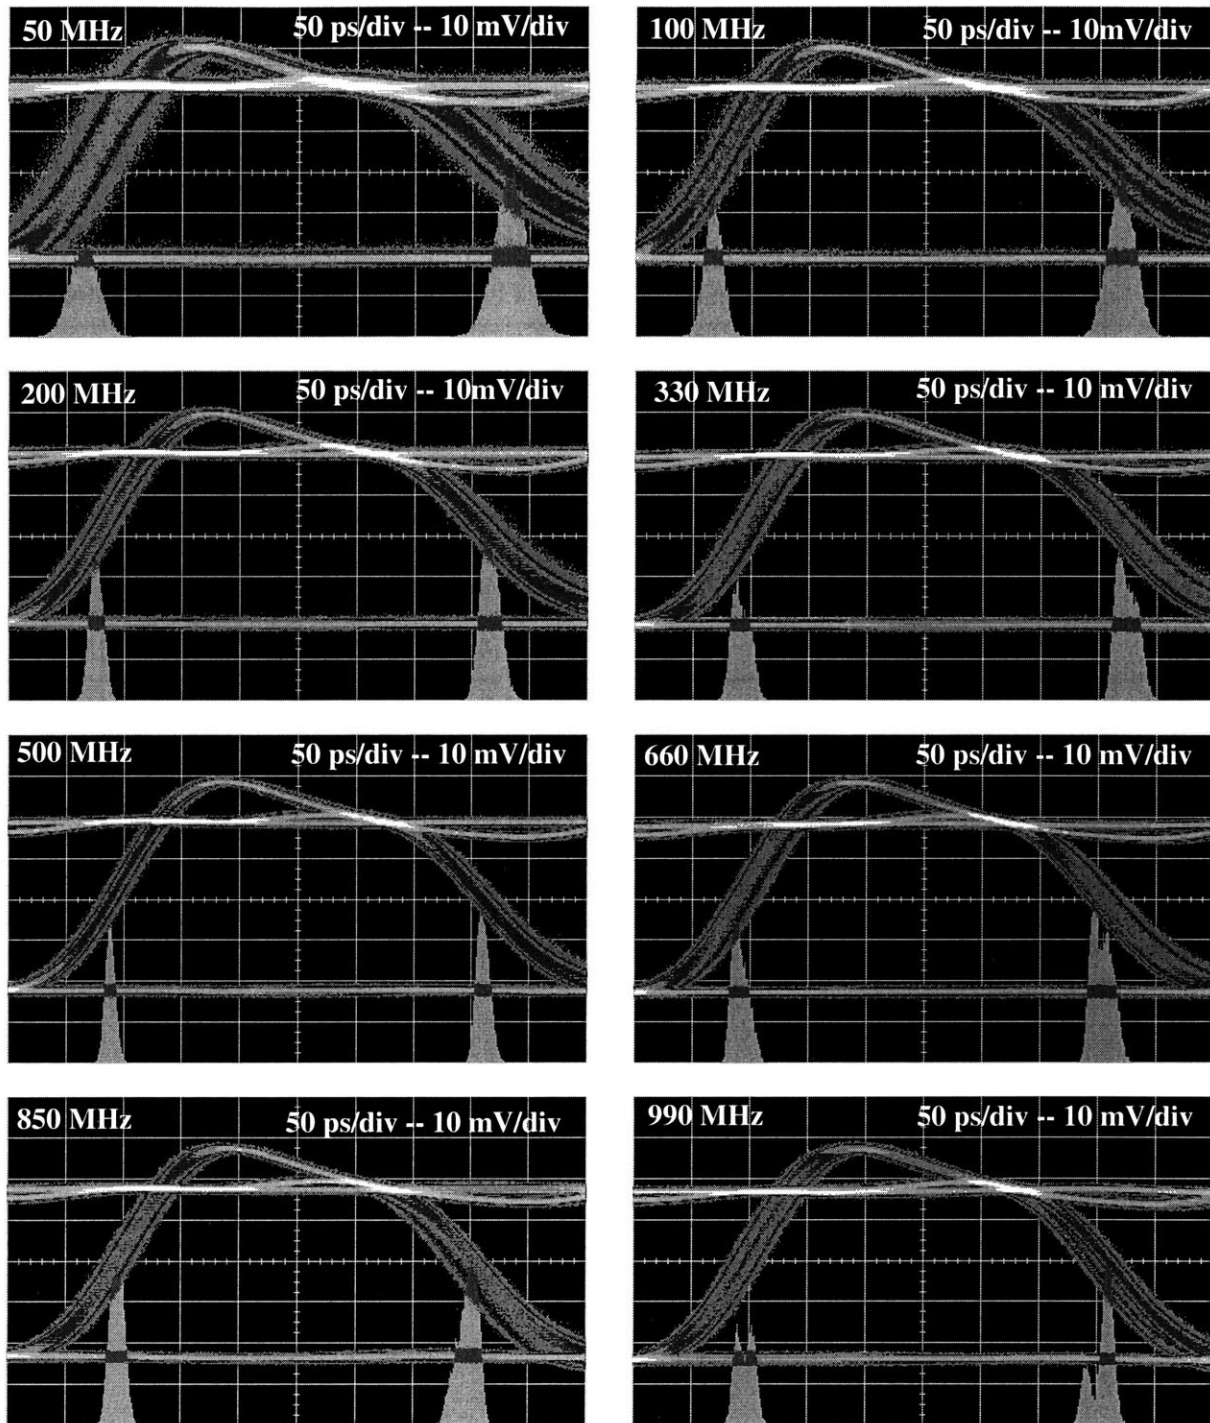


Figure 6.40: Edge timing vs. bit-rate with $V_{swg}=1.5\text{ V}$

These waveforms are inverted relative to the eye diagrams of Figure 6.31-Figure 6.33. The histograms indicate the crossing time of the waveform through a 1 mV window at the transition midpoint.

F_{CLK} (MHz)	T_{CLK} (ns)	Rising Jitter Std. Dev (ps)	Rising Jitter Pk.-to-Pk. (ps)	Falling Jitter Std. Dev. (ps)	Falling Jitter Pk.-to-Pk. (ps)	Delay Edge-to-Edge (ps)
50	20	11.0	81.1	11.7	87.8	194
100	10	8.2	65.6	7.3	53.3	191
200	5	5.3	36.7	4.3	34.4	189
330	3	5.3	34.4	5.9	35.6	186
500	2	4.7	28.9	3.4	23.3	188
660	1.5	5.6	28.9	5.8	32.2	179
850	1.2	5.6	32.3	5.4	31.1	179
990	1	8.6	45.6	6.3	32.2	186

Table 6.4: Summary of input jitter and delay asymmetry characteristics at various frequencies

Summary of data extracted from Figure 6.39. The input rising and falling entries refer to the turn-on and turn-off cycles of the laser driver, that is, the entries have been exchanged in order to match the corresponding edge of the output. Peak-to-peak and standard deviation jitters were determined by the HP 5475A oscilloscope on the basis of histogram data. The peak-to-peak value is the difference between the highest and lowest time entries in the histogram. Edge-to-edge delay is the difference between the histogram means.

F_{CLK} (MHz)	T_{CLK} (ns)	Rising Jitter Std. Dev (ps)	Rising Jitter Pk.-to-Pk. (ps)	Falling Jitter Std. Dev. (ps)	Falling Jitter Pk.-to-Pk. (ps)	Delay Edge-to-Edge (ps)
50	20	14.1	111.1	12.5	108.9	369
100	10	11.4	86.7	7.9	64.4	353
200	5	8.4	68.9	5.9	46.7	341
330	3	9.0	61.1	7.7	47.8	333
500	2	5.7	37.8	4.4	33.3	323
660	1.5	8.6	47.8	7.0	42.2	313
850	1.2	8.9	52.2	6.2	42.2	303
990	1	9.1	52.2	7.6	45.6	309

Table 6.5: Summary of output jitter and delay asymmetry characteristics at various frequencies

Summary of data extracted from Figure 6.40. The input rising and falling entries refer to the turn-on and turn-off cycles of the laser driver. Peak-to-peak and standard deviation jitters were determined by the HP 5475A oscilloscope on the basis of histogram data. The peak-to-peak value is the difference between the highest and lowest time entries in the histogram. Edge-to-edge delay is the difference between the histogram means.




Figure 6.41: Delay asymmetry vs. bit rate for $V_{swg}=1.5$ V

Based on Table 6.4 and Table 6.5, the delay asymmetry is found as the difference between output and input edge-to-edge delays. The trend is consistent with Figure 6.37 and can be explained by the frequency dependence of the stored charge on floating node 103.

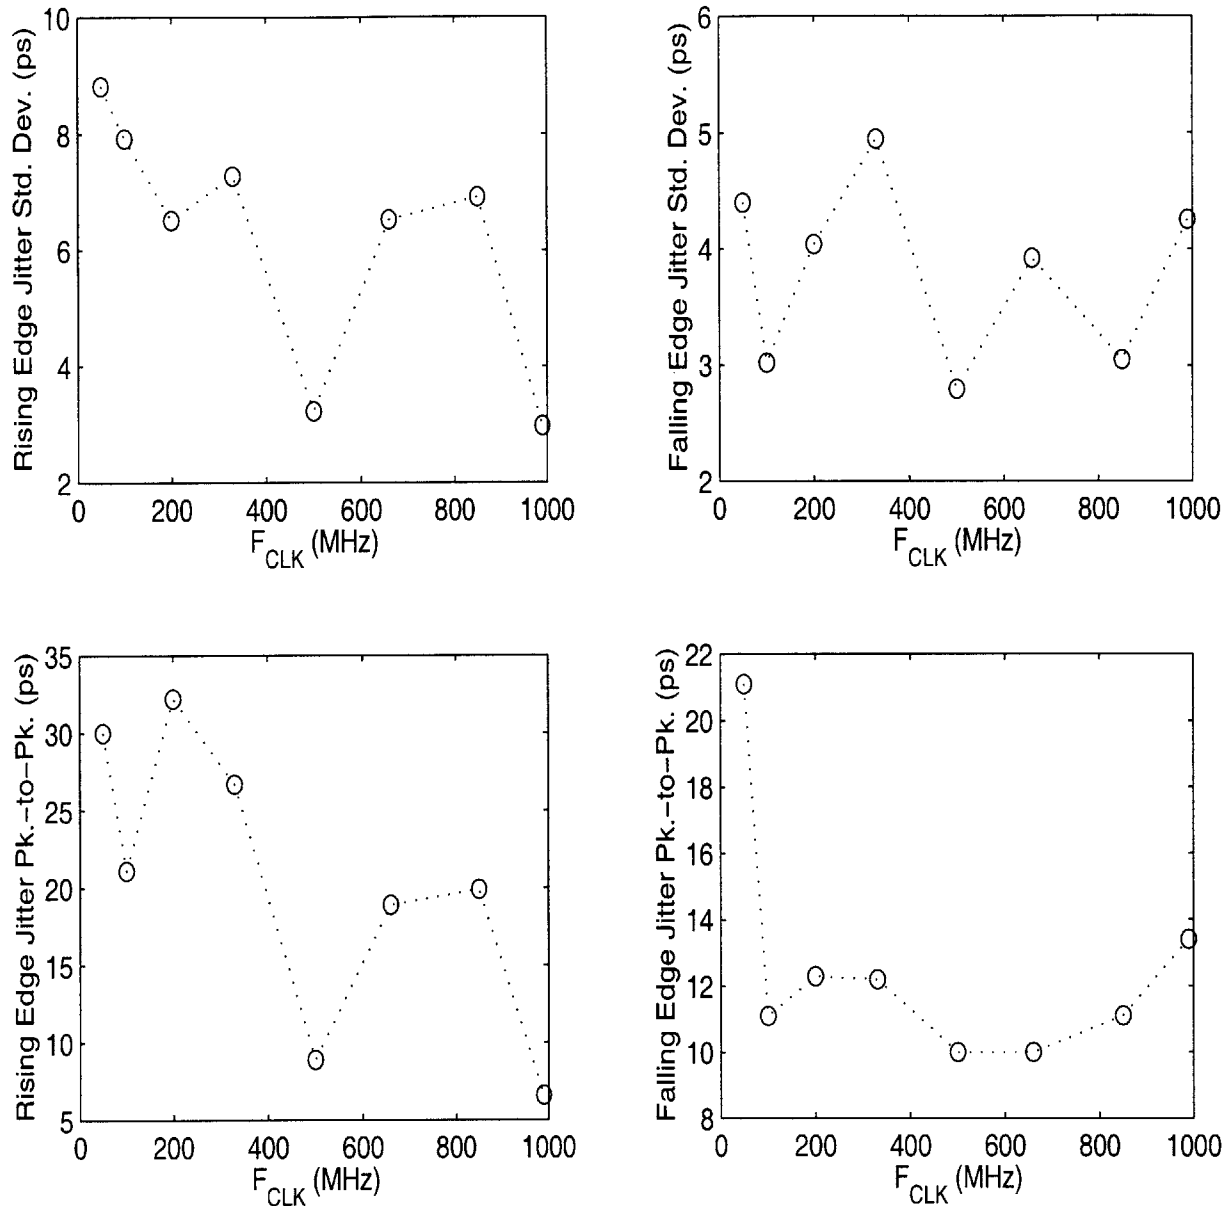


Figure 6.42: Jitter vs. bit-rate for $V_{swg}=1.5$ V

Based on Table 6.4 and Table 6.5, the standard deviation jitters are calculated as the root-difference-square of the input and output jitters. The peak-to-peak jitters are direct difference. The standard deviation data must be used with caution above 330 Mb/s since the input timing distributions become bimodal. The falling edge jitters do not have a strong frequency dependence. The falling edge peak-to-peak data point at 50 Mb/s is inconsistent with the standard deviation figure and appears to be a measurement error. The rising edge jitter becomes smaller as the bit rate is increased. These trends are consistent with Figure 6.38 and point to the charge storage on node 103 as the source of the jitter vs. bit rate variation.

6.7 Combined Laser Driver and VCSEL Characteristics

The previous sections examined the electrical characteristics of the laser driver. In particular, they established the laser drivers jitter characteristics, and its switching behavior in conjunction with a low impedance, resistive load. This section examines the performance of the laser driver when coupled with a VCSEL. Ultimately, the VCSEL is to be monolithically integrated with the laser driver using the EoE process. At this stage, however, the test are conducted by wire bonding the VCSEL n-type contact to the laser driver output as shown in Figure 6.12. Simulations have shown that the ~ 1 nH bond wire inductance does not limit the performance at the of the laser driver and VCSEL combination when operating at the 1 Gb/s bit rate.

Static characteristics

Static measurements were made on the laser driver and VCSEL combination in analogy with those of Section 6.4. The VCSEL in use during these measurements had similar static characteristics to those shown in Figure 6.13, but with a threshold current of 1.6 mA.

Figure 6.43 and Figure 6.44 show the dependence of the operation of the current mirrors on the supply voltage applied to the VCSEL. The results are similar to the corresponding measurements made by varying the voltage at the laser driver output (Figure 6.19). The applied voltage, however, must be raised above the 1.2 V diode drop before laser current begins to flow.

The mirroring functionality of the off-state and modulation current mirrors are shown in Figure 6.45 and Figure 6.46, respectively. The off-state circuit operates near its 10:1 mirroring ratio when the laser is supplied with 3.5 V. The modulation current mirror operates at its 4x gain with a 2 V laser supply. In the dynamic measurements shown below, the laser supply voltage was fixed at 2.5 V.

Figure 6.47 shows the combined operation of both the off-state and modulation current mirrors. As demonstrated in Figure 6.47(b), the laser driver may be configured to switch the laser between two above-threshold operating points by setting V_{off} above 1.5 V (for this particular laser) to determine the low state and then selecting V_{swg} to determine the high state.

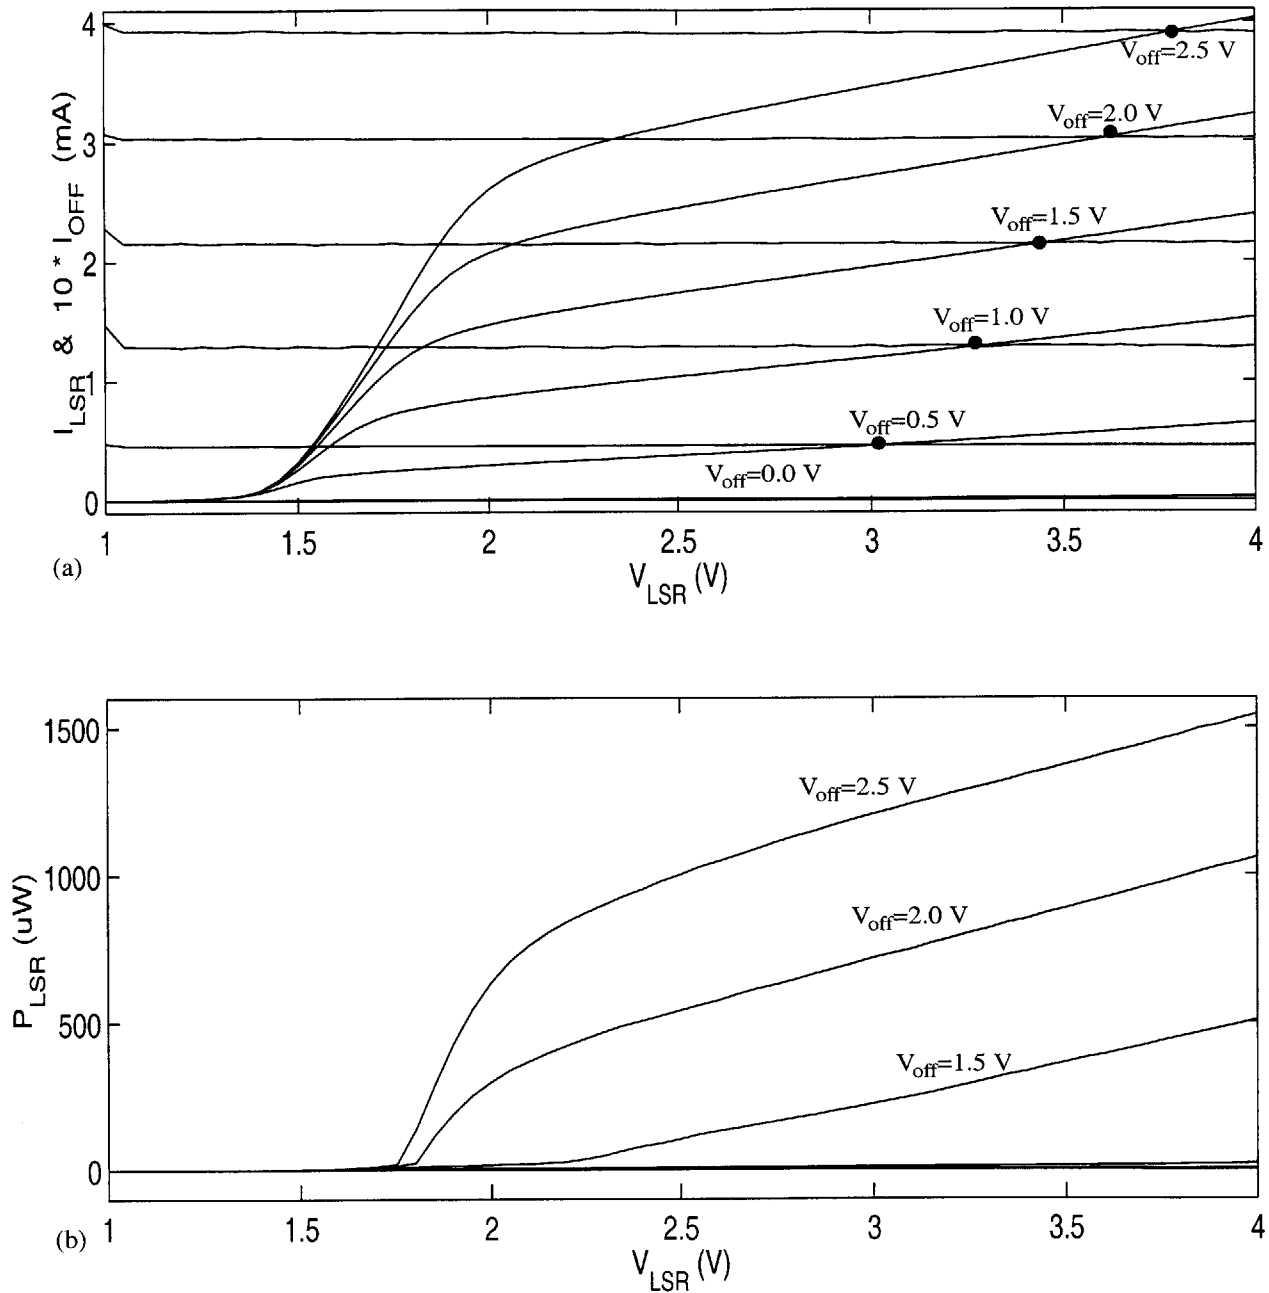


Figure 6.43: Dependence of off-state current mirror on VCSEL supply voltage

V_{LSR} , the VCSEL supply voltage, was swept while the off-state bias level was stepped. (a) Shows the measured current characteristics, which are similar to Figure 6.19(a), but are shifted to ramp up at $V_{LSR} = 1.5$ V because of the voltage drop across the VCSEL. A laser supply of around 3.5 V is needed to operate the off-state current mirror at its designed gain. (b) Shows the fiber-coupled light output. An off-state bias, V_{off} , of greater than 1.5 V is needed to operate the VCSEL above threshold. The modulation current mirror was off during this measurement.




Figure 6.44: Dependence of modulation current mirror on VCSEL supply voltage

V_{LSR} , the VCSEL supply voltage, was swept while the modulation bias level was stepped. (a) Shows the measured current characteristics, which are similar to Figure 6.19(b), but are shifted to ramp up at $V_{LSR} = 1.5$ V because of the voltage drop across the VCSEL. A laser supply of around 2 V is needed to operate the off-state current mirror at its designed gain. (b) Shows the fiber-coupled light output. The off-state current mirror off was during this measurement.

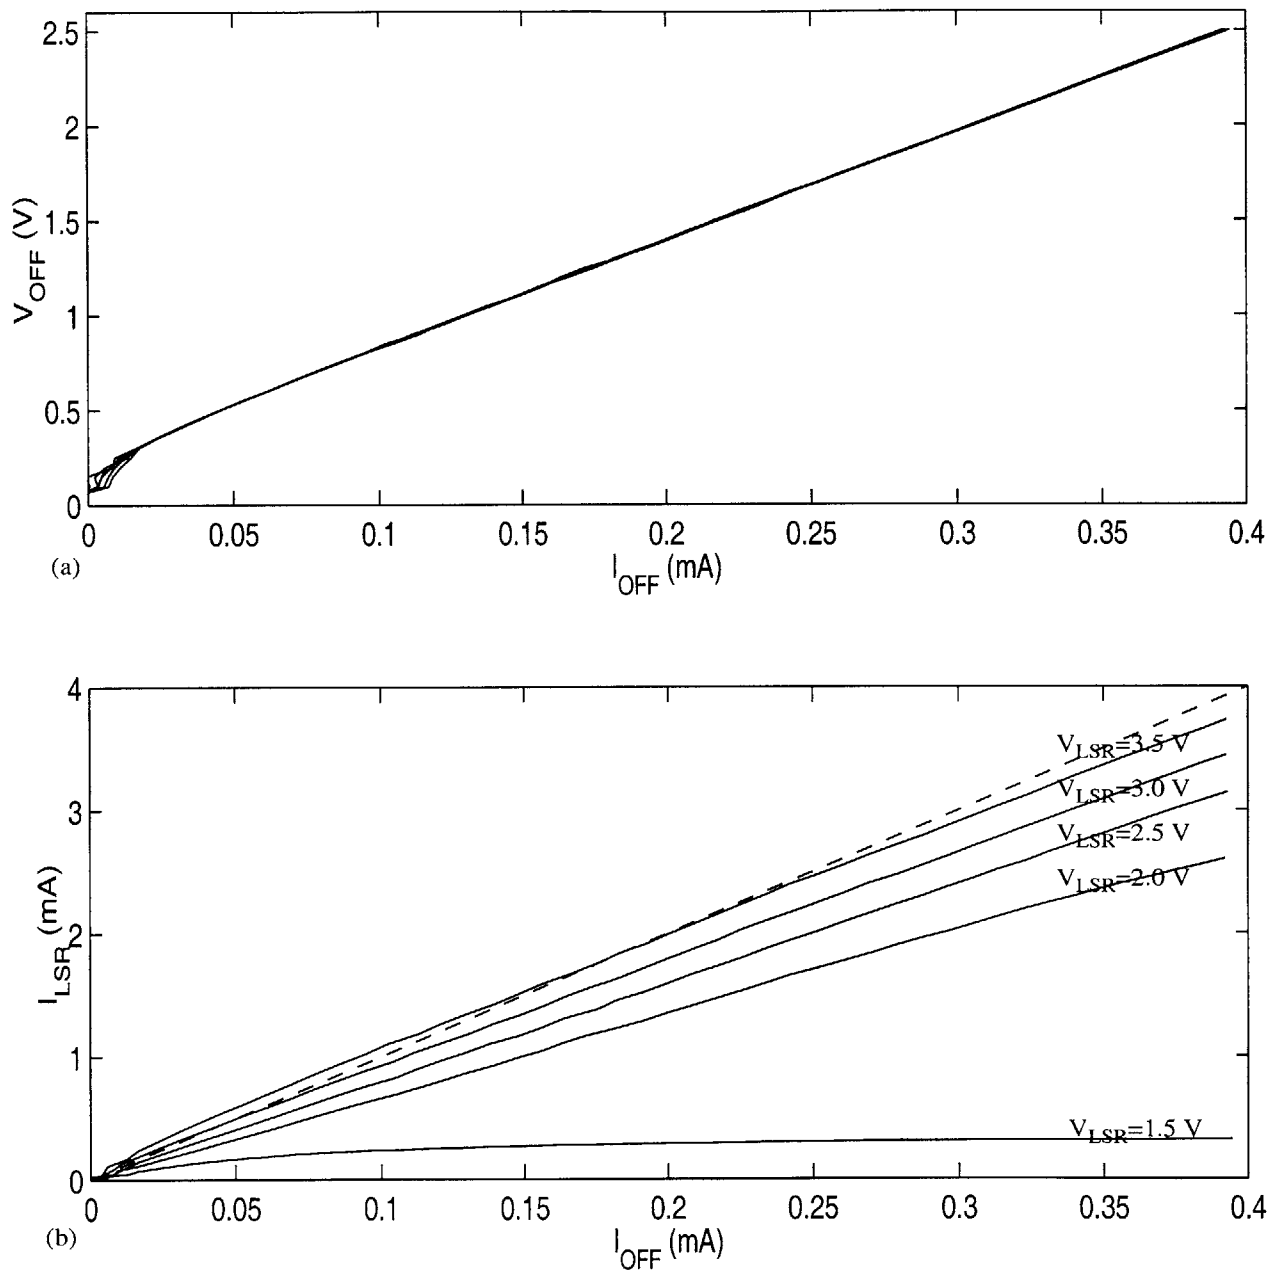


Figure 6.45: Static characteristics of off-state current mirror with VCSEL

The reference voltage, V_{off} , was swept while the reference current, I_{off} , and the laser current were measured. The laser supply voltage, V_{LSR} , was stepped from 1.5 V to 3.5 V in 0.5 V increments. This is analogous to Figure 6.20. (a) Shows the V_{off} - I_{off} relationship, which does not depend on V_{LSR} . (b) Shows the laser current as a function of the reference current. The dashed line has a slope of 10 representing the designed current mirror gain. The modulation current mirror is kept off during this measurement.




Figure 6.46: Static characteristics of modulation current mirror with VCSEL

As in Figure 6.21, the reference voltage, V_{SWG} , was swept while the reference current, I_{SWG} , and the laser current were measured. The laser supply voltage, V_{LSR} , was stepped from 1.5 V to 3.5 V in 0.5 V increments. (a) Shows the V_{SWG} - I_{SWG} relationship, which does not depend on V_{LSR} . (b) Shows the laser current as a function of the reference current. The dashed line has a slope of 4 representing the designed current mirror gain. The off-state current mirror is kept off during this measurement.




Figure 6.47: Combined output of off-state and modulation current mirrors with VCSEL

The laser current and fiber-coupled power were measured with both the off-state and modulation current mirrors active. V_{SWG} was swept and V_{off} was stepped while V_{LSR} was maintained at 2.5 V. Correct addition of the off-state and modulation current is seen.

Turn-on delay and relaxation oscillations

The low state operating point must be maintained above the lasing threshold in order to avoid any turn-on delay when modulating the laser. Turn-on delay results when the a laser is stepped from a current level below threshold to above threshold. The delay is the time required to build up the carrier concentration in the active region to the point where there is sufficient optical gain to support lasing.

In addition to turn-on delay, the laser step response can include relaxation oscillations. This is a result of the resonance formed by the coupling of the active carrier population and the photon population. The amount of ringing that results from this resonance depends on bias level.

The behavior of the laser driver + VCSEL combination, including turn-on delay and ringing, was observed by operating at various off-state levels while maintaining a constant modulation level. The VCSEL used in these measurements was similar to that shown in Figure 6.13 but with a threshold current of 0.8 mA. Table 6.6 summarizes the operating parameters under which the switching characteristics were measured. The DSC50 13 GHz photodetector was used to record the switching response of the laser; a 50 MHz clock was input to the laser driver. The response of the DCS50 was calibrated to match DC measurements of the fiber-coupled power made with a calibrated optical power meter. The turn-on and turn-off switching responses are shown in Figure 6.48 and Figure 6.49, respectively.

Both the turn-on delay and relaxation oscillations are visible in the turn-on waveforms of Figure 6.48. The relaxation oscillations are seen to be damped out at higher bias levels. This behavior underlines the need for work on designing VCSELs for specific interconnect applications in terms of their time domain response, and the need for accurate large signal models for use in the design of the laser driver. The turn-on delay drops from around 1 ns at zero off-state bias to zero delay above threshold. The delay values are summarized in Table 6.6 and plotted in Figure 6.50. The plot includes a theoretical curve of the turn off delay based on an established physical model [212].

V_{OFF} (V)	I_{OFF} (μ A)	$I_{LSR,initial}$ (mA)	$I_{LSR,final}$ (mA)	$\langle I_{LSR} \rangle$ (mA)	$P_{LSR,initial}$ (mW)	$P_{LSR,final}$ (mW)	$\langle P_{LSR} \rangle$ (mW)	Turn-On Delay (ps)
0.0	0	4.2×10^{-3}	1.83	0.871	0	0.60	0.25	961
0.2	3.48	5.3×10^{-2}	1.87	0.919	9.4×10^{-6}	0.62	0.27	848
0.4	24.8	0.287	2.08	1.51	4.8×10^{-4}	0.74	0.34	504
0.6	54.3	0.583	2.36	1.43	2.0×10^{-3}	0.91	0.43	246
0.8	86.3	0.894	2.65	1.74	0.055	1.07	0.56	20
1.0	119	1.21	2.95	2.05	0.23	1.25	0.73	8
1.2	153	1.53	3.25	2.36	0.41	1.44	0.90	0
1.4	187	1.85	3.55	2.66	0.60	1.61	1.08	4
1.6	222	2.16	3.84	2.97	0.78	1.76	1.26	10
1.8	256	2.46	4.13	3.27	0.97	1.90	1.43	8
2.0	291	2.76	4.41	3.55	1.14	2.03	1.59	0

Table 6.6: Operating parameters for optical switching measurements at various off-state levels

V_{off} was varied as shown while V_{swg} was maintained at 1.5 V ($I_{swg} = 377 \mu$ A). The laser supply voltage was 2.5 V. Fiber-coupled powers were measured with a calibrated power meter. ‘Initial’ and ‘Final’ refer to the low and high states, respectively. $\langle I_{LSR} \rangle$ and $\langle P_{LSR} \rangle$ are time-average values while switching at 50 MHz. The turn on delay of the laser switching waveforms shown in Figure 6.48 are listed. The delay time was measured from the mid point of each transition to that of the waveform corresponding to $V_{off}=2$ V. The delay times are plotted in Figure 6.50.




Figure 6.48: Laser rising-edge dependence on off-state level

The turn-on waveform of the laser was measured while varying the off-state level and keeping the modulation level fixed. The various parameters for these measurements are summarized in Table 6.6. (a) Shows the waveforms in a waterfall plot, and in (b), all of the waveforms are overlaid for easier comparison. A turn-on delay is observed in cases where the off-state level is below the laser's threshold. Ringing is also evident as a result of the relaxation oscillations.

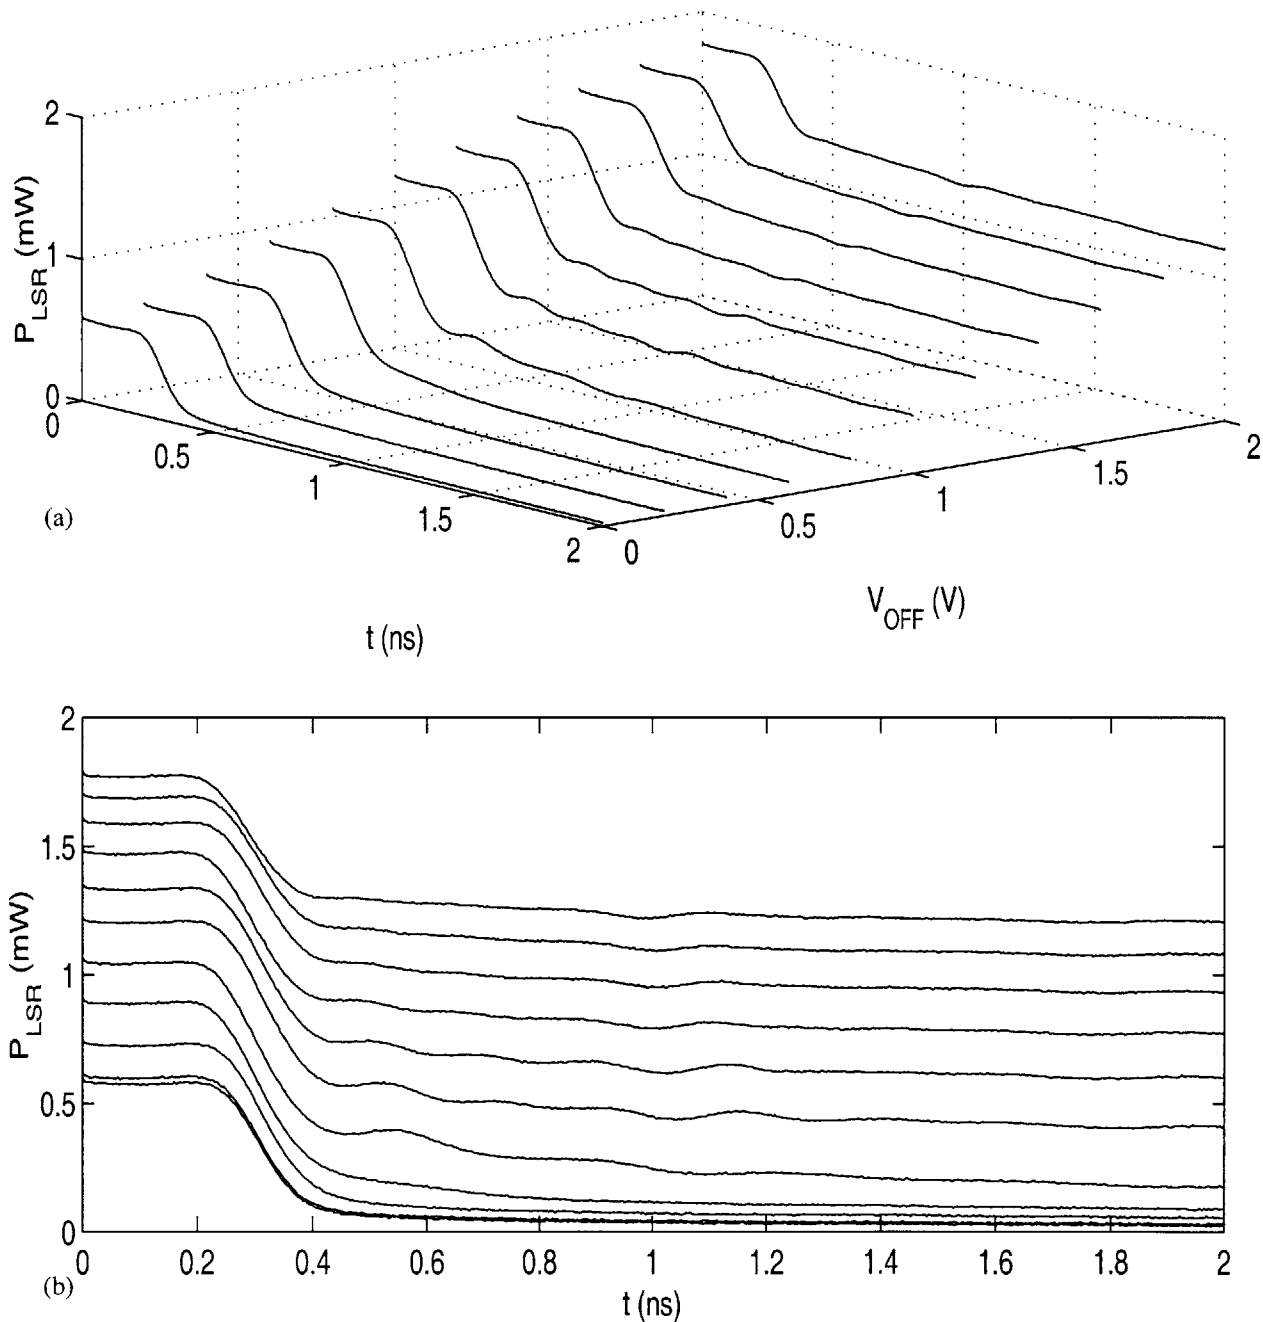


Figure 6.49: Laser falling-edge dependence on off-state level

The turn-off waveform of the laser was measured while varying the off-state level and keeping the modulation level fixed. The various parameters for these measurements are summarized in Table 6.6. (a) Shows the waveforms in a waterfall plot, and in (b), all of the waveforms are overlaid for easier comparison. A small amount of ringing due to relaxation oscillations is also present in the above-threshold results.




Figure 6.50: Turn-on delay dependence on off-state level

Measured time delays were extracted from Figure 6.48 based on the midpoint of the rising edge and are summarized in Table 6.6. The solid line is based on a theoretical formula given in [212]. $I_{initial}$ and I_{final} are the low- and high-state laser currents. In the theoretical curve, $I_{final} - I_{initial}$ was fixed at 1.7 mA, the average value determined from the measurements, and the numerical value of 0.8 mA was extracted for I_{th} from the L-I curve of the laser. τ_f is the differential spontaneous lifetime at I_{final} in the absence of lasing and is determined in terms of the active volume of the laser, V_{act} , the internal quantum efficiency, η_i , and the bimolecular (radiative) recombination coefficient, B . Since numerical values for these parameters were not available, $\sqrt{(qV_{act})/(\eta_i B)}$ was set to $6.5 \times 10^{-11} \text{ ns/A}^{-1/2}$ to fit the measured data. The resulting values of 1.3-1.6 ns are in line with expected values for τ_f .

Above-threshold switching characteristics

The measurements made above at various off-state levels can be used to select an off-state level which allows the laser to be modulated without incurring the effects of turn-on delay or relaxation oscillations. The off-state level corresponding to $V_{\text{off}}=1.4$ V was selected for the following measurements made at various modulation levels. Table 6.7 summarizes the parameters used in the measurements. The turn-on waveforms are plotted in Figure 6.51 and the turn-off waveforms are plotted in Figure 6.52. To allow easier

V_{SWG} (V)	I_{SWG} (μA)	$I_{\text{LSR},\text{initial}}$ (mA)	$I_{\text{LSR},\text{final}}$ (mA)	$\langle I_{\text{LSR}} \rangle$ (mA)	$P_{\text{LSR},\text{initial}}$ (mW)	$P_{\text{LSR},\text{final}}$ (mW)	$\langle P_{\text{LSR}} \rangle$ (mW)
0.00	-36	1.85	1.98	1.92	0.61	0.69	0.65
0.25	9.8	1.86	2.11	1.99	0.61	0.77	0.69
0.50	76.9	1.87	2.33	2.10	0.62	0.89	0.76
0.75	150	1.89	2.59	2.25	0.63	1.05	0.84
1.00	224	1.90	2.87	2.39	0.64	1.21	0.93
1.25	300	1.92	3.15	2.54	0.65	1.37	1.02
1.50	376	1.94	3.43	2.70	0.66	1.52	1.10
1.75	452	1.95	3.71	2.84	0.67	1.68	1.19
2.00	529	1.97	3.98	2.98	0.68	1.83	1.26
2.25	606	1.98	4.25	3.13	0.69	1.96	1.34
2.50	682	2.00	4.51	3.27	0.70	2.08	1.41

Table 6.7: Operating parameters for optical switching measurements at various modulation levels

V_{SWG} was varied as shown while V_{off} was maintained at 1.4 V ($I_{\text{off}} = 187 \mu\text{A}$). The laser supply voltage was 2.5 V. Fiber-coupled powers were measured with a calibrated power meter. ‘Initial’ and ‘Final’ refer to the low and high states, respectively. $\langle I_{\text{LSR}} \rangle$ and $\langle P_{\text{LSR}} \rangle$ are time-average values while switching at 10 MHz. A small increase in the off-state laser current and power are seen as the modulation level in increased. This trend was not observed in the static measurements since the off-state current mirror was exercised with the other modulation mirror fully disabled--i.e. laser driver input set to high and $V_{\text{SWG}}=0$. The effect is not detrimental to the operation of the laser driver.





Figure 6.51: Laser rising-edge dependence on modulation level

The turn-on waveform on the laser was measured while varying the modulation level and keeping the off-state level fixed. The various parameters for these measurements are summarized in Table 6.7. (a) Shows the waveforms in a waterfall plot, and in (b), all of the waveforms are overlaid for easier comparison.

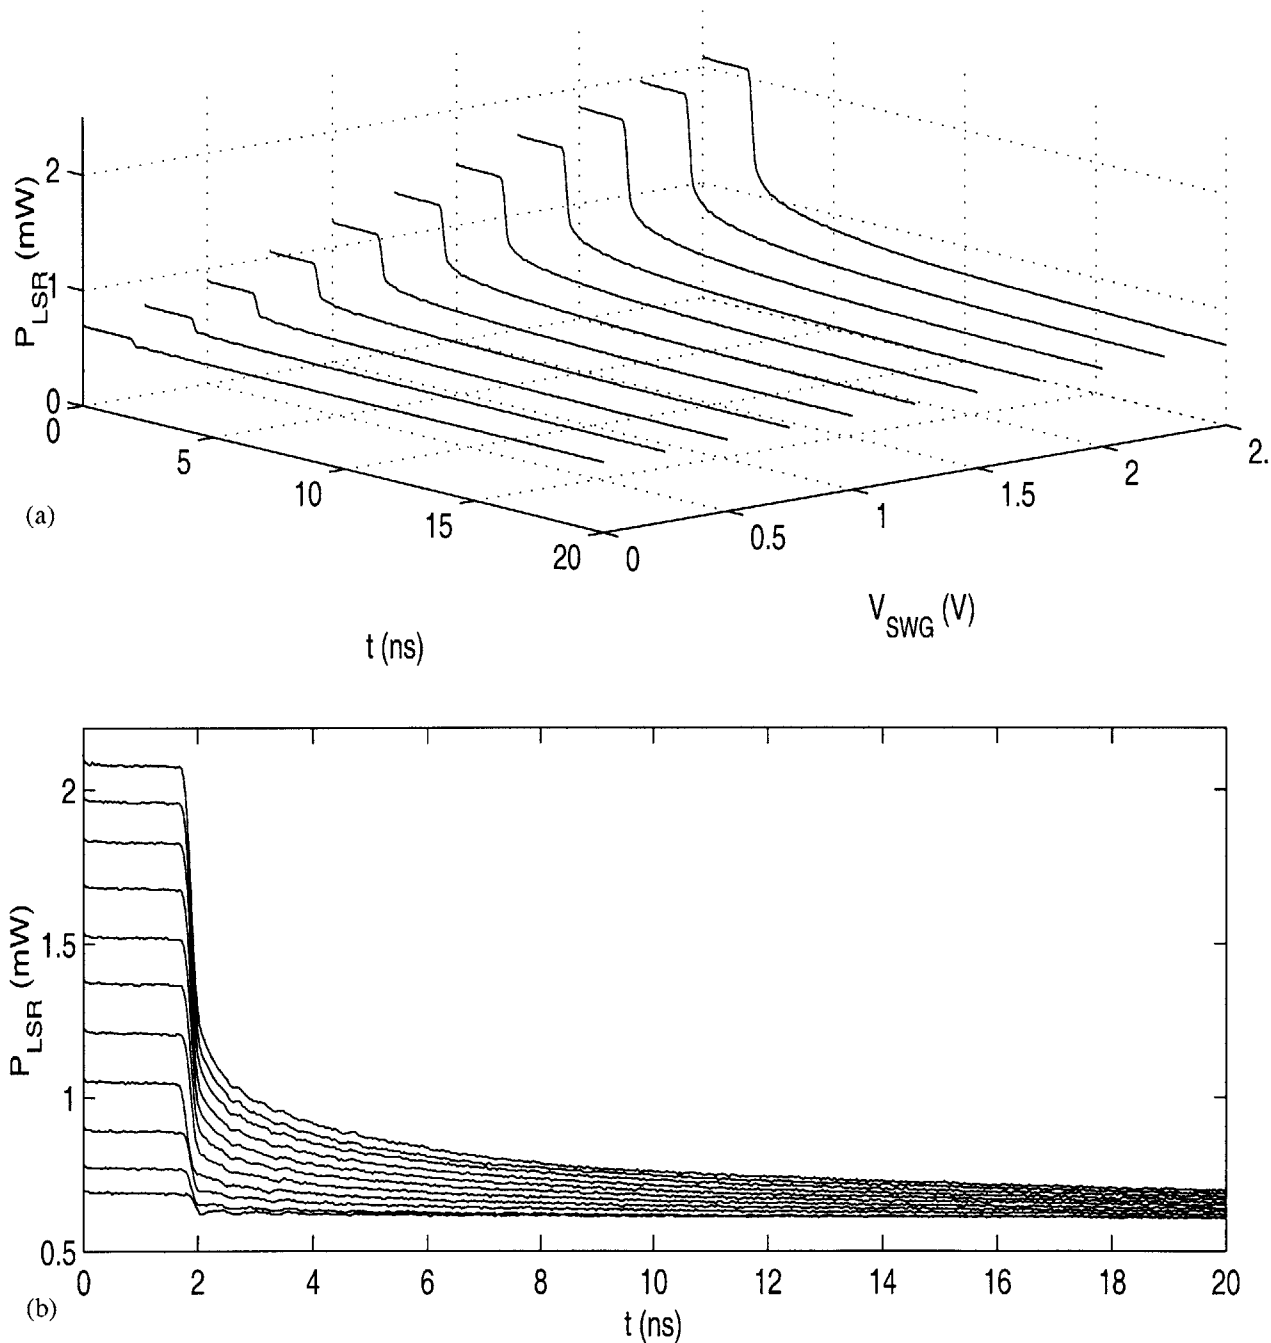


Figure 6.52: Laser rising-edge dependence on modulation level

The turn-off waveform on the laser was measured while varying the modulation level and keeping the off-state level fixed. The various parameters for these measurements are summarized in Table 6.7. (a) Shows the waveforms in a waterfall plot, and in (b), all of the waveforms are overlaid for easier comparison.

V_{SWG} (V)	$P_{LSR,initial}$ (mW)	$P_{LSR,final}$ (mW)	Amplitude (mW)	Rise Time 10%-90% (ns)	Rise Time 20%-80% (ns)	Fall Time 90%-10% (ns)	Fall Time 80%-20% (ns)
0.00	0.61	0.69	0.08	10.51	8.09	5.90	1.76
0.25	0.61	0.77	0.16	8.67	6.52	4.10	1.05
0.50	0.62	0.89	0.27	6.99	3.79	4.10	1.21
0.75	0.63	1.05	0.42	5.12	1.76	3.95	0.82
1.00	0.64	1.21	0.57	4.49	1.13	3.79	0.86
1.25	0.65	1.37	0.72	4.41	1.13	3.71	0.98
1.50	0.66	1.52	0.86	5.51	1.13	3.91	0.98
1.75	0.67	1.68	1.01	4.10	1.13	3.75	1.02
2.00	0.68	1.83	1.15	4.02	1.09	4.02	1.13
2.25	0.69	1.96	1.27	4.38	1.09	3.75	1.21
2.50	0.70	2.08	1.38	4.49	1.13	4.10	1.29

Table 6.8: Transition times of optical switching waveforms at various modulation levels
Rise and fall times were extracted from Figure 6.51 and Figure 6.52.

comparison, the waveforms have been normalized and overlaid in Figure 6.53. The rise and fall times have also been extracted and are summarized in Table 6.8 and plotted in Figure 6.54.

The turn-on waveforms corresponding to $V_{swg} = 0, 0.25,$ and 0.5 V exhibit the slow turn-on behavior observed in the electrical switching waveforms of Figure 6.23. At the higher biases, however, the normalized rise and fall waveforms are essentially identical and have a 10%-to-90% transition time of 4 ns and a 20%-to-80% transition time of 1 ns. By contrast, the electrical 10%-to-90% transition times were below 540 ps, the 20%-to-80% rise time was below 200 ps, and the 20%-to-80% fall time was around 50 ps.




Figure 6.53: Normalized optical transition waveforms at various modulation levels

Waveforms measured at 11 modulation levels summarized in Table 6.7 have been normalized and overlaid for comparison. (a) and (c) show the turn-on waveforms on 20 ns and 2 ns time scales, respectively. Similarly, (b) and (d) show the turn-off transitions on 20 ns and 2 ns time scales. In the turn-on data, the waveforms corresponding to $V_{\text{swg}}=0, 0.25,$ and 0.5 V exhibit the slow turn off effect observed in Figure 6.23.

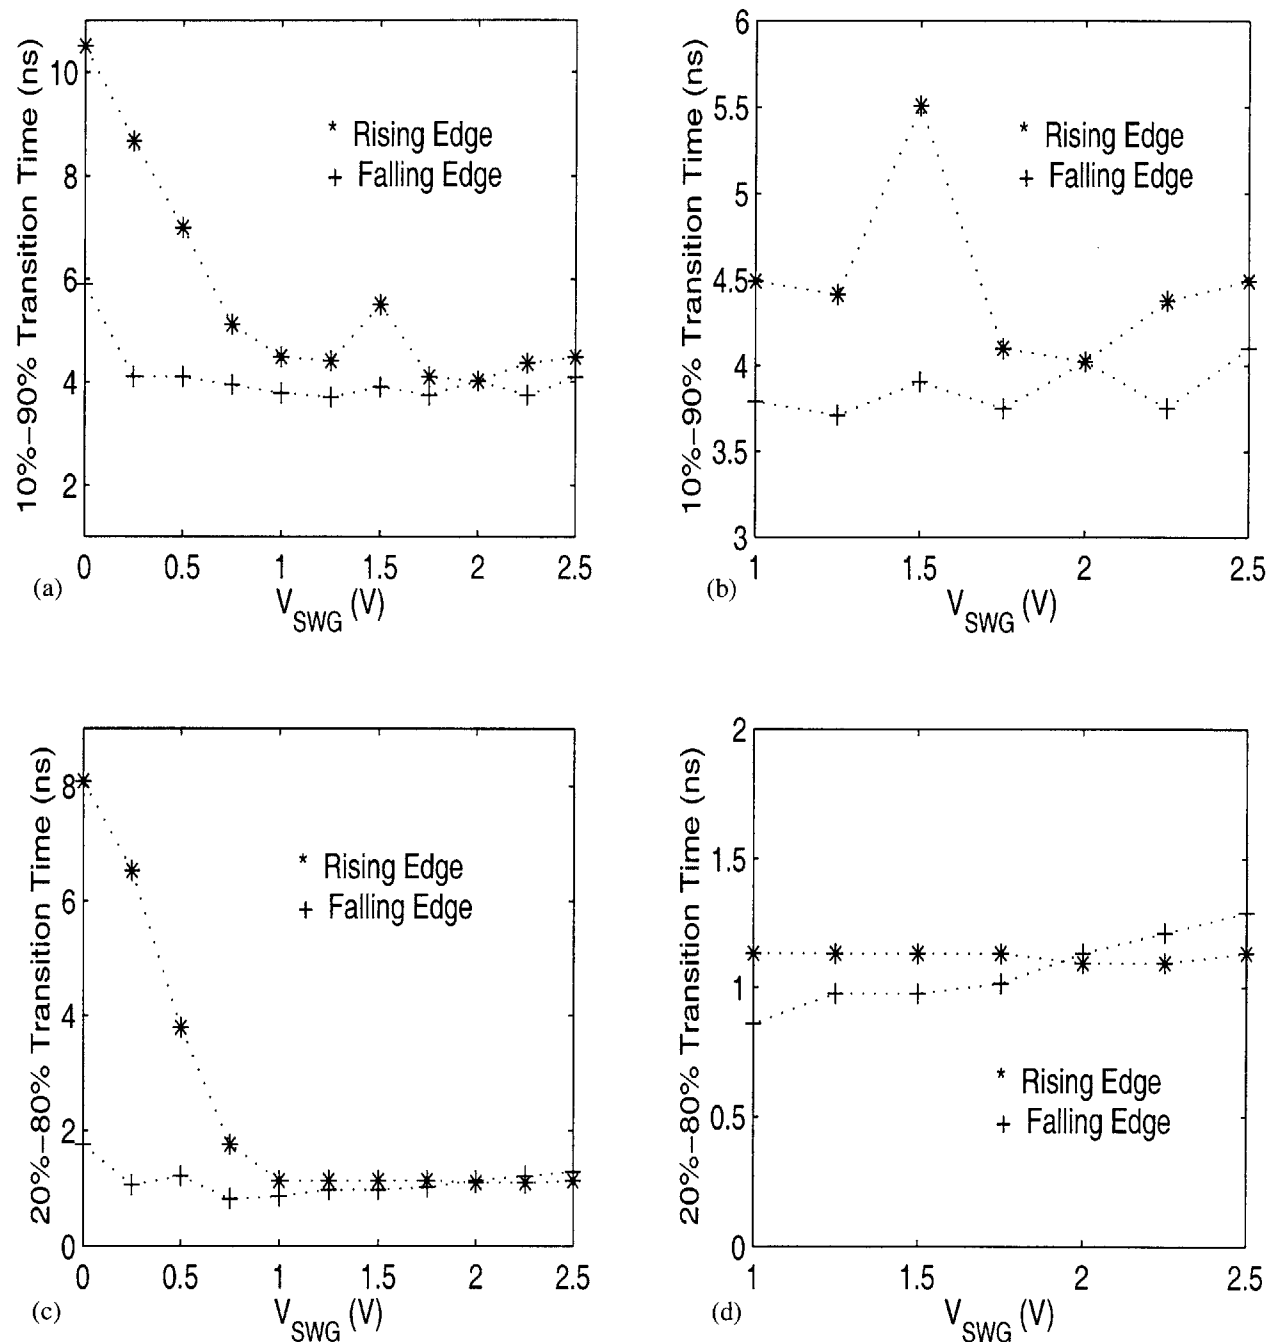


Figure 6.54: Transition times of optical switching waveforms measure at various modulation levels

The rise and fall times extracted from Figure 6.51 and Figure 6.52 and summarized in Table 6.8. (a) and (b) are the 10%-to-90% transition times. (b) excludes the data points at low modulation levels which exhibit the slow turn on effect. Similarly, (c) and (d) show the 20%-to-80% transition times.

Drain lag on the output node

The culprit, once again, is drain lag--this time on the output FETs, E111 and E211. Figure 6.55 is the result of simulating the laser driver with the equivalent input circuit of Figure 6.16 used as the load. The sharp transition of the V_{111} waveform indicates that the internal switching of the laser driver is not the cause of the poor transition characteristics at the output. The long response tails seen in Figure 6.51 and Figure 6.52 are reproduced, though the effect is worse in the simulation. This is because drain lag is actually overestimated in the Vitesse MESFET models. Both the modulation current and the off-state current are seen to contribute to the long response tail in the output current.

To further isolate the problem as drain lag, and not the consequence of the parasitic capacitance of the VCSEL, Figure 6.56 repeats the above simulation using only a 500 Ω resistance as the load. (The equivalent input circuit of Figure 6.16 reduces to a 498 Ω resistance at low frequencies.) The simulation results are found to be quite similar to those above.

If the input impedance of the VCSEL was due primarily to its series resistance, 85 Ω for the VCSEL shown in Figure 6.13, drain lag would be greatly reduced. Figure 6.57 shows the result of simulating the laser driver with an 85 Ω load. The magnitude of the extended response tail is at a tolerable level.

Finally, for comparison with the experimental results of the previous two sections, Figure 6.58 shows simulated waveforms when using a 25 Ω load. The magnitude of the slow response tail is now very small, and is similar to the measured electrical output waveforms.




Figure 6.55: Simulation of laser driver loaded by VCSEL equivalent input circuit shown in Figure 6.16

The results of Figure 6.51 and Figure 6.52 are reproduced in this simulation by loading the laser driver with the equivalent input circuit determined from S_{11} measurements on a VCSEL. The sharp transition of the V_{111} waveform indicate that the internal switching of the laser driver is not the cause of the poor transition characteristics at the output. Both the modulation current and the off-state current are seen to contribute to the long response tail in the output current.




Figure 6.56: Simulation of laser driver loaded by a 500 Ω resistor

The equivalent input circuit of Figure 6.16 reduces to a 498 Ω resistance at low frequencies. In this simulation, the laser drivers is loaded with a 500 Ω resistance and is found to have very similar transition characteristics to the simulation results using the full equivalent circuit. Evidently, the capacitive loading of the VCSEL does not play a significant role in creating the long response tails.




Figure 6.57: Simulation of laser driver loaded by an 85 Ω resistor

The 500 Ω resistor of the simulation in Figure 6.56 is replaced with an 85 Ω resistor. The transition is seen to improve greatly, though the long response tail is still present.




Figure 6.58: Simulation of laser driver loaded by an 25 Ω resistor

The previous simulation is repeated with a 25 Ω load resistance. This is the output loading present the in laser driver characterization in Section 6.5 and Section 6.6. The magnitude of the long response tail is reduced to the point where were it does not have a significant effect on the performance of the laser driver.

Redesign

Since drain lag is a consequence of signal coupling between the drain of a MESFET at its backgate, the effect may be reduced by using a cascode technique to reduce the amplitude of the drain signal. This approach is reviewed in Appendix E where drain lag is examined more carefully. A quick redesign of the laser driver to include cascoded output FETs is shown in Figure 6.59. Careful optimization of this circuit will be necessary to regain the level of performance shown in the original design, but as seen in Figure 6.60, the modified circuit is able to drive the $500\ \Omega$ load with output waveforms similar to that of the original circuit driving a load between 25 and $85\ \Omega$.

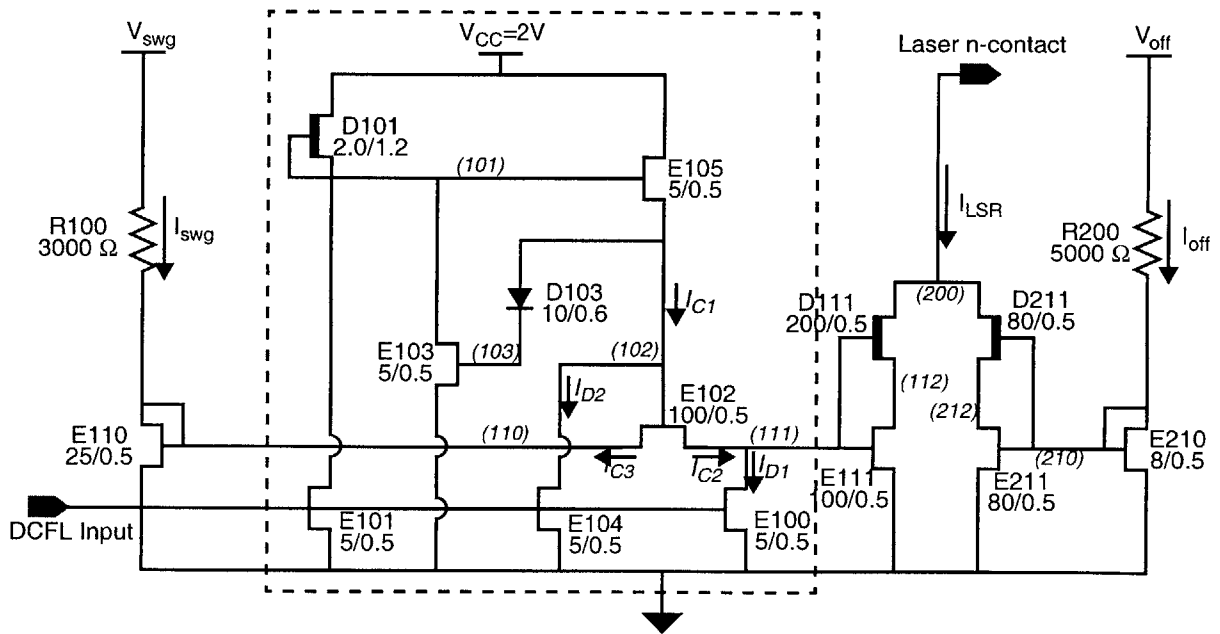


Figure 6.59: Modification of the laser driver to allow large resistive loads

The laser driver shown in Figure 6.3 and studied throughout this chapter can be modified to support large resistive loads by adding the cascode transistors D111 and D211. The circuit has not been re-optimized following this change, however, since D111 added a significant additional capacitance to node 111, E110 and E111 were scaled by $1/2$ so as not to overburden the switching circuit.




Figure 6.60: Modified laser driver simulated with 500 Ω load

The modified laser driver of Figure 6.59 was simulated with a 500 Ω load. The magnitude of the long response tail is seen to be greatly reduced, and is comparable to the case of a 25 Ω load on a original laser driver. Note that the initial edge rate has been reduced compared to Figure 6.58 because of the greater capacitive loading on node III. Re-optimization of the modified circuit should improve these transition characteristics.

Summary

The importance of having an accurate VCSEL model available in the design of the laser driver was painfully demonstrated in this section. The VCSEL's turn-on delay and relaxation oscillations must be considered in designing the VCSEL, the laser driver, and the entire optical link. This is particularly true because they help determine the minimum bias level which can be used in the off-state. An even more immediate issue was seen to be the effect of the electrical characteristics of the VCSEL on the behavior of the laser driver. A relatively simple modification to the existing laser driver can cope with the unexpectedly large input impedance of the VCSEL.

6.8 Third-generation Laser Driver Designs

The laser driver analyzed in the previous sections goes a long way towards meeting the requirements for a compact, low power transmitter for use with an integrated VCSEL in a 1 Gb/s digital optical interconnect. However, a number of shortcomings have been identified in its design:

- Drain-lag on the output node results in a long response tail when driving the large resistive load of a VCSEL.
- The internal switching dynamics do not track the modulation bias level. At high modulation levels, the primary switching mechanism only partially charges the gate of the drive FET resulting in long response tail as the charging is completed by the reference resistor.
- A floating node at a critical point in the switching circuit results in a bit-rate dependence of the rising waveform, and accompanying pattern dependent jitter.
- Drain-lag on internal node 101 results in long response tail under low modulation level
- There is a delay asymmetry between the rising and fall output transitions due to the extra delay through the turn-on switching circuit.

This section will put forth two new laser driver design which overcome these limitations. The first is an improved switched current mirror design while the second returns to the idea of a differential current switch. The latter incorporates key observations from the current mirror laser drivers and from the fully-differential receiver of Appendix H to arrive at a compact and low power design. These designs are of a preliminary nature and are intended to demonstrate their underlying concepts. They have not been optimized, nor have they been simulated over process and temperature spreads.

An improved active switched current mirror laser driver

A schematic of the new design is shown in Figure 6.61. This design is conceptually similar to the original circuit shown in Figure 6.3, but in addition to cascoding the outputs as in Figure 6.59, it includes a turn-on switching mechanism which tracks the modulation bias level and eliminates critical floating nodes. A delay is also included in the turn-off signal path in order to reduce the delay asymmetry. Note that the cascode on the off-state current mirror has its gate grounded instead of connected to node 210 as in the modulation current mirror (as shown in Appendix E). Since V_{210} is not being modulated, not making this connection is not detrimental. Furthermore, grounding the gate of D211 reduces the capacitive coupling of output voltage variation to node 210. Simulation results from the new laser driver are shown in Figure 6.62.

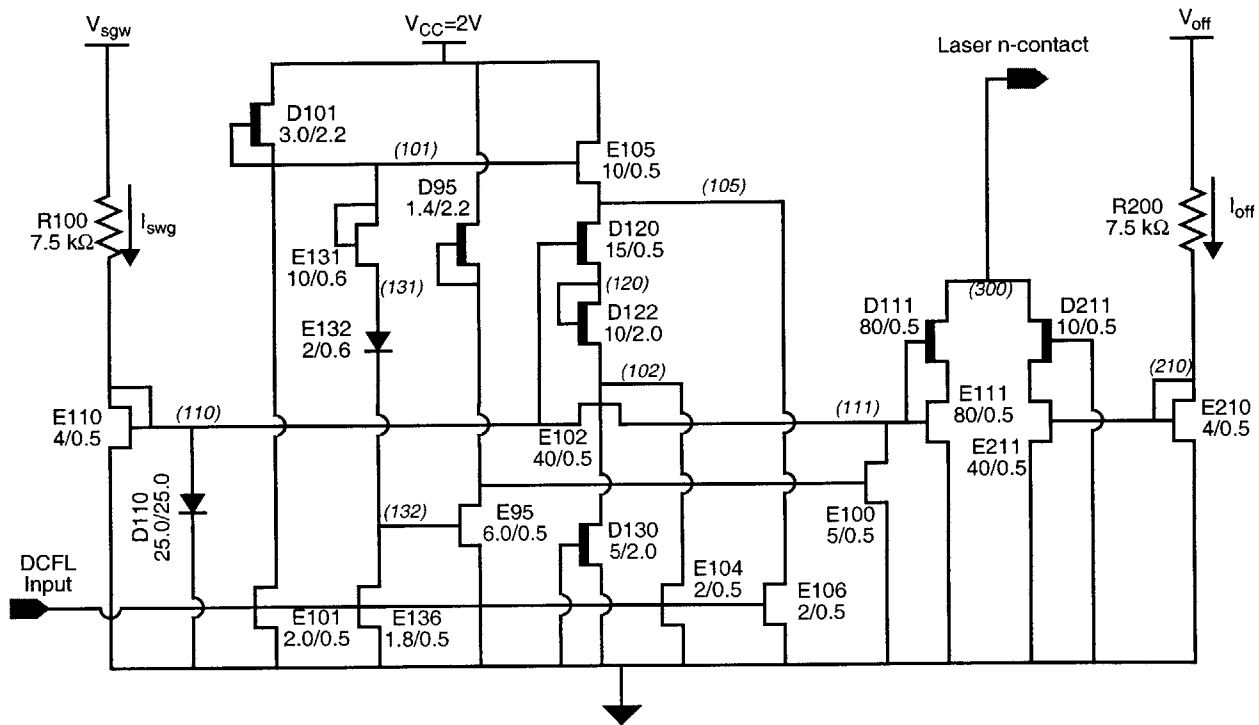


Figure 6.61: An improved active switched current mirror laser driver

This design is conceptually similar to the original circuit shown in Figure 6.3, but in addition to cascoding the outputs as in Figure 6.59, it includes a turn-on switching mechanism which tracks the modulation bias level and eliminates critical floating nodes. A delay is also included in the turn-off signal path in order to reduce the delay asymmetry.




Figure 6.62: Simulation of the improved active switched current mirror laser driver shown in Figure 6.61. The upper curve shows the static characteristics of the modulation current mirror verifying that the new switching mechanism allows correct mirroring. Note that, unlike the previous design, the output now tracks the reference down to zero modulation current. The middle curves show a transient simulation of the internal nodes involved in the switching operation, and the lower curve shows output current. The simulation was carried out with $V_{swg}=V_{off}=1$ V, and the VCSEL equivalent input circuit of Figure 6.16 was used to load the output of the laser driver.

An important addition to the circuit is the diode-connected DFET, D110. This $L=W=25.0\ \mu\text{m}$ DFET is used to provide bypass capacitance on node 110. This prevents node 110 from being perturbed by the large, rapid signal swings on the nodes to which it is capacitively coupled. Without the low AC impedance provided by the bypass capacitance, node 110, and in turn node 111 are left to equilibrate through R100, which is a slow process. By stabilizing node 110 in this way, many of the constraints on the rest of the switching circuit may be lifted. Furthermore, the capacitance on node 110 allows node 111 to actually be charged by node 110 in addition to charge delivered through the gate of the pass transistor.

A further benefit of the bypass capacitance may be found when using a large number of these laser drivers in an OEIC. In this case, the reference level would, presumably, be set globally by a reference circuit which compensates for temperature and process variations. Local bypass capacitance on node 110 then allows multiple laser drivers to be connected to a common reference voltage without interacting.

The primary difference in the new design is the switching circuit. As before, E101 and D101 form an inverter which drives E105. Turning on E105, following the falling edge of the input, begins the turn-on transient. With node 120 and 102 initially at 0 V, significant current flows through D120 and D122, pulling up the gate of the pass transistor. Since node 111 is also initially at 0 V, the gate of the pass transistor momentarily becomes forward biased and contributes to the charging of node 111. Node 111 is also charged through the pass transistor by the D110 capacitance. A bias current is established through D130 which sets the voltage drop across D122 and the gate-to-source voltage of D120. The tracking of the reference voltage, V_{110} , occurs as a result of the action of D120 as a positive level shifting source follower. D120, D122, and D130 are sized to hold V_{102} around 300 mV above V_{110} . In this way, the pass transistor is on but there is not significant gate voltage to distort the operation of the current mirror. Tracking the V_{110} reference in this way should allow operation over a wider range of modulation currents, temperatures, and process variations.

The static curves in Figure 6.62 confirm that the new switching circuit does not interfere with the operation of the current mirror. It is seen that, unlike the previous design, the output now tracks the reference down to zero modulation current. The simulated node voltages show the rapid and accurate switching of V_{102} to a point $\sim 300\ \text{mV}$ above V_{110} . The output current is seen to have rise and fall times of below 200 ps. The simulation was carried out with the laser driver loaded by the equivalent input circuit shown in Fig-

ure 6.16. The cascodes are seen to be effective in holding the magnitude of the drain-lag induced extended response tail to a acceptable value.

To eliminate drain lag on node 101, E131 and E132 form a diode chain along with the gate of E95 in order to clamp V_{101} at 1.5 V. This is the same way in which DCFL logic gates function without exhibiting an extended response tail.

Also, to reduce the difference in the input-to-output delay between the turn-on and turn-off cycles, E100, the EFET which is used to discharge node 111, is driven by inverter E95/D95. In turn, E95 is driven by the E101/D101 inverter through the clamping diode chain. E136 is provided to discharge node 132. the input to E95, which would otherwise remain charged and floating when V_{101} is driven low.

Note that no discharge path has been provided for node 131 in diode the chain used to clamp node 101. Node 131 is allowed to remain charged during the turn-off cycle. Careful simulation is required to determine if this has any negative consequence. If so, an additional EFET may be used to discharge it.

The average power dissipation of the laser driver, not including the reference currents or the output current, is 0.5 mW. This is somewhat lower than the original design which dissipated 0.6 mW. The modulation current mirror gain is also larger in this circuit--10:1 at output currents up to around 1.5 mA then dropping somewhat as D130 (the "capacitor") begins to leak¹. The larger gain means lower power overhead in setting the modulation level. The current gain on the off-state mirror has also been increased to 10.

A low power differential laser driver

Because of their very low power consumption, the active switched current mirror laser drivers show promise for use in optical interconnect applications requiring thousands of optical outputs. One of their drawbacks, however, is the subtlety of their design. Careful device sizing is required to control the turn-on transient over the required range of modulation levels, though this task is simplified if the design is intended for use at one specific modulation level. The sensitivity to device sizing may lead to potentially large performance variation over process and temperature spreads.

Traditional current-steering laser driver designs, which use a differential pair as a current switch, avoid many of these design subtleties since the modulation current may be entirely determined by a simple cur-

1. The problem of gate leakage in D₁₃₀ may be eliminated by reversing its polarity. The capacitance will drop as V_{110} is increased, but this is probably acceptable.

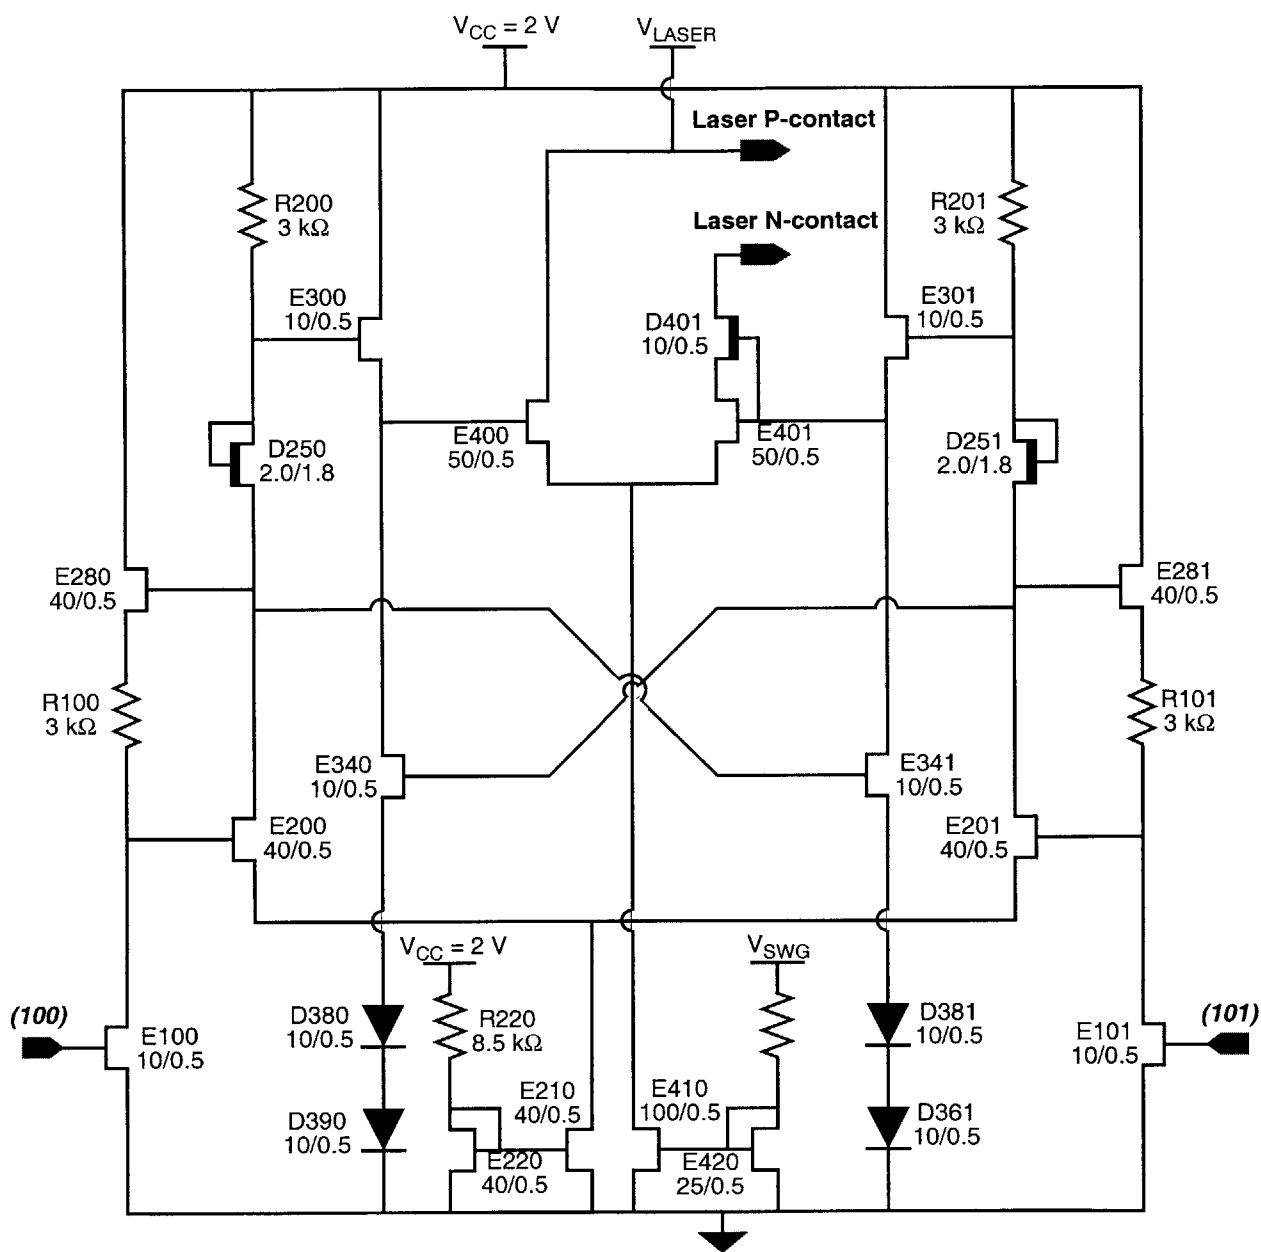


Figure 6.63: Schematic of a low-power differential laser driver

This design consists of three states: conversion from imperfectly timed complementary logic inputs to synchronized differential signal, an efficient driver, and the differential pair current switch. These blocks are detailed in Figure 6.67, Figure 6.66, and Figure 6.65. A circuit for generating the complementary inputs from a single DCFL signal is shown in Figure 6.64. Simulation results of this circuit are shown in Figure 6.68. The off-state current mirror has not been included in this schematic but may be implemented as in the switched current mirror laser driver designs.




Figure 6.64: Circuit for generation of complementary logic signals

The DCFL input is inverted by E50/D50. The input and its inverse are used to drive the other two stages so that the signal path to each of the outputs is roughly equal.

rent mirror. Such a design is also advantageous in that it maintains a constant current level through the supply and ground lines, thereby reducing noise generation. However, initial attempts at a differential-pair based laser driver were thwarted by the large power overhead associated with implementing the conventional designs. The overhead included the conversion of a single-ended DCFL logic signal into a suitable pair of differential signals, and the conventional use of power-hungry source-followers to drive the final differential-pair.

Observations made in the design of the switched current mirror laser drivers, and in the design of the receiver discussed in Appendix H, can be used to overcome the overhead associated with the conventional approach. Figure 6.63 is the preliminary design of a differential-pair based laser driver with 1.8 mW of power consumption. The inputs to this circuit, “(100)” and “(101)”, are a pair of complementary DCFL logic signals. It is desirable to have the input closely timed (i.e. small phase error between them), but perfect timing is not necessary. The complementary signals may be generated from a single input using the circuit of Figure 6.64. (This circuit is included in the 1.8 mW power figure, and contributes significantly to it.)

The circuit in Figure 6.63 may be broken into three parts: conversion from imperfectly timed complementary logic inputs to a pair of well-synchronized differential signal, an efficient driver, and the differential-pair current switch. They are most easily understood in reverse order. The final stage is shown in Figure 6.65 with the remainder of the circuit deleted. The modulation current is set by the current mirror,

E410 and E420, and is switched between the two sides of the differential-pair by the outputs of the buffer stage, V_{DRIVE} and its complement. The output is cascoded by D401 as shown previously. Even without this cascode, the differential-pair design actually performs comparably to the cascoded switched current mirror designs. This is because E401 of the differential-pair can actually be thought of as cascoding the current mirror's E410 when all of the modulation current is flowing through the output. The addition of D401 further reduces the sensitivity to drain lag¹.

The key to reducing the power consumption of the laser driver is in the buffer stage. It is shown isolated in Figure 6.66². To lower the overall power consumption of the laser driver, both the power consumption of the buffer and the size of the load it presents to the previous stage must be minimized. The buffer shown in Figure 6.66 does this by using EFETs in a digital switching mode and operating them at their maximum drain-to-source current level³. E300 and E301 are used to charge the outputs and are controlled by the differential signals, $V_{\text{DIFF,A}}$ and its complement. Level-shifted copies of these signals, $V_{\text{DIFF,B}}$ and its complement, control E340 and E341 which are used to discharge the outputs. To drive up one of the outputs, the upper EFET connected to that output (E300 or E301) is turned on while the corresponding lower EFET (E340 or E341) is turned off. Current flows through the upper EFET while the output charges up. This continues until the gate-to-source voltage of the EFET drops to the threshold voltage, at which point the EFET turns off and current flow ceases (with the exception of sub-threshold conduction). To output a low level, the upper EFET is turned off and the lower one is turned on. The output will drop to the point set by the pair of diodes sitting below the lower EFET. The operation of the buffer is thus seen to be similar to that of CMOS logic, but is implemented without the benefit of complementary transistors. Like CMOS, this buffer stage operates with no static power dissipation. The lack of static power dissipation and the small FET sizes are the keys to the lowering the power consumption of the laser driver.

The differential signals used to drive the buffer stage are generated by the first stage of the laser driver based on a set of complementary logic-level inputs. It is important that the signals used to drive the final

-
1. It would be interesting to reverse the order of the cascoding. If this proves to be successful, it would allow lower voltage operation.
 2. Designers at Vitesse have recently filed for patent protection on similar buffer circuits which they have independently developed.
 3. By contrast, to maintain analog operation in a source follower, the gate-to-source voltage must be kept below the Schottky gate turn-on voltage. The FET must then be made wider to achieve the same level of drive.
-

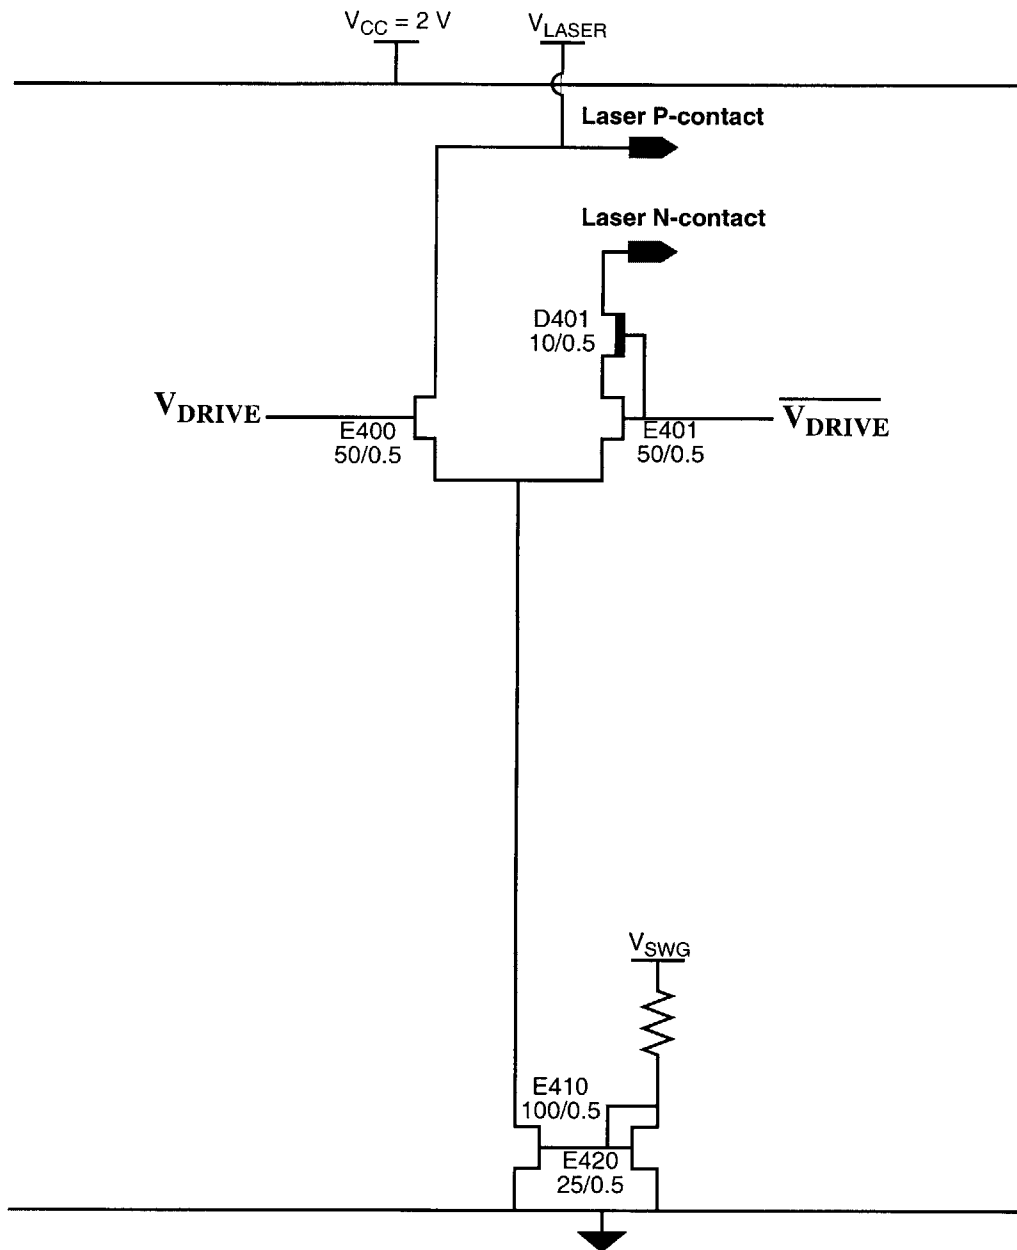


Figure 6.65: Third stage of the differential laser driver

The third stage of the laser driver is the differential current switch used to modulate the output current. The modulation current level is set by the current mirror and is fully switched between the two sides of the differential pair on the basis of the differential input signals. The output side of the diff-pair is cascoded by D401 to reduce the effects of drain lag resulting from the large resistive load of the VCSEL.

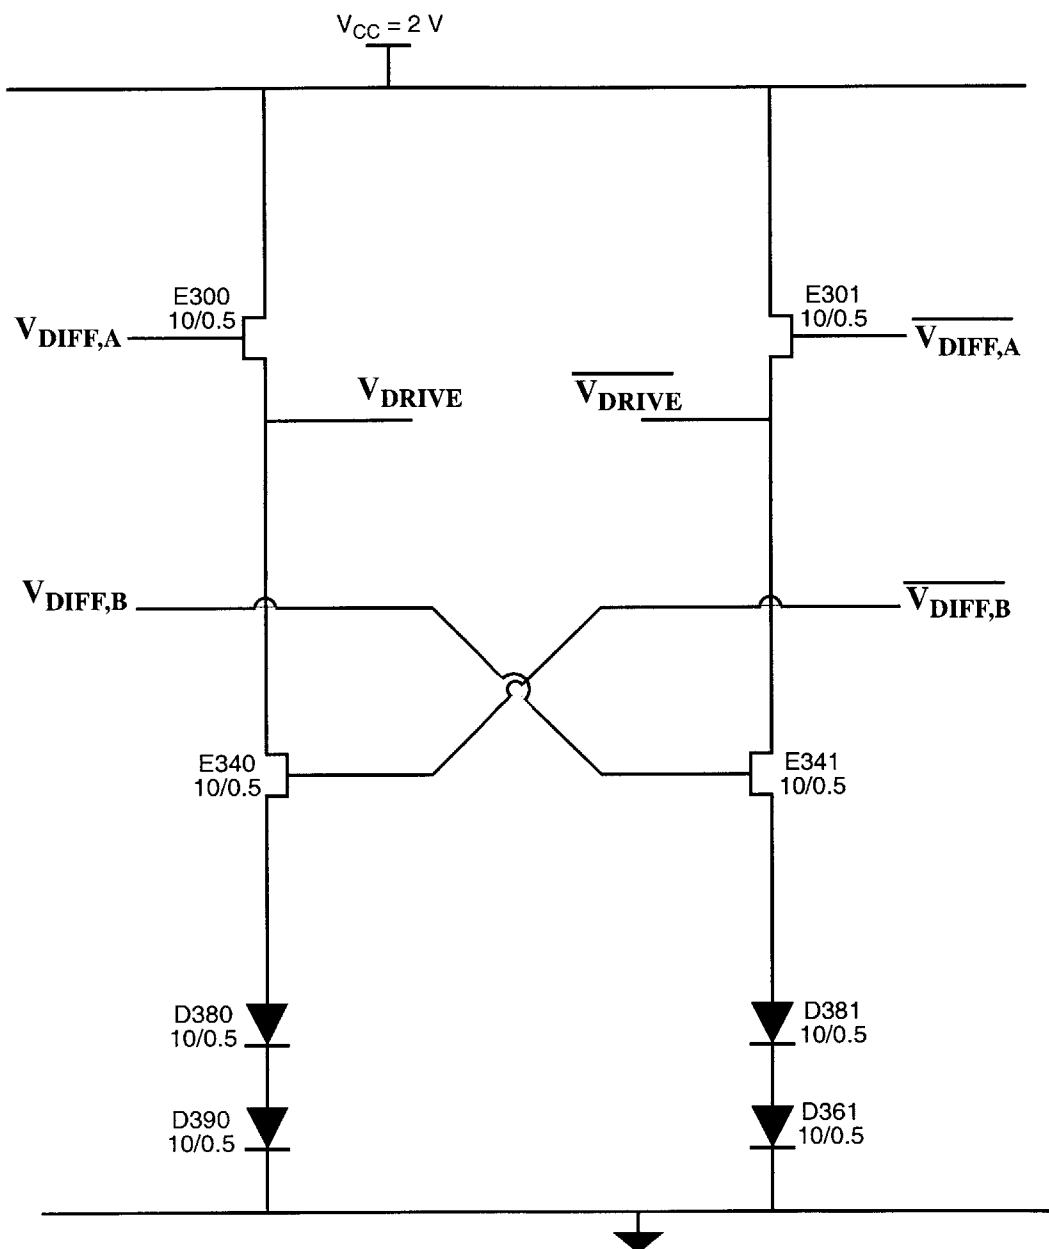


Figure 6.66: Second stage of the differential laser driver

The key to reducing the power consumption of the differential laser driver is to drive the capacitive load of the diff-pair with a very efficient driver stage. The operation of the stage shown here is analogous to that of CMOS logic, but without the benefit of complementary transistors. It has no static power consumption. It minimizes the loading on the previous stage by using the EFETs in a purely digital mode. That is, the Schottky gates are turned on fully in order to allow maximum current flow through the channel.

current switch, and thus the buffer stage, be well synchronized. A delay between them will result in both E400 and E401 being momentarily off. This causes a sharp glitch in the modulation current which appears at the output and causes ground bounce. The goal of the first stage is thus to re-time the logic inputs and to form differential signals of the appropriate amplitudes and bias voltages. The stage is shown in Figure 6.67. Its operation is based on the cascade of transconductance and transimpedance amplifiers, as used in the receiver discussed in Appendix H. Here, E200 and E201 form a differential-pair voltage amplifier with load resistances R200 and R201. D250 and D251 provide level shifting in order to generate a second set of differential outputs. The lower, “B”, outputs are buffered by E280 and E281 and fed back to the inputs of the gain stage through R100 and R101. This forms a fully-differential transimpedance amplifier (TIA). E100 and E101 are then used to generate complementary current inputs to this TIA. Now, the bias current on the voltage amplifier can be made relatively small because of the small size of the load presented to it by the buffer stage. Thus, glitches in the bias current of this stage which result from imperfect timing at its inputs are of little consequence.

The static characteristics of the full circuit shown in Figure 6.63 and a transient simulation of the current output is shown in Figure 6.68. The circuit was loaded by the VCSEL equivalent input circuit of Figure 6.16. Output drain lag effects are seen to be negligible, and the transitions have a rise/fall time of 200 ps with a 300 ps delay with respect to the input.

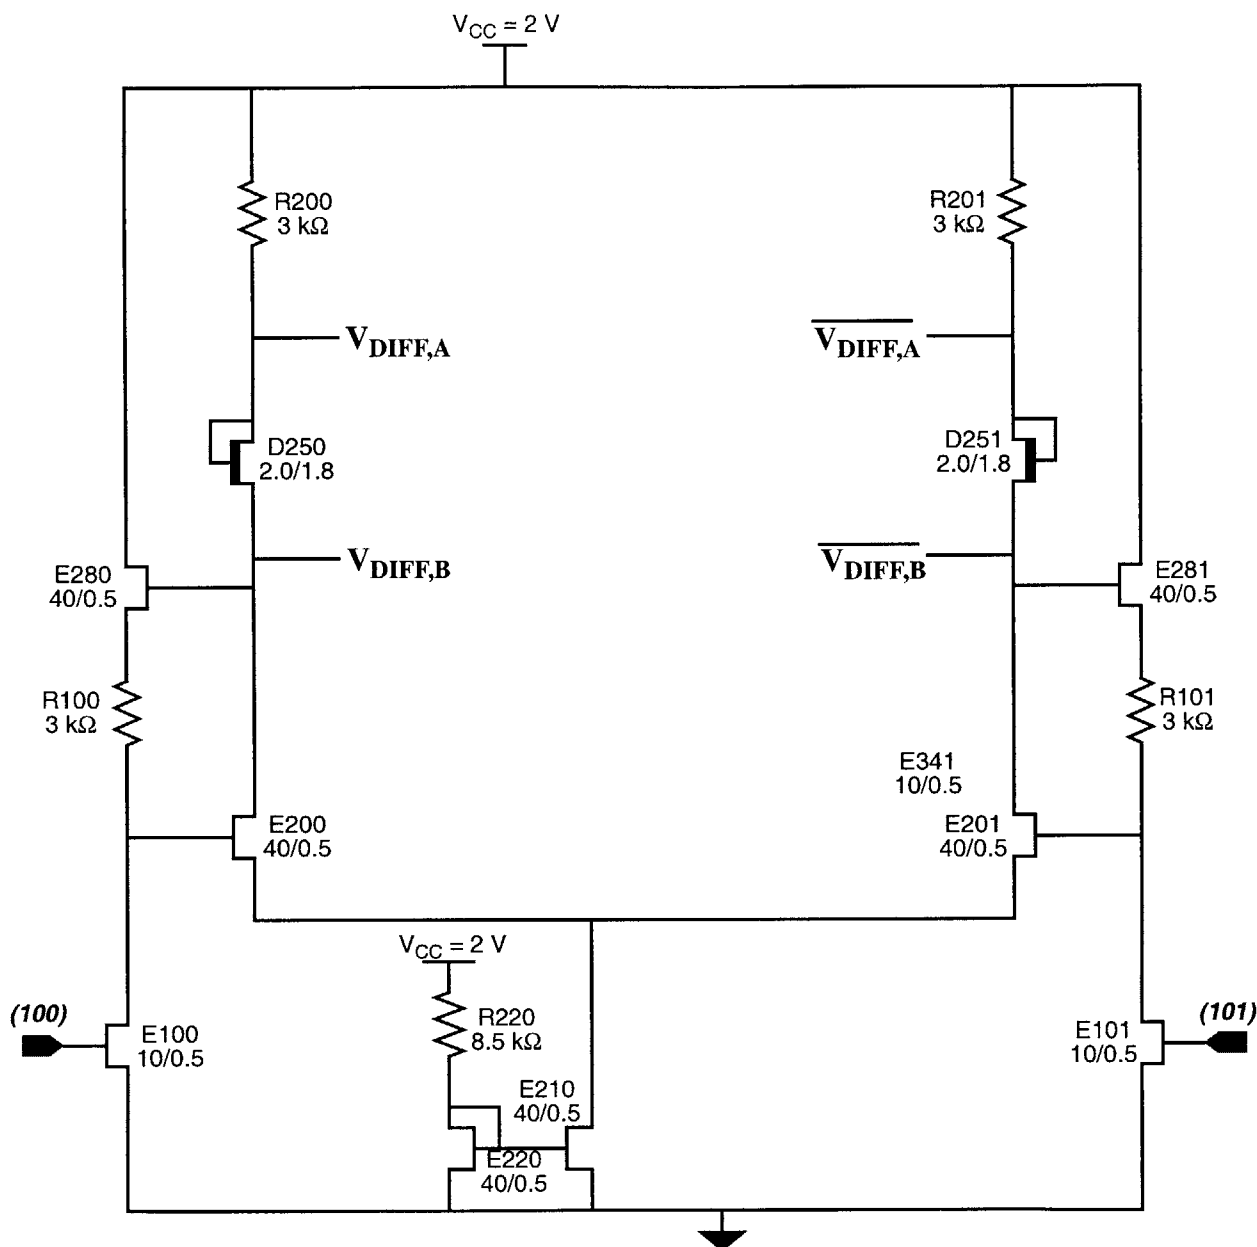


Figure 6.67: First stage of the differential laser driver

Complementary logic signals produced by straightforward means, such as the circuit shown in Figure 6.64, are not adequately synchronized to drive the differential current switch of Figure 6.65 without producing a large current glitch when both sides of the differential pair are momentarily off. The circuit shown here uses a cascade of transconductance (E100 and E101) and transimpedance (the rest of the circuit) stages to generate synchronized, differential outputs from logic inputs. The small loading of the second stage (Figure 6.66) allows the bias current on the transimpedance amplifier (TIA) can be made small. The glitch in the TIA's bias current which results from the imperfect timing of its inputs is thus reduced to a tolerable level.




Figure 6.68: Simulation of the differential laser driver shown in Figure 6.63

The static characteristics are shown along with a transient simulation of the output current. The input signal is shown for reference. The output was loaded by the VCSEL equivalent input circuit shown in Figure 6.16, and V_{swg} was set to 1 V. An off-state current mirror was not included in the circuit in this simulation.

6.9 Conclusion

This chapter has explained the design of laser drivers based on a switched current mirror. A low-power, compact laser driver of this type was demonstrated which is suitable for use in a 1 Gb/s digital optical interconnect. A number of limitations in the design were identified, and an improved version of the circuit was introduced. The switched current mirror laser drivers are particularly attractive in applications which call for hundreds or thousands of optical outputs because of their low power consumption. In addition, they minimize the laser supply voltage that is required.

An additional laser driver, which uses a differential pair as a current switch, was also described. This circuit achieves a power savings relative to conventional laser drivers of its type by making use of a novel driver stage featuring zero static power dissipation. This laser driver may not achieve as low a power consumption as the switched current mirrors, but avoids many of the subtleties in their design. This leads to a more robust circuit which is likely to be preferable in applications requiring a smaller number of optical outputs. The design concepts demonstrated should also scale to reduce the power consumption of laser drivers used for telecommunication-grade lasers. These lasers require modulation currents of 50-80 mA rather than ~ 1 mA required for a VCSEL in a short-haul link.

7

A Low Power Optical Receiver

The purpose of the optical receiver is to generate a logic output on the basis of an optical input. To this end, the photocurrent generated by a detector must be compared to a reference level to determine if the optical signal is high or low. For experimental purposes, the reference should be externally variable, though in practice it may be hard wired. This chapter deals with a continuous time receiver in which an amplifier/comparator continually produces a valid logic output reflecting its input. Another type of receiver, which latches its input value in response to a clock signal, is examined in Appendix I.

Reports on receiver design, both in Si and GaAs electronics, are most often concerned with optical communication systems and are not directly applicable here [213-221]. Receiver design for use in optical interconnects have also been reported based on CMOS electronics [201,222]. Again, this does not directly transfer to the GaAs-based EoE technology. Furthermore, these efforts rely exclusively on single-ended receiver designs and do not address the issues of supply and substrate noise which are abundant in the VLSI environment.

7.1 Design of a Low-Power Continuous-Time Optical Receiver

The goal of a continuous time optical receiver in a digital interconnect is to convert the small photocurrent at its input into a valid logic-level output, following a minimal propagation delay, and to maintain this output state until the next input transition. It may be characterized by a number of parameters including delay, delay variation, supply noise immunity, sensitivity, power dissipation, and area. The delay, measured, for

instance, between the midpoint of the input and output transitions, should generally be minimized.¹ Delay variation may be static, arising from the variations in processing and in operating temperature, or they may result dynamically based on the input signal and on the noise level. The dynamic delay fluctuations are referred to as jitter. Noise may be present in the input signal, inherent in the receiver circuit, or may be coupled in from the power supplies. The latter source of noise is a particular concern here since the receiver is intended for use in a VLSI environment. Sensitivity refers to the minimum signal power (or amplitude) that is required to prevent the output from being perturbed from its steady state value by the effects of additive noise². Finally, power dissipation and area must both be minimized. A power budget of 5 mW was established for the receiver in Chapter 5.

Design overview

Figure 7.1 is a block representation of the receiver examined in this chapter. It contains three gain stages to amplify the photocurrent from a detector at one of its inputs. The amplified difference between the input current and a reference current applied at the other receiver input is used to produce a direct-coupled FET logic (DCFL) output. The details of the receiver are shown schematically in Figure 7.2.

A key factor in the design of this receiver is the desire to reject supply noise while minimizing power dissipation and area. A common approach to supply noise rejection is to use fully differential amplifiers throughout the receiver. This was the strategy taken in the initial receiver design described in Appendix F. To implement this approach, level shifting source followers were needed between each of the differential gain stages. The source followers consumed a substantial amount of power and area (roughly as much as the amplifiers). The current design compromises some of the noise immunity of the fully-differential approach in order to dramatically reduce power and size. The first two stages are built as matched pairs of single-ended amplifiers while a differential amplifier is retained as the third stage. In this arrangement, the

1. In some applications the actual delay, or latency, is not a key design objective. For example, a pipelining technique may be used to allow for a large latency. In such cases, however, the delay variation is still very important.

2. That is, once the output has settled to the logic state corresponding to the transmitted bit value, it should remain there until the next input transition. The probability that this condition is not met at any given instant in time after the output has settled corresponds to the bit-error-rate (BER) calculated in a communication system context. In the latter case, only discrete-time samples of the received signal are used to determine the transmitted bit sequence, whereas in a digital interconnect, the input and output of the receiver are continuous-time objects.

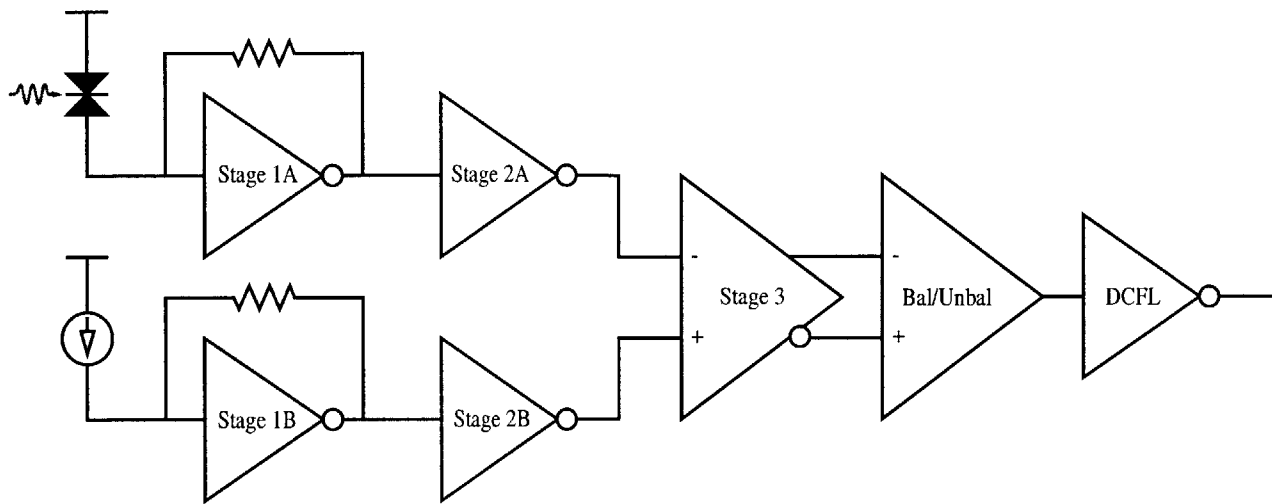


Figure 7.1: Block representation of a low-power continuous time receiver

Block representation of the low power optical receiver detailed in Figure 7.2. The receiver uses three gain stages. The third stage is fully-differential while the first two consist of pairs of matched single-ended amplifiers. The first stages are configured as transimpedance amplifiers. The third stage differential outputs are combined to form a single-ended signal which drives a DCFL inverter to generate the final logic output.

bias levels of each stage can be set to allow direct coupling of the stages without any intervening source followers. Although supply noise is readily coupled into the output of the first and second stages, it appears as a common mode signal and is rejected by the third, differential stage. The outputs of the third stage are then combined to form a single ended signal. At this point, the signal level is assumed to be large compared to the additive noise.

The receiver's final output is produced by the DCFL inverter formed by E500 and D500. The receiver may be conceptualized as a nearly ideal comparator (the DCFL inverter) preceded by the linear gain stages. The inverter has a nominal trip point of 400 mV, as seen in the simulated static characteristics plotted in Figure 7.3. The balanced-to-unbalanced conversion circuit is designed to produce $V_{411}=400$ mV when the third stage outputs are equal ($V_{400}=V_{401}$). This corresponds to equal signal and reference input currents. Process and temperature variations cause relative changes between the V_{411} bias point and the DCFL inverter trip point. This leads to variation in the propagation delay of the receiver. To minimize this static delay variation, D400/D401 and E410/E411 of the bal/unbal circuit are sized similarly to D500 and E500 of the DCFL inverter. Additional delay variations resulting from changes in the bandwidths of the gain stages.

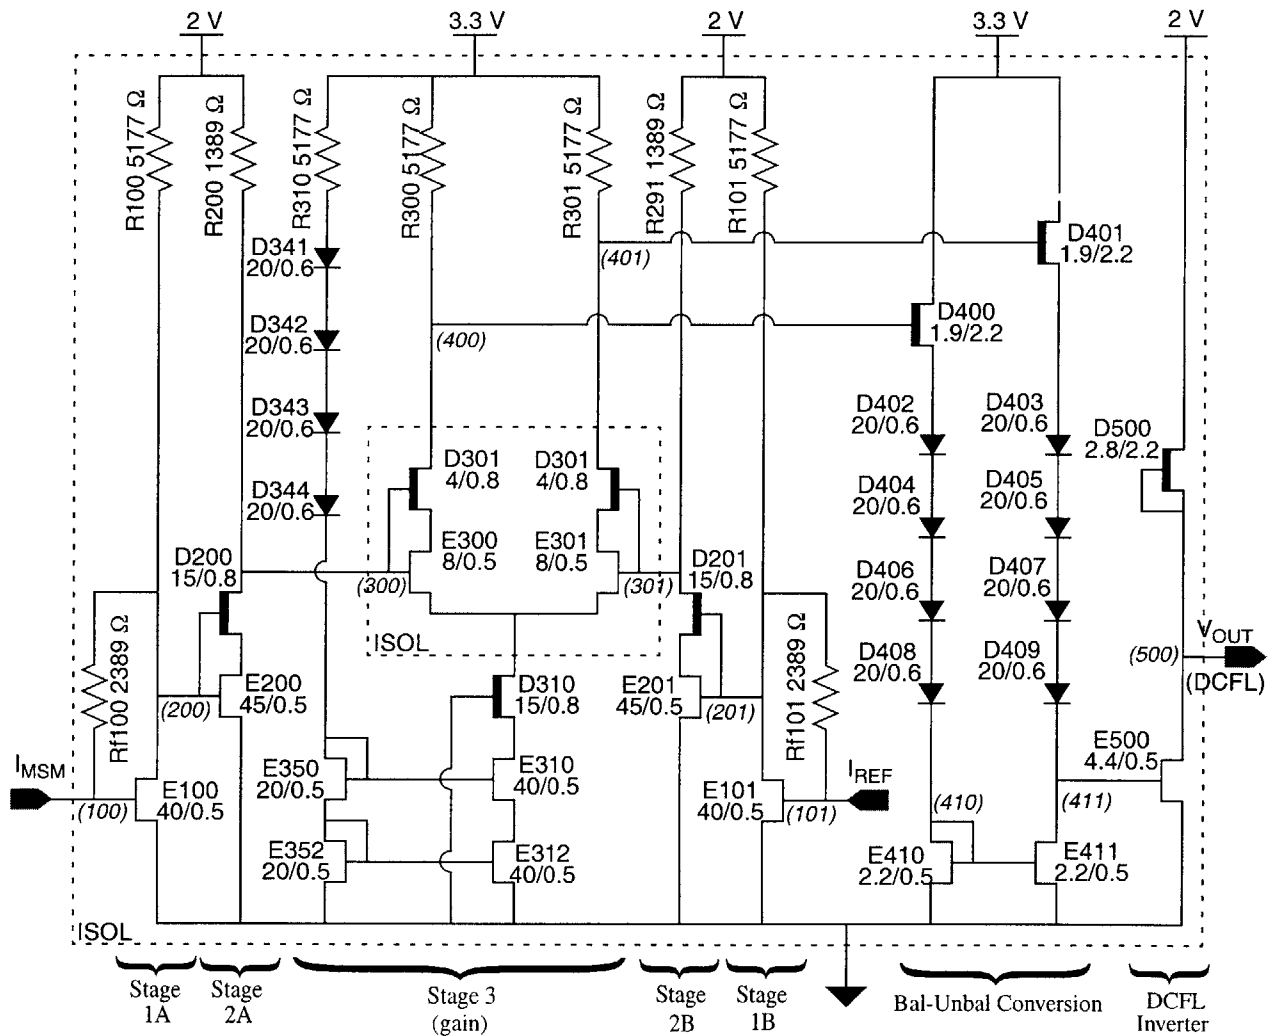


Figure 7.2: Schematic of a low-power continuous time receiver

Continuous-time optical receiver circuit shown symbolically in Figure 7.1. The first stage is a symmetric pair of single-ended transimpedance amplifiers. They are followed by a pair of symmetric voltage amplifiers. The third stage is a differential voltage amplifier. Its balanced outputs are converted to an unbalanced signal which drives a DCFL inverter. Node identification numbers are given in parentheses for reference. The resistor values were extracted from the circuit layout.




Figure 7.3: Simulated static characteristics of low power receiver

The reference current was set to zero, and the input current was swept. The output was loaded by a DCFL inverter.

Intersymbol interference and jitter

Assuming the input amplitude is chosen so as to meet the desired static error probability requirement¹, the key performance consideration is the jitter at the output of the receiver. Appendix D describes two principal mechanisms by which jitter appears in the receiver output: intersymbol interference and additive noise. Considering the receiver output as the product of a comparator (the DCFL inverter), a simple model gives the standard deviation and peak-to-peak jitter as

$$\sigma_O^2 = \sigma_{411}^2 + \frac{\tau^2}{SNR} + \sigma_D^2 \quad (7.1)$$

$$\Delta T_O = \Delta T_{411} + \frac{\Delta V}{A}\tau + \Delta T_D = \Delta T_{411} + \Delta V \left(\frac{d}{dt} V_{411}(t) \Big|_{V_{411} = 400mV} \right)^{-1} + \Delta T_D \quad (7.2)$$

where σ_O is the standard deviation of the output jitter, σ_{411} is the standard deviation of the jitter in V_{411} , τ is the rise/fall time of V_{411} , SNR is the signal-to-noise ratio of V_{411} , σ_D is the standard deviation of the delay through the inverter, ΔT_O is the output peak-to-peak jitter, ΔT_{411} is the peak-to-peak jitter in V_{411} , $\Delta V/A$ is the peak-to-peak noise normalized to the peak-to-peak amplitude of V_{411} , and ΔT_D is peak-to-peak variation in the delay through the inverter. These expressions show how additive noise is converted to jitter by the action of the comparator.

The jitter in V_{411} includes the jitter present in the transmitted signal in addition to a term resulting from intersymbol interference (ISI). In a digital interconnect context, intersymbol interference occurs when a signal does not settle within one bit period. The accumulated effect of the response tails leads to eye closure. Appendix D approximates the peak-to-peak jitter caused by intersymbol interference as

$$T_\Delta = \frac{V_\Delta}{A}\tau \quad (7.3)$$

The ratio V_Δ/A is the fractional eye closure in V_{411} measured immediately prior to the transition.

These expressions show what is required to minimize jitter in the receiver output: 1) V_{411} must settle within one bit period in order to prevent ISI, and 2) V_{411} must make a rapid transition past the inverter trip

1. In a digital interconnect, the error probability, or bit-error-rate, should be vanishingly small. The meaning of vanishingly small depends on the bit rate and the required mean-time-to-error. For example, if the link is included in a 1 GHz computer which is to operate for at least 1 year between hardware errors, the bit error rate on the link must be significantly less than 3×10^{-17} , that is, one divided by the number of bits transmitted in one year.

point in order to reduce the influence of noise on the output. In practice, these requirements must be compromised along with the other design goals.

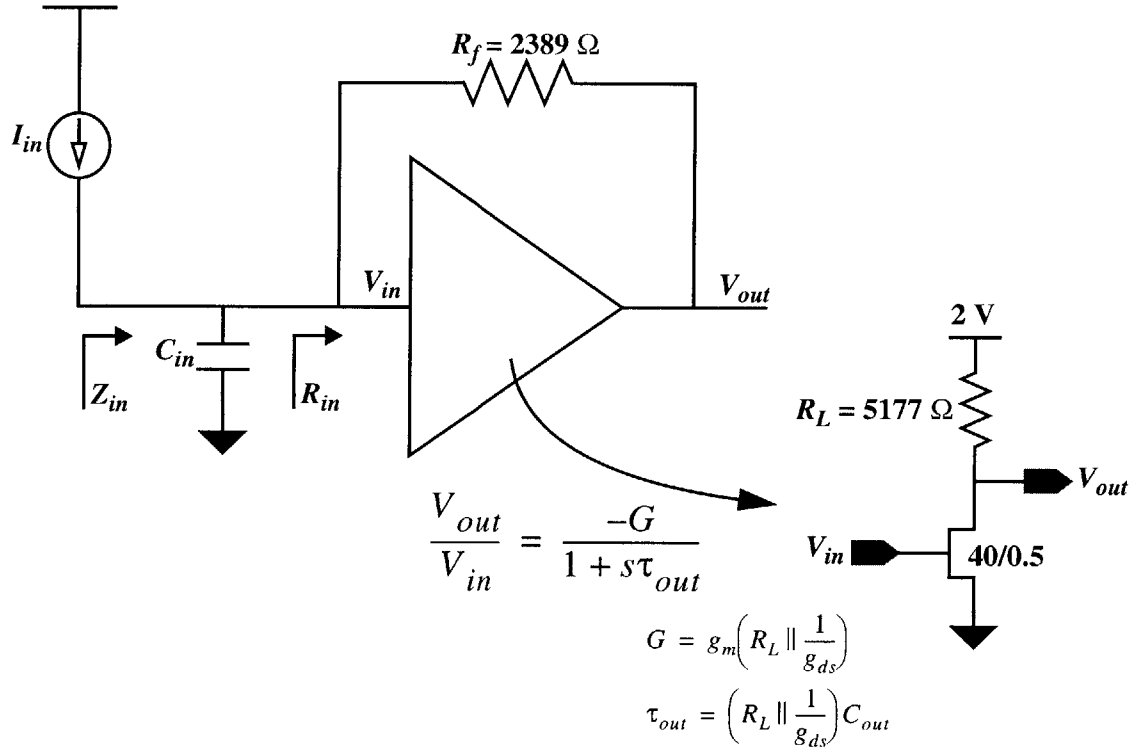
First stage: transimpedance amplifier

A simplified analysis of the first stage, a transimpedance amplifier (TIA), is given in Figure 7.4. The common-source gain stage, around which the TIA is built, is modeled as single pole transfer function with a bandwidth determined by the RC combination at its output. The use of feedback in the TIA results in an input resistance of $R_{in} = R_f / (1 + G)$. This sets the RC time constant at the input to $\tau_{in} = C_{in} R_{in} = C_{in} R_f / (1 + G)$, where C_{in} is the sum of the detector and amplifier capacitances. An additional time constant, τ_{out} , appears at the output of the amplifier. However, it is approximately the unity gain bandwidth¹ of the amplifier, $1 / \tau_{GBW} = (G + 1) / \tau_{out}$, that determines the TIA bandwidth along with τ_{in} . The TIA response becomes critically damped by making τ_{in} approximately equal to τ_{GBW} . In this case the TIA bandwidth is approximately determined by the sum of these time constants².

An additional point of interest is the use of a single stage amplifier rather than a two stage design. This choice leads to an adequate phase margin to ensure the stability of the circuit. A two stage design would have greater gain, and thus a lower input resistance. However, maintaining adequate phase margin in such a design was found to be very difficult in the Vitesse process.

The simulated frequency response of the TIA is shown in Figure 7.5. In Vitesse's MESFET process, a single-pole model of the common source amplifier is only approximately valid due to the presence of an effect known as "drain lag". As reviewed in Appendix E, drain lag may be explained in terms of capacitive coupling of the drain signal to the backgate which is grounded through a finite resistance. At high frequencies, the drain signal drives the backgate which adds to the drain current. This results in an apparent increase in the drain conductance. In a basic amplifier stage this translates into a drop in voltage gain. This effect is seen in Figure 7.5(c). The amplifier has a low frequency gain of 14.2 dB (5.1) but drops for around a decade starting at 2.5 MHz and remains level at its high frequency gain of 11.6 dB (3.8) up to the 813 MHz bandwidth determined by the output pole. In spite of these non-idealities, the single pole model is

1. The time constant $\tau_{GBW} = (G + 1) / \tau_{out}$ does not correspond exactly with the unity gain bandwidth. The unity-gain bandwidth would be $(G - 1) / \tau_{out}$. Since the gain is not that large, the distinction is not irrelevant. Conceptually, however, it is appropriate to think about the amplifier in terms of its unit-gain bandwidth or its gain-bandwidth product, G / τ_{out} .
2. More precisely, the critically damped TIA bandwidth is $(1 / (R_f C_{in}) + 1 / \tau_{out})$. Conveniently, this is nearly equal to $1 / (\tau_{GBW} + \tau_{in})$!



$$\frac{V_{out}}{I_{in}} = -\left(\frac{G}{1+G} R_f \right) \left\{ \frac{1}{1 + s[\tau_{in} + \tau_{GBW}] + s^2 [\tau_{in} \tau_{GBW} (1+G)]} \right\}$$

$$\tau_{in} = R_{in} C_{in} \quad \tau_{GBW} = \frac{\tau_{out}}{1+G} \quad R_{in} = \frac{R_f}{1+G}$$

$$Z_{in} = \left(R_{in} \frac{1 + s\tau_{out}}{1 + s\frac{\tau_{out}}{1+G}} \right) \parallel \left(\frac{1}{sC_{in}} \right)$$

Figure 7.4: Simplified analysis of transimpedance amplifier used as the first stage.

The transimpedance gain and input impedance are derived assuming a single-pole model for the common-source gain stage.




Figure 7.5: Simulated frequency response of first stage of low power receiver

(a) The transimpedance gain (dB Ω) and phase (degrees). (b) The input impedance magnitude (dB Ω) and phase (degrees). (c) The voltage gain (dB) and phase (degrees) of the amplifier used to build the TIA. In input capacitance of 100 fF was assumed.

useful in understanding the TIA. Note, also, that even with the inclusion of drain lag, the single stage amplifier has adequate phase margin to build a feedback circuit without running into stability problems.

Three techniques are used to control drain lag effects in the receiver design. The first of these, feedback, arises naturally into the design of the TIA. As seen in Figure 7.5(a), the transimpedance gain is flat at $66 \text{ dB}\Omega$ ($2 \text{ k}\Omega$) out to a bandwidth of 1.5 GHz ¹. The drop in the gain of the voltage amplifier shows up only in the input impedance, which rises from $51.8 \text{ dB}\Omega$ (389Ω) to a peak of $58 \text{ dB}\Omega$ (794Ω) before dropping under the influence of the input capacitance.

The above simulation assumed a capacitance of 100 fF for the detector. This is a conservative estimate of the capacitance of a $75 \mu\text{m}$ diameter MSM photodetector [148].

Second stage: voltage amplifier

Because of the limited intrinsic gain of Vitesse's MESFETS (particularly in light of drain lag)², the transimpedance gain of the first stage ($\sim R_f$) can not be made large enough (while maintaining the required bandwidth) to directly drive the final comparator. Additional voltage gain stages are required, and the effect of drain lag on these stages must also be controlled. The frequency response of Figure 7.5(c), for instance, would lead to a large amount of intersymbol interference. The use of feedback on the voltage gain stages would require buffering between stages in order to drive the resulting low input impedances. Instead, the effects of drain lag are reduced by a second technique, the use of a cascode structure described in Appendix E. The stage-2 voltage amplifier is highlighted in Figure 7.6 and its simulated frequency response is plotted in Figure 7.7. The low frequency gain of this stage is 15.6 dB (6.03) which drops by 0.7 dB to a high frequency value of 14.9 dB (5.56). The bandwidth, measured relative to the low frequency gain, is 3.4 GHz .

Third stage: differential voltage amplifier

The third gain stage is shown isolated in Figure 7.8 and its simulated frequency response is plotted in Figure 7.9. This stage again uses a cascode to control the drain lag. It has a low frequency gain of 9.52 dB (2.99) and a high frequency gain of 8.88 dB (2.78). The inverted output (with respect to the input driven by

1. The unity-gain bandwidth of the voltage amplifier is 3.5 GHz and the 3 dB bandwidth of the input impedance is 4.4 GHz . This gives an estimate of the TIA bandwidth of 1.9 GHz which is close to the simulated value of 1.5 GHz .

2. The typical usable range amplifier gain is around 2 to 6 [209].




Figure 7.6: Isolated schematic of the low power receiver's second stage voltage amplifier. A cascode structure is used to control the effects of drain lag.

the receiver input signal) has a bandwidth of 1.5 GHz measured relative to the low frequency gain and the non-inverting output has a bandwidth of 2.2 GHz. The difference arises because of a small difference in loading conditions between of the two inputs of the balanced-to-unbalanced conversion circuit which follows the differential amplifier.

To improve the rejection of common mode signals, the current source used to bias the differential pair is cascoded. A DFET cascode similar to the amplifier itself is used in addition to a conventional cascode current mirror design.

An important design feature of this amplifier is the use of resistive loads matched to the resistor which determines the bias current. With a current mirror gain of two, the bias current in each leg of the differential pair is the same as the current through R310. Thus, the bias voltage on nodes 400 and 401 are effectively set by the sum of the voltage drops across the diodes (D341-344) and the gate-to-source voltages of E350 and E352. The dependence of these voltage drops on the bias current is relatively weak. As a result, the bias voltages on nodes 400 and 401 are fairly stable with respect to changes in supply voltage. The fabricate circuit was operated with the 3.3 V supply varied from 3.0 V to 5 V. The speed and sensitivity of the receiver increases at higher supply voltages as a result of the increase in stage-3's gain at the higher bias current.




Figure 7.7: Simulated frequency response of second stage of low power receiver (a) is the magnitude (dB) and (b) the phase of the stage-2 frequency response.

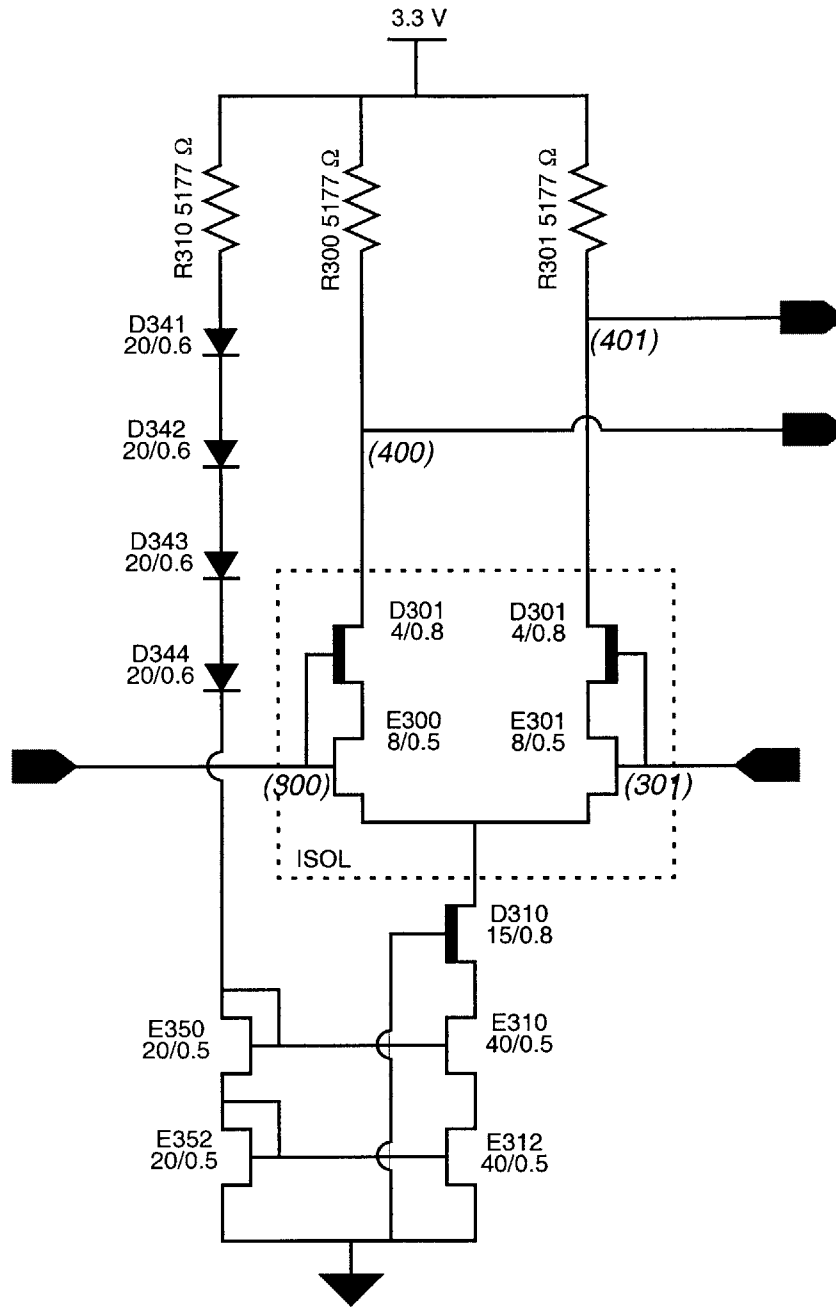


Figure 7.8: Isolated schematic of the low power receiver’s third stage differential amplifier

A cascode structure is used to control the effects of drain lag. The current source is also cascoded to improve common-mode rejection. The resistive loads are matched to the current reference for good bias stability over process and temperature variations.




Figure 7.9: Simulated frequency response of third stage of the low power receiver (b) is the magnitude (dB) and (a) the phase (degrees) of the stage three differential amplifier frequency response. An input is applied to node 300 while node 301 is constant. The solid lines are the response from node 300 to node 400 and the dotted lines are the response from node 300 to node 401.

Single-ended linear output and hard limiting

The balanced-to-unbalanced conversion circuit, isolated in Figure 7.10, combines the differential signals on nodes 400 and 401 to produce the single ended signal, V_{411} , which drives the final DCFL inverter. The simulated frequency response between the input current and the voltages on nodes 400, 401, and 411 is plotted in Figure 7.11. The low frequency gain at nodes 400 and 401 is $91.1 \text{ dB}\Omega$ ($35.9 \text{ k}\Omega$) and the high frequency gain is $89.3 \text{ dB}\Omega$ ($29.2 \text{ k}\Omega$). The bandwidth is 860 MHz on node 400 and 1.1 GHz on node 401. Ideally, the bal/unbal circuit should add the two signals, resulting in an additional 6 dB of gain. However, a significant amount of signal is lost due to the finite resistance of the level shifting diodes. The final low frequency gain on node 411 is $93.9 \text{ dB}\Omega$ ($49.5 \text{ k}\Omega$), and the high frequency value is $91.1 \text{ dB}\Omega$ ($39.9 \text{ k}\Omega$), a difference of 2.8 dB. Measured relative to the low frequency gain, the bandwidth is 280 MHz, but relative the high frequency gain, the bandwidth is 1.0 GHz.

As will be seen momentarily in transient simulations, the bandwidth on nodes 400 and 401 are adequate to avoid significant intersymbol interference effects. This is not true, however, of node 411. The two-tiered frequency response on this node means that the step response of V_{411} would be expected to rise quickly ($\sim 350 \text{ ps}$) to a value determined by its high frequency gain, then continue to rise on a longer time scale to its final, static value. The 2.8 dB difference between the low and high frequency gains may be expected to cause $\sim 40\%$ eye closure!

A third technique is used to avoid this drain-lag induced ISI: hard limiting. The high frequency gain results in a 0.5 V step in V_{411} in response to a $13 \mu\text{A}$ input. V_{411} can not grow beyond this amplitude because it is clamped above by the gate of E500 and below by ground. Once it is driven to its limits, node 411 ceases to behave as a linear system. It “forgets” the signal history, eliminating ISI¹.

An important consequence of the need to hard limit V_{411} is that it sets a minimum input amplitude. This is independent of the steady-state sensitivity discussed above, and is very likely to be greater than it. In effect, this is a “dynamic” sensitivity which accounts for the receiver's tendency to create jitter at high bit-rates.

1. This is, in fact, how DCFL logic circuits function, in the presence of drain lag, without experiencing ISI effects.

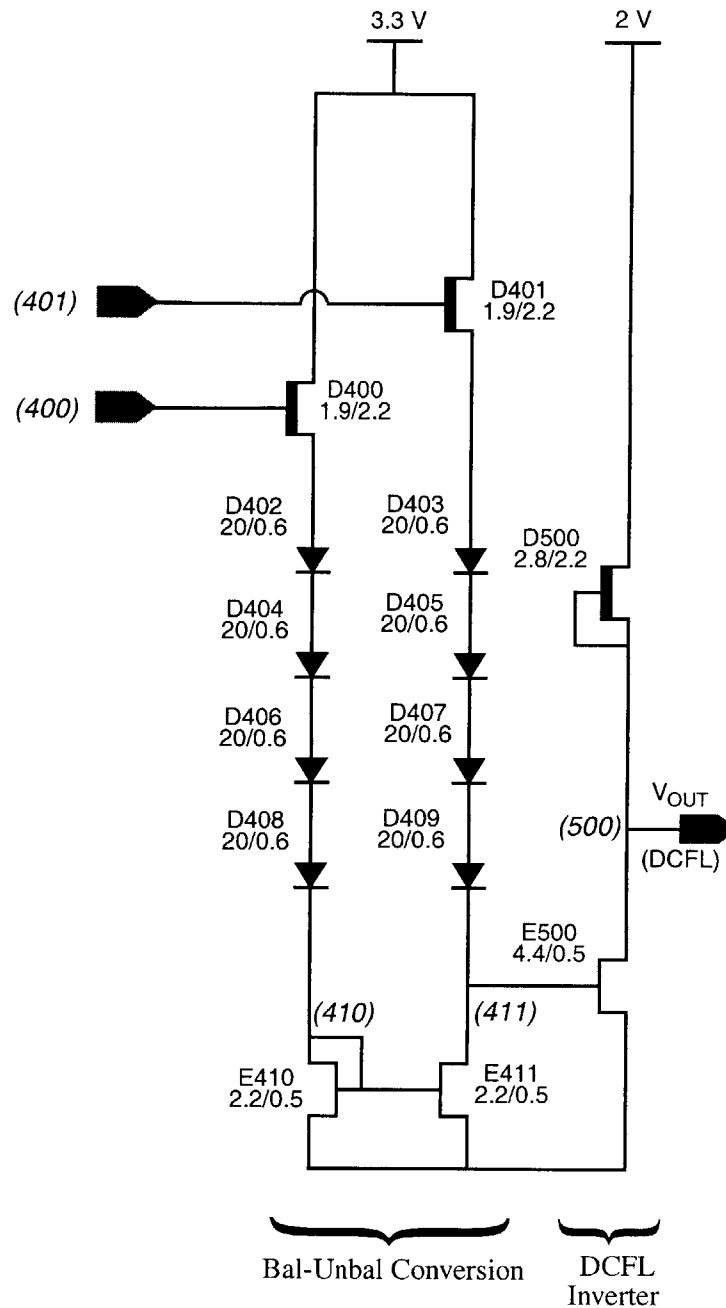


Figure 7.10: Isolated schematic of balanced-to-unbalanced conversion circuit and DCFL inverter

The bal/unbal circuit combines its inputs by mirroring the signal from node 400 and adding it to that of node 401. With $V_{400} = V_{401}$, the circuit is designed to set V_{411} equal to the DCFL inverter's trip point. The transistors in this circuit are sized to nearly match the DCFL inverter so that the bias point on node 411 tracks the trip point of the inverter across process and temperature variations.




Figure 7.11: Simulated frequency response after first three stages and bal/unbal conversion

The magnitude of the frequency response is given in dBΩ. The data in (a) is shown more closely in (b). The solid line represents node 411 and the dashed and dotted lines are nodes 400 and 401.

The use of hard limiting also effects how the gain and bandwidth of the receiver are selected. For a given input amplitude and bandwidth, increasing the gain beyond the point at which V_{411} hard limits effectively decreases the V_{411} rise time, and along with it the receiver delay and jitter. For the receiver presented here, increasing the gain, at the expense of bandwidth, was found to be more optimal with regard to minimizing power consumption.

DCFL inverter at the output

Since V_{411} is designed to make a full DCFL swing, the final DCFL inverter is not entirely necessary. However, it does serve to buffer node 411, sharpens the output transition, and adds minimally to the power, delay, and area of the circuit.

Transient response

A simulation of the receiver's transient response is shown in Figure 7.12. The input signal steps between 0 and 16 μA in 200 ps and the reference current is set to 8 μA . An input capacitance of 100 fF is assumed, and the output is loaded by a DCFL inverter. Details of the input-rising and -falling transitions are shown in Figure 7.13 and Figure 7.14 respectively. These plots show that the receiver is expected to exhibit only a small amount of ISI (see the drain-lag induced response tail in V_{400} and V_{401}). The 16 μA input amplitude is adequate to hard limit V_{411} so that the large drain-lag predicted by the frequency response is countered. In regard to the ISI on nodes 400 and 401 and the 16 μA dynamic sensitivity, the transient simulation is slightly more pessimistic than the frequency responses, which only give information about the quiescent operating point of the receiver. In spite of the rather limiter bandwidth, however, the receiver is predicted to have a propagation delay of 350 ps.

Delay and static delay variation

Under typical process conditions and at 85°C, the propagation delay is 350 ps for both transition polarities. Over a combined temperature variation of 25°C to 125 °C and a process spread of two standard variations, the input-rising-to-output-falling delay varies from 290 ps to 410 ps and the input-falling-to-output rising delay varies from 310 ps to 440 ps.

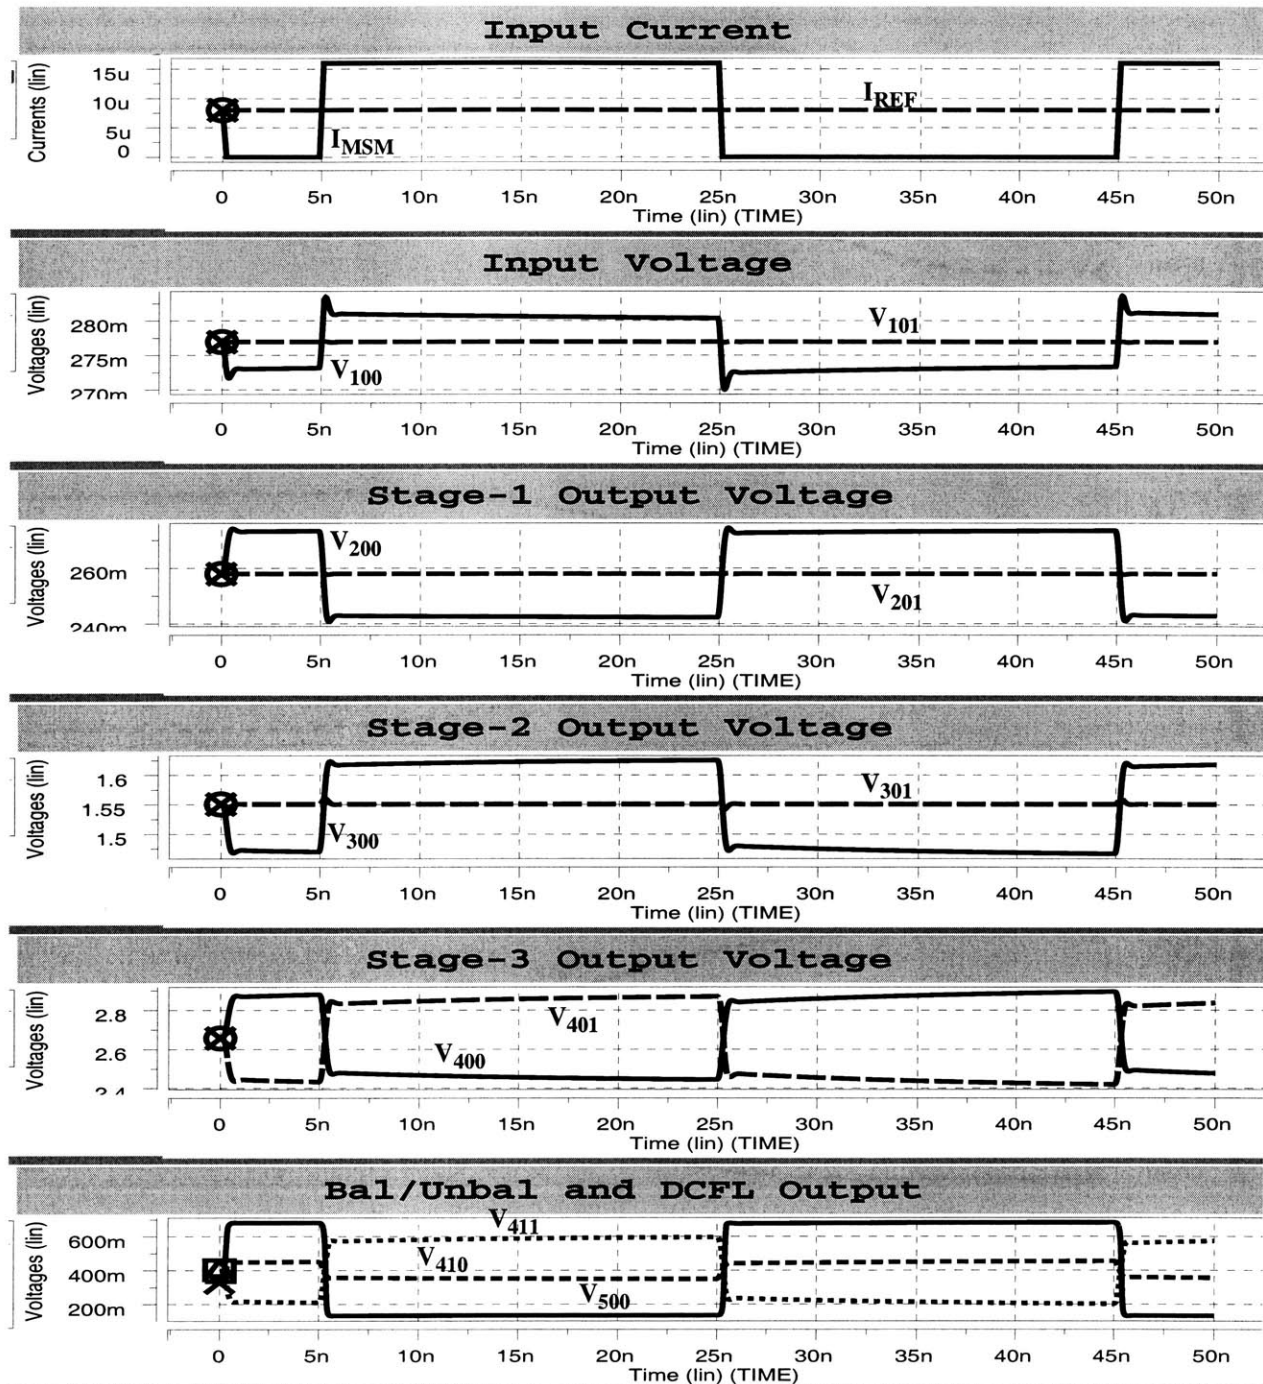


Figure 7.12: Simulated transient response of the low power receiver

The input current ramps from 0 to $16 \mu\text{A}$ in 200 ps while the reference current is held at $8 \mu\text{A}$. A detector capacitance of 100 fF is assumed at the input and the output is loaded by a DCFL inverter.




Figure 7.13: Close-up of the receiver transient simulation for a rising input transition

The simulation results in Figure 7.12 are replotted to show the details of the waveforms following the input rising edge.




Figure 7.14: Close-up of the receiver transient simulation for a falling input transition. The simulation results in Figure 7.12 are replotted to show the details of the waveforms following the input falling edge.

Power dissipation

Simulated under typical process conditions and 85°C operating temperature, the receiver dissipates 4.5 mW. The power dissipation varies between 2.2 mW and 5.9 mW over a combined temperature variation of 25°C to 125 °C and a process spread of two standard variations.

Noise analysis

Vitesse's HSPICE models do not include a noise analysis component. In lieu of such an analysis, and assuming the total receiver noise is dominated by that of the first stage, the input referred noise current may be written [158,194-196,223-226]

$$\langle i_{in}^2 \rangle = \frac{4kT}{R_f} B + 2q(I_{gate} + I_{dark})B + \frac{4kT\Gamma}{g_m} \left\{ \frac{1}{R_f^2} B + [2\pi(C_{det} + C_{gs} + C_{gd})]^2 \frac{B^3}{3} \right\} \quad (7.4)$$

The first term is the thermal noise of the feedback resistor, the second term is the shot noise due to gate leakage and detector dark current, and the third term is the thermal noise in the FET channel¹. Here, $B \sim 1$ GHz is the bandwidth, $R_f = 2389 \Omega$ is the feedback resistor, $I_{gate} = 30$ nA is the gate current of the input FET (E100), $I_{dark} < 1$ nA is the dark current (and may be neglected), $g_m = 3.9$ mS is transconductance and $C_{gs} = 23$ fF and $C_{gd} = 13$ fF are gate capacitances of E100, $C_{det} \sim 100$ fF is the capacitance of the detector, and Γ is the FET noise factor and which has a value of around 1.8 for a GaAs MESFET [158]. Thus, the feedback resistor's thermal noise variance (first term) is computed as $6.7 \times 10^{-15} \text{ A}^2$, the shot noise variance (second term) is $9.6 \times 10^{-18} \text{ A}^2$, and the channel thermal noise variance (third term) is $1.9 \times 10^{-15} \text{ A}^2$. The total RMS input referred noise is thus dominated by the thermal noise of the feedback resistor and has a value of 93 nA.

The input RMS noise current is amplified by the receiver to produce 3.7 mV of RMS noise on node 411. There is a similar contribution from the reference-side TIA bringing the total RMS noise on node 411 to 7.4 mV. With a 0.5 V, 100 ps transition on this node (see Figure 7.13 and Figure 7.14), (7.1) gives an

1. The MESFET drain current has a noise spectral density of $4kT\Gamma g_m$ [223,224]. This may be referred to the gate as an noise voltage source with spectral density $4kT\Gamma/g_m$. This source may be used in conjunction with the model shown in Figure 7.4 to evaluate the noise spectral density at the output of the TIA. This result may then be referred back to the TIA input current and integrated up to a bandwidth of B in order to find the third term in (7.4) in agreement with previous reports [225,226]. The more recent reference, [158], ignores the $1/R_f^2$ term since it is relatively small. Both the formula used here and in [158] ignore $1/f$ noise. Also, note that the channel thermal noise also induced a gate current with spectral density $4kT\Gamma'(2\pi C_{gs})^2 f^2/g_m$, but this term is significantly smaller than that of the first one and is thus ignored [224].

RMS jitter of 1.4 ps. Assuming the peak-to-peak noise is ten times the RMS value, (7.2) gives 14 ps for the peak-to-peak jitter.

The input optical signal also includes a noise component. For the 16 μA signal that is required to hard limit V_{411} , the optical signal's shot noise contribution to the receiver's RMS input noise current is 72 nA. This is comparable to the receiver's own noise contribution. The laser signal includes additional noise, however, arising, for instance, from the laser driver and from optical reflections into the laser. This can add substantially to the total noise. Thus, because of the receiver's dynamic sensitivity limitation, the optical input is actually expected to dominate the total noise in the system¹.

Supply noise rejection

Since this receiver is intended for use in a digital VLSI environment, the coupling of supply noise into the signal path, and in particular into V_{411} , is an important issue. As explained above, the receiver architecture is based on a compromise between supply noise rejection and power dissipation.

The magnitude of the simulated frequency response to fluctuations in the 2 V supply is shown in Figure 7.15. The effect on the first stage is relatively weak at -26 dB, but the second stage has only 3 dB of supply rejection. However, the supply noise output of the second stage amplifiers appears symmetrically on nodes 300 and 301. Thus, it is attenuated by the third stage differential amplifier to which it appears as a common-mode signal. The third stage's rejection of this common-mode signal is seen to be quite good at low frequencies, but worsens starting at 1 MHz. This is true even though a multiply cascoded current source is used to bias the differential pair. Additional work in the future may improve the common-mode rejection. Fortunately, the balanced-to-unbalanced conversion circuit provides additional common mode rejection which significantly reduces the supply noise on the critical V_{411} signal. To better gauge the effect of noise transients in the 2 V supply, Figure 7.16 plots the receiver's simulated response to 0.1 V steps on the 2 V supply with 10 ps rise/fall times. The effect on V_{411} is a 10 mV peak-to-peak glitch occurring on each supply voltage step.

The simulated frequency response to noise originating in the 3.3 V supply is shown in Figure 7.17. As pointed out above, the differential pair uses load resistors which are matched to the reference resistor

1. By contrast, in a communication system receiver there is no dynamic sensitivity issue, and the static sensitivity is determined by the input referred noise and the signal-to-noise ratio required to meet the specified bit-error-rate.




Figure 7.15: Simulated frequency response to variations in the 2 V supply

A detector capacitance of 100 fF is assumed at the receiver input. The input and reference currents are zero.




Figure 7.16: Simulated transient response to variations in the 2 V supply

The 2 V supply is ramped between 2.0 V to 2.1 V in 10 ps. A detector capacitance of 100 fF is assumed at the receiver input. The input and reference currents are set to zero.




Figure 7.17: Simulated frequency response to variations in the 3.3 V supply
The input and reference currents are zero.




Figure 7.18: Simulated transient response to variations in the 3.3 V supply

The 3.3 V supply is ramped between 3.3 V to 3.4 V in 10 ps. The input and reference currents are zero. The 100 mV of noise in the supply is seen to translate into 30 mV of noise on node 411.

which sets the bias current. This tends to stabilize the V_{400}/V_{401} bias. However, cascoding the current mirror in order to improve the amplifiers common-mode rejection tends to degrade the supply rejection. At low frequencies, it results in greater supply dependence because of the current dependence of V_{gs} on E350 and E352. The cascode also degrades the current mirror's frequency response, which defeats the ability of the biasing arrangement to stabilize the V_{400} and V_{401} at higher frequencies. A compromise is required between common mode and supply rejection which must take into account the relative magnitude of noise on the two supplies.

Once again, the balanced-to-unbalanced conversion circuit improves the noise rejection. The transient response to 0.1 V, 10 ps rise/fall, steps in the 3.3 V supply is shown in Figure 7.18. In this case, the noise amplitude on V_{411} is around 30 mV.

For comparable noise levels on the two supplies, the effect of noise coupling into the third gain stage is seen to be greater than the contribution of the first two stages. This indicates that the strategy of using symmetric single-ended amplifiers for the first two stages does not substantially effect the overall supply noise rejection. In a practical design in which corresponding stages are not perfectly matched, there may be some additional contribution from the first two stages.

Physical design

The receiver was design in Vitesse's four metal layer H-GaAs IV process. A layout plot of the circuit is shown in Figure 7.19. The cell occupies a $230\ \mu\text{m} \times 80\ \mu\text{m}$ area, and requires an external connection to the n^+ isolation ring which surrounds it.

The source and drain regions have been lengthened to $4\ \mu\text{m}$ in order to minimize contact resistance. According to the H-GaAs IV Technology Reference Manual, the nominal contact resistance drops from $750\ \Omega\cdot\mu\text{m}$ at a contact length of $1.2\ \mu\text{m}$ down to $204\ \Omega\cdot\mu\text{m}$ for contact lengths exceeding $2.8\ \mu\text{m}$, and contact lengths of at least $3\ \mu\text{m}$ are recommended when contacts are shared [208]. However, the junction capacitance increases as the source/drain regions are lengthened. Subsequent to the completion of the above design, it was learned that a contact length of around $2\ \mu\text{m}$ is believed to be optimal [209].

Resistors were implemented as source/drain implanted regions which have a nominal sheet resistance of $300\ \Omega/\square$. The temperature coefficient of variation of these resistors is reported to be zero [208], and the resistors are typically matched to within 2% across a die. The width of $3.0\ \mu\text{m}$ was chosen for the resistor




Figure 7.19: Layout plot of low power receiver

The inputs and output are connected on metal-2 at the points indicated. Metal-3 and metal-4, not shown, are present over the entire area and are used for the 2 V supply and ground, respectively. An n^+ isolation ring surrounds the entire cell and the third-stage differential pair. The isolation structure should be connected (to the most positive supply available) at the perimeter of the cell using metal-1.

implants in order to minimize across-die variability. However, there is around a 20% run-to-run variation in resistance. The ohmic contact length to the resistors was set at 4 μm . Whereas in the source/drain sizing the contact area is a performance issue, here it is one of predictability (i.e. by the extraction tool) and uniformity.

Liberal use was made of p-contacts in order to minimize the resistance to the backgate and to limit susceptibility to substrate noise. Each gain stage was surrounded by a p-contact ring. The differential-pair of the third stage was also surrounded by its own isolation ring. The p-contact in this isolated region was connected to the source of the differential-pair EFETs in order to reduce backgating effects¹. The p-contacts of the remainder of the circuit were grounded, and the entire circuit was surrounded by an n^+ isolation ring.

Around one-third of the layout in Figure 7.19 is occupied by diodes. In many cases, these may be shrunk in order to reduced the cell area without severe performance loss. It may also be possible to eliminate some of the diodes through a redesign of the circuit.

1. It was not known during the implementation of this receiver that Vitesse design rules require the substrate of circuits which include p-contacts at above-ground potentials to be left floating. Leaving the substrate floating may lead to difficulties in the presence of photogenerated carriers. In future designs, it may be advisable to ground all p-contacts. Grounding the backgate of the differential-pair transistors results in only a modest performance loss relative to connecting the backgate of these transistors to their source.

7.2 Measured Receiver Characteristics

This section presents measure static and dynamic characteristics of the low power receiver. The receiver was coupled with an integrated photodetector fabricated by Vitesse and detailed in Chapter 4. Unfortunately, this detector is found to limit the dynamic performance of the receiver, so that the receiver's actual performance limits are not known.

Experimental setup

The experimental setup used in characterizing the receiver is described in Appendix C. The receiver was tested in conjunction with the MSM27 photodetector detailed in Chapter 4. An on-chip ~ 200 k Ω resistor was connected to node 101 of the receiver in order to set the reference current. A Vitesse ECL output driver was used to buffer the receiver's DCFL output¹. The polarity of the ECL output is the same as the DCFL output of the receiver (node 500). The light source used to drive the receiver was an 850 nm VCSEL driven by an ECL signal. The characteristics of the VCSEL were given in Section 6.3. The same module was used to characterize the photodetectors in Chapter 4. Its output was coupled into a 62.5 μm graded index core multimode fiber. The amplitude of the optical signal was adjusted by attenuating it, and the fiber was butt-coupled onto the 75 μm diameter detector integrated with the receiver circuit. To examine the receiver's at-speed functionality, the optical source was modulated by a 2^{23} -1 bit pseudorandom sequence. The data source is the same one used in characterizing the laser driver in Chapter 6 (see Figure 6.8 and Appendix C). Additional details of the optical signal characteristics will be given as needed below.

Static characteristics

The decision level of the receiver is set by applying voltage V_{REF} to the ~ 200 k Ω reference resistor connected to node 101. The resulting reference current, I_{REF} , is plotted in Figure 7.20. V_{REF} can not be raised above ~ 6 V without increasing the probability of damaging the circuit².

Figure 7.21 shows the static input-output transfer characteristics of the receiver. The voltage applied to reference resistor, V_{REF} , was stepped from 1 to 6 V and the optical power was swept to generate the family

1. The driver operates from 0 to 2 V rather than from -2 V to 0. The output, which is end terminated by a 50 Ω resistor to ground, has a nominal swing of 0 to 1 V.
2. V_{REF} was set as high as ~ 20 V without immediate damage to the receiver, but this circuit subsequently failed.




Figure 7.20: Setting V_{REF} and I_{REF} for receiver characterization

The receiver decision level is set by the reference current, I_{REF} , flowing through a ~ 200 k Ω resistor connected to node 101. Here, I_{REF} is plotted as a function of the applied reference voltage, V_{REF} .

of curves plotted in Figure 7.21(a). The output voltage that is plotted is a time-average value. When the optical signal is far above or below the reference level, the output remains stationary at the output low or high values, respectively. As the mean optical power level comes near the decision level set by V_{REF} , the output begins to fluctuate rapidly between its high and low levels in response to noise. On a sampling oscilloscope, this is observed as “snow” filling the voltage range between the high and low output levels. The data plotted in Figure 7.21 is the average of this snow at each mean optical power. The width of the transition of this average output voltage from the high to the low level gives an estimate of the peak-to-peak noise amplitude referenced to the input optical power. Figure 7.21(b) thus plots both the midpoint and the beginning and end points of the transition.

The difference between the optical power at the beginning and the end of the transition curves is around 62 μ W for all but the $V_{REF}=1$ V data point. This is around an order of magnitude larger than the

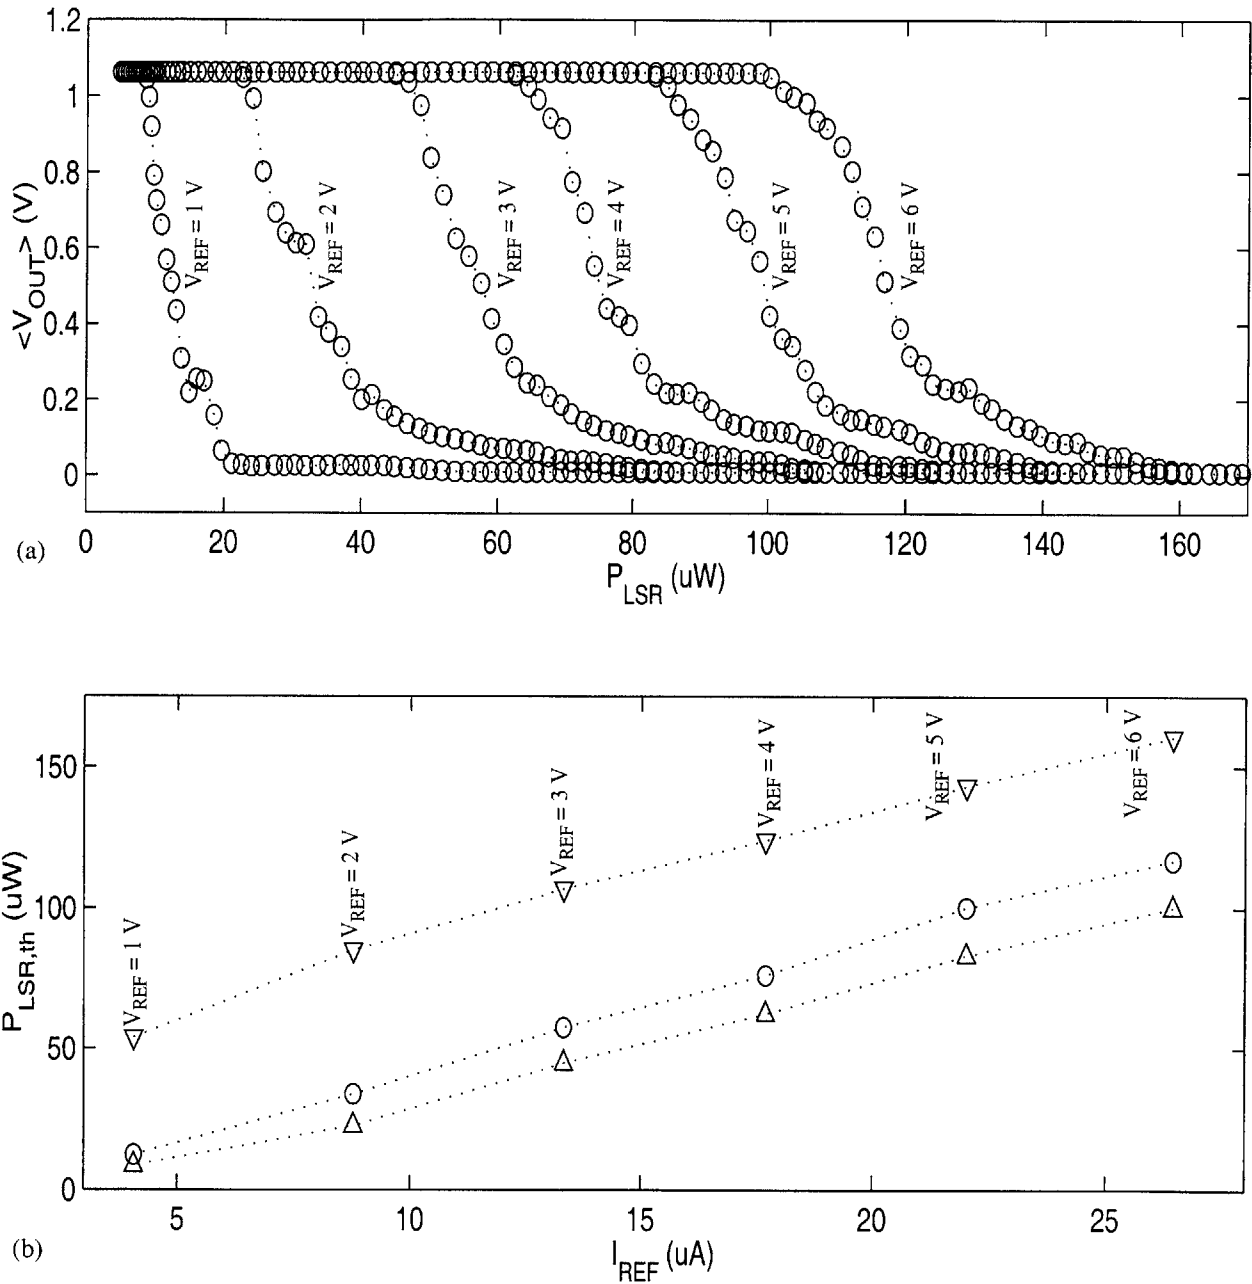


Figure 7.21: Time-average output voltage vs. input power

V_{REF} was set to the voltages indicates and the 850 nm optical input signal was swept. The ECL output signal was averaged by the HP 4145, which also set V_{REF} , measured I_{REF} and set the bias current of the VCSEL used as the optical source. (a) Shows the family of output transfer curves. (b) Summarizes the thresholding characteristics of the receiver by plotting the beginning (bottom curve), mid-point (middle curve), end (top curve) of each transition. The beginning and end points of the transition are determined as the first and last points having values which differ from the HI and LO output values at the extreme low and high power levels, respectively.

static sensitivity predicted above based on the noise in the receiver and shot noise from the laser¹. In fact, the noise in the optical signal was found to far exceed the shot noise estimate. The optical signal was measured with a Discovery Semiconductor DSC50 13 GHz photodetector. For a 100 μW optical input, and with the oscilloscope bandwidth limited to 12.4 GHz, the standard deviation in the output voltage of the detector was found to be ~ 0.1 mV after subtracting off noise due to the oscilloscope². The DSC50 was also calibrated to determine its conversion efficiency at 850 nm to be 0.24 W/V. The 0.1 mV measurement thus translates into a standard deviation of 24 μW in the optical signal. The receiver bandwidth is estimated to be approximately 1 GHz, so the 24 μW figure may be divided by $(12.4)^{0.5}$ to give the noise standard deviation in a 1 GHz band as 6.8 μW . Finally, it is observed that the peak-to-peak variation is typically around a factor of ten larger than the standard deviation. The peak-to-peak optical noise level is thus estimated to be approximately 68 μW . This agrees with the large transition width found in Figure 7.21. The noise in the optical signal will have to be reduced significantly in order to determine that actual sensitivity of the receiver³.

Figure 7.21(b) indirectly measures the efficiency of photodetector. $P_{\text{LSR}}/I_{\text{REF}}$ has a value of around 0.23 A/W. This is in agreement with the MSM27 source-current efficiency found in Chapter 4. The input bias point of the receiver, which determines the source voltage of the detector, was 0.325 V. This was measured at the open-circuited reference resistor terminal. The drain of MSM27 was set to 5 V.

Dynamic characteristics

Although the receiver itself can operate with much larger input current, the 6 V upper limit on V_{REF} sets the maximum optical power level that can be used in the present setup. For the following experiment, V_{REF} was set to 5.420 V corresponding to $I_{\text{REF}}=23.7$ μA . The optical signal amplitude was adjusted to give the best eye pattern at this reference level. The MSM27 source current resulting from this optical sig-

-
1. Note that this is a peak-to-peak value rather than a standard deviation value estimated above. The peak-to-peak variation is typically a factor of ten larger than the corresponding standard deviation.
 2. The oscilloscope signal had a ~ 0.23 mV standard deviation with no optical signal applied. This figure includes the dark noise of the DSC50 detector. The standard deviation was remeasured with 100 μW optical signal applied. The dark standard deviation value was subtracted from the latter measurement in a root-difference-square manner.
 3. Sources of noise include the ECL signal driving the laser, the 50 Ω termination resistor used in the laser module, optical reflections from the fiber facet into the VCSEL (which couple mechanical noise into the system) and mechanical noise in the attenuator (which consisted of a loosely connected FC-to-FC fiber coupler placed in the optical path).
-

nal was measured on the setup of Chapter 4 and found to have time average values of $\langle -I_S \rangle = 33 \mu\text{A}$, $\langle -I_B \rangle = 22 \mu\text{A}$, and $\langle I_D \rangle = \langle -I_S \rangle + \langle -I_B \rangle = 55 \mu\text{A}$.

Chapter 4 developed a black-box model of the detector which consisted of a current source in parallel with a impedance. Because of the non-linear and non-time-invariant behavior of the detector, the current source must be directly measured. In what follows, the MSM27 dynamics will be found to limit the performance of the receiver. Before examining these results, it is useful to consider the effect of the MSM27 terminal impedance.

In Chapter 4 (Figure 4.24) an equivalent circuit was determined for the terminal impedance of MSM27. The MSM27 step response was also measured (Figure 4.28). These elements may be directly implemented in HSPICE in order to observe their interaction with the receiver. Figure 7.22 and Figure 7.23 show transient simulations of the receiver input voltage in response to the rising and falling optical signals. The MSM27 measured source current step response is used as the input current source in the simulation. Instead of the assumed 100 fF detector capacitance used in the design phase, the equivalent MSM27 terminal impedance circuit was placed at the receiver input. To match the receiver measurements below, the MSM27 step response corresponding to $\langle -I_S \rangle = 31 \mu\text{A}$ was used. The component values used in the equivalent circuit were optimized for a measurement in which $-I_S = 15.2 \mu\text{A}$, however, the component values depend only weakly on the optical power level.

The key observation to be made from Figure 7.22 and Figure 7.23 is that the filtering effect of the parallel combination of the MSM27 terminal impedance, Z_{MSM} , and the receiver input impedance, Z_{IN} , does not limit the receiver response. That is, the bandwidth of the MSM27 photocurrent source is less than the bandwidth of $Z_{\text{MSM}} \parallel Z_{\text{IN}}$. This is apparent from the figures which shows that V_{100} closely tracks the variation in the input current. The observation that the MSM27 source current alone limits the dynamics at the receiver input allows a direct measurement of the MSM27 eye pattern to be used in assessing the effects of the detector on the receiver's performance.

Figure 7.24 shows the eye pattern of the optical signal measure with the DSC50 photodetector. The ground-signal-ground probing setup described in Chapter 4 was used to measure the photocurrent eye patterns shown in Figure 7.25 and Figure 7.26. Although the optical eye pattern is seen to be degraded by the effect of reflections in the signal path driving the laser, it is apparent that the photocurrent eye pattern is not limited by the optical signal. In viewing the MSM27 eye patterns, it should be noted that the noise band-




Figure 7.22: Transient simulation of receiver input with the MSM27 modeled rising edge

The MSM27 source current rising edge step response from Figure 4.28, with $\langle -I_S \rangle = 31 \mu\text{A}$, was used in conjunction with the equivalent circuit representation of the MSM27 terminal impedance (Figure 4.24) to simulate the interaction of the photodetector and the receiver.




Figure 7.23: Transient simulation of receiver input with the MSM27 modeled falling edge

The MSM27 source current falling edge step response from Figure 4.28, with $\langle -I_S \rangle = 31 \mu\text{A}$, was used in conjunction with the equivalent circuit representation of the MSM27 terminal impedance (Figure 4.24) to simulate the interaction of the photodetector and the receiver.

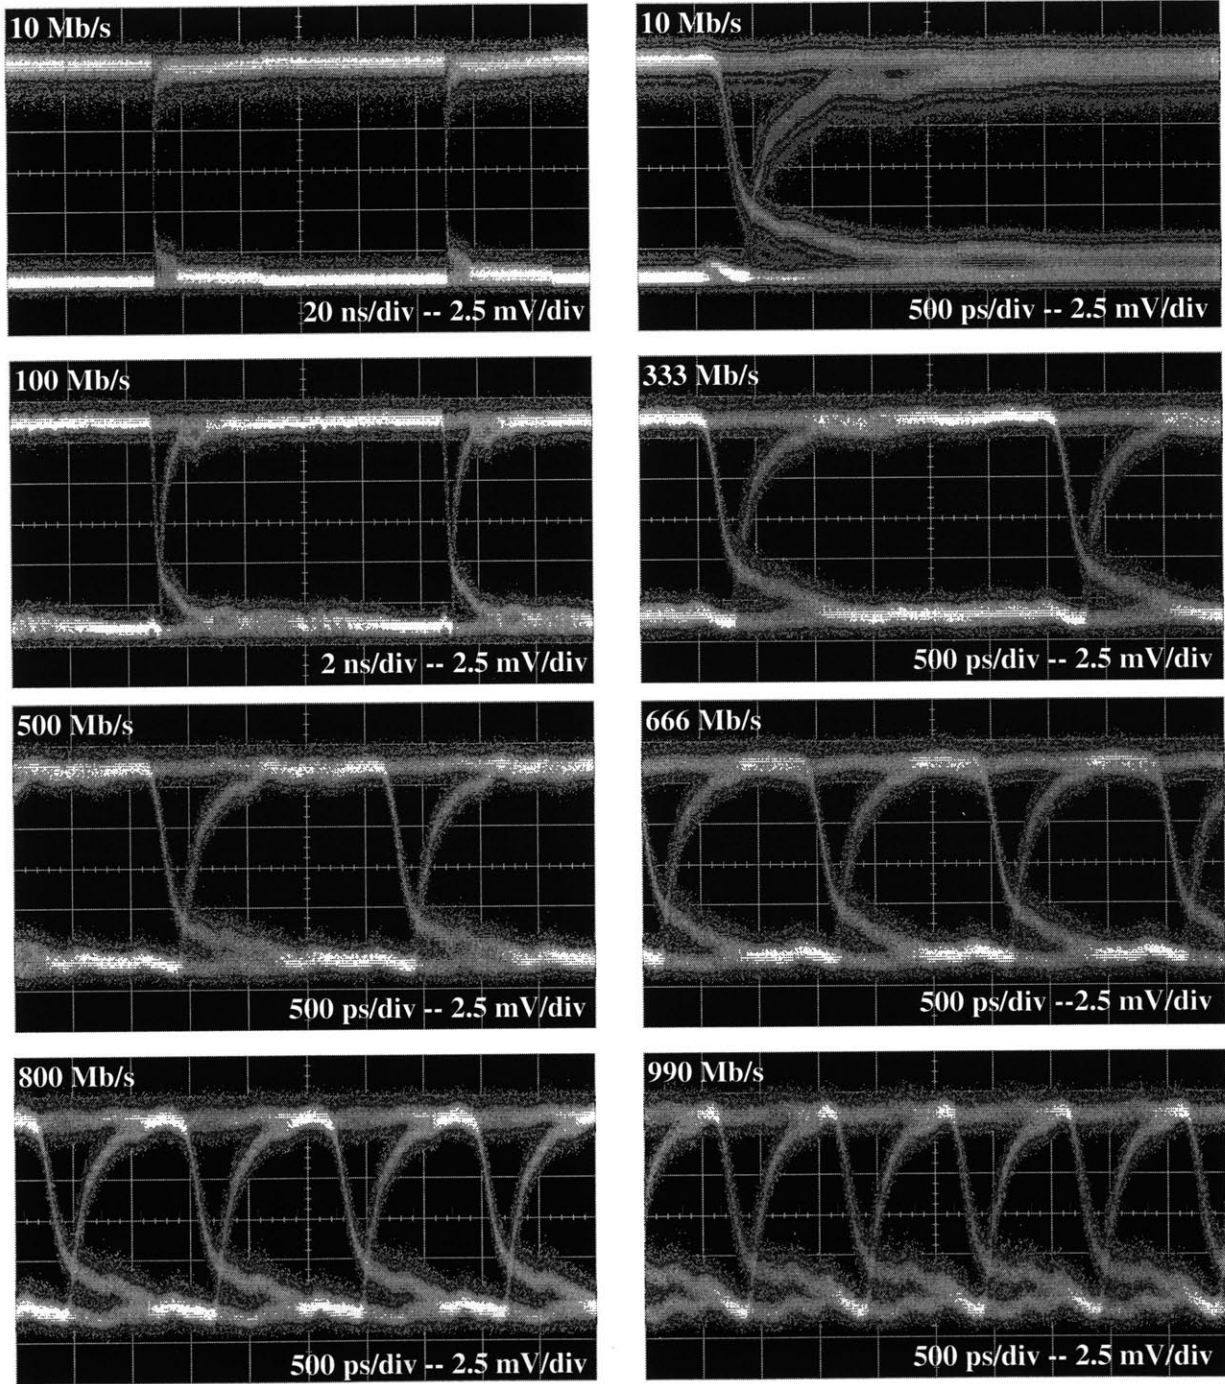


Figure 7.24: Eye pattern of optical signal used to test the receiver

The signal shown here was attenuated before butt-coupling it into the MSM27 photodetector that was used in the receiver or into a ground-signal-ground probe test structure.




Figure 7.25: MSM27 source current eye patterns -- 1 of 2

The optical signal that was used to test the receiver was also applied to an MSM27 that was connected to ground-signal-ground probe pads. This setup is described in Chapter 4. The test signal produced an average source current of $\langle -I_S \rangle = 33 \mu\text{A}$ and $\langle -I_B \rangle = 22 \mu\text{A}$. The source current eye patterns are continued in Figure 7.26.




Figure 7.26: MSM27 source current eye patterns -- 2 of 2
Continuation of Figure 7.25.




Figure 7.27: Receiver output eye patterns -- 1 of 2

Eye patterns were formed from the ECL output of the low power receiver in response to the MSM27 source photocurrent. The MSM27 response limits the performance of the receiver. Eye patterns of the MSM27 photocurrent are shown in Figure 7.25 and Figure 7.26. The receiver eye patterns are continued in Figure 7.28.

width in this measurement is 12.4 GHz as set by the oscilloscope. The receiver has a bandwidth of ~ 1 GHz so that the noise level in the receiver will be around $\sim 28\%$ of that in the eye patterns shown here.

The ECL output eye patterns from the receiver are shown in Figure 7.27 and Figure 7.28. The receiver essentially tracks the MSM27 photocurrent, which appears at its input, and recovers whatever open eye width is available in this signal.

To more closely examine the receiver response, transient simulations were carried out using the rising and falling MSM27 current waveforms of Figure 7.22 and Figure 7.23. The MSM27 terminal impedance equivalent circuit was again used at the receiver input, and a current of $23.7 \mu\text{A}$ was set at the reference input. Figure 7.29 and Figure 7.30 show the transient simulations corresponding to the rising and falling inputs, respectively. As in Figure 7.22 and Figure 7.23, the output of each linear stage is seen to track the

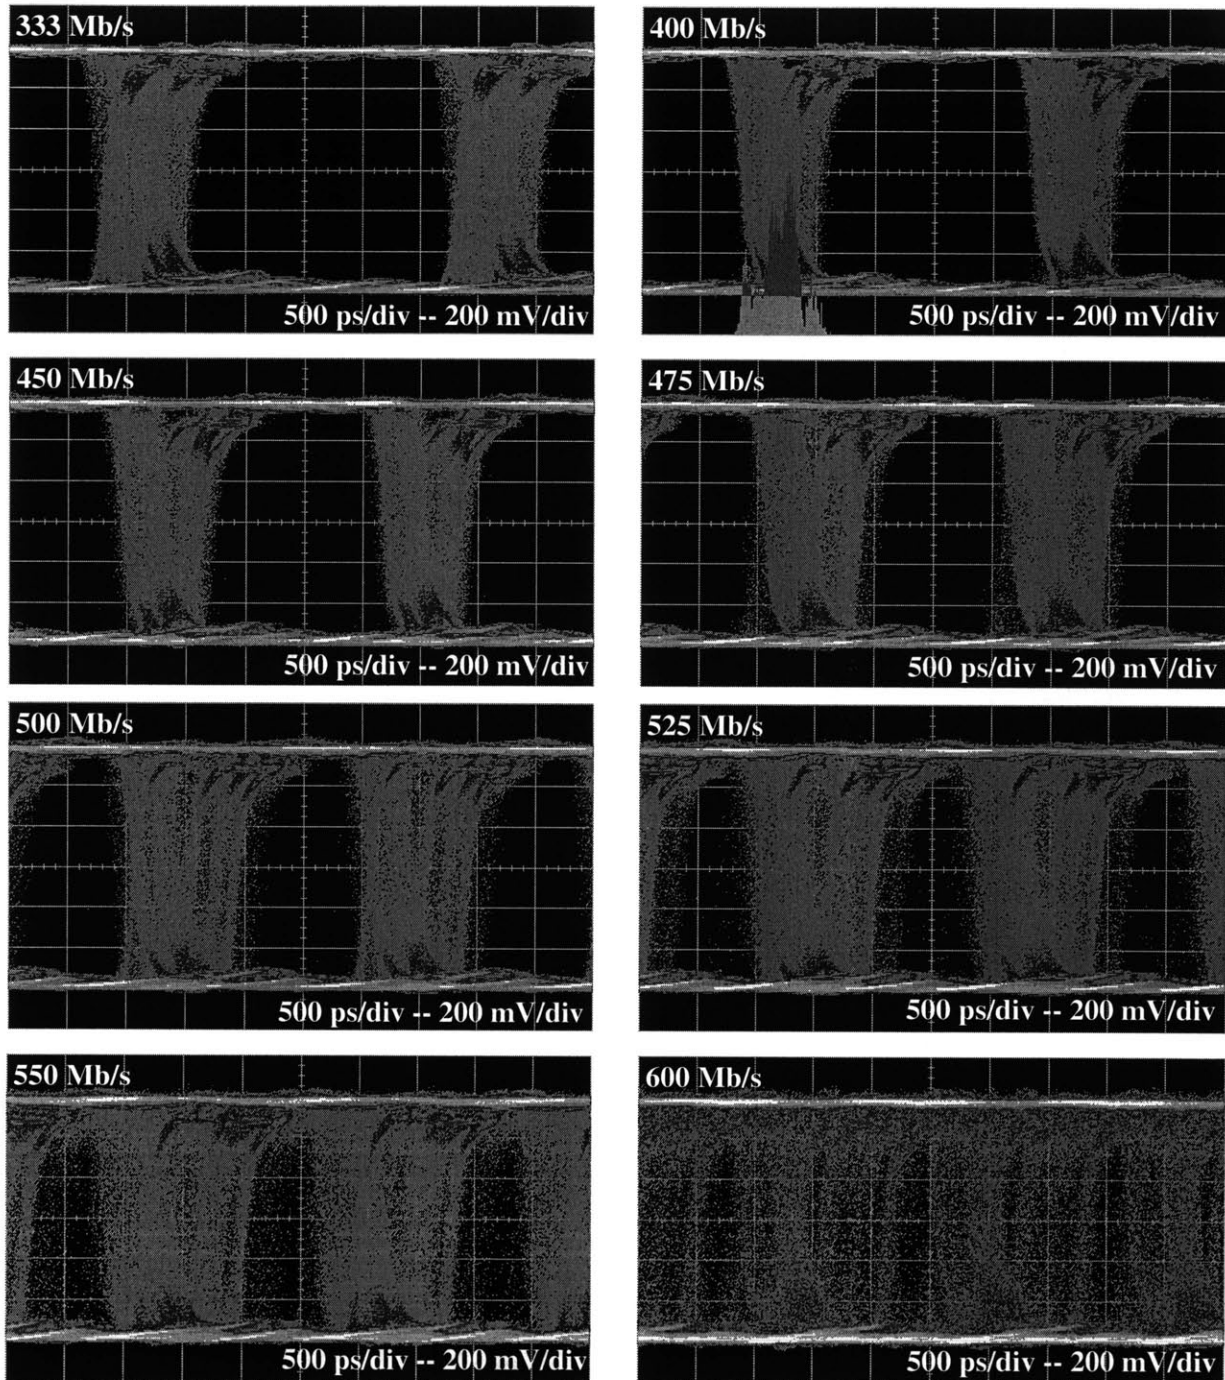


Figure 7.28: Receiver output eye patterns -- 2 of 2
Continuation of Figure 7.27.




Figure 7.29: Simulated transient response with the MSM27 modeled rising edge

The MSM27 source current rising edge step response from Figure 4.28, with $\langle -I_S \rangle = 31 \mu\text{A}$, was used in conjunction with the equivalent circuit representation of the MSM27 terminal impedance (Figure 4.24) to simulate the receiver transient response.




Figure 7.30: Simulated transient response with the MSM27 modeled falling edge

The MSM27 source current falling edge step response from Figure 4.28, with $\langle -I_S \rangle = 31 \mu\text{A}$, was used in conjunction with the equivalent circuit representation of the MSM27 terminal impedance (Figure 4.24) to simulate the receiver transient response.

input current waveforms, with only a slight degradation of the initial, rapid step and the addition of a level of drain lag consistent with the more-idealized simulation of Figure 7.12. The primary fault mechanism appears in the V_{411} waveform of the input-rising simulation shown in Figure 7.29. Because of the poor turn-on characteristics of the photodetector, V_{411} does not make a concise transition between its hard-limited low and high values. Rather, V_{411} “lingers” near the 400 mV DCFL trip-point, resulting in a large amount of output jitter. The V_{411} rising waveforms reaches its upper hard limit 6 ns after it begins to rise. This corresponds to a bit-rate of 167 MHz. For bit-rates above ~200 MHz, the V_{411} eye pattern would be expected to close significantly, and with it the output eye pattern. This is, in fact, what is observed in Figure 7.28.

7.3 Conclusion

This chapter has reviewed the development of a receiver suitable for use in a 1 Gb/s digital optical interconnect in a VLSI environment dissipating less than 5 mW. The receiver architecture balances power dissipation and area with supply noise rejection by using a combination of symmetric single-ended amplifiers followed by a differential amplifier. The elimination of level shifting between gain states reduces the power consumption by nearly one half. With perfect symmetry between the two sides of the circuit, there is essentially no degradation of the supply noise rejection as compared to a fully differential design such as that examined in Appendix H. In practical designs there will be some mismatch between the two signal paths and this may degrade the supply noise rejection to some extent.

As a digital building block, the primary performance metrics for the receiver, once correct logical functionality is established, are delay and jitter. The latter arises primarily from noise and ISI. Control of ISI was hampered by the effects of drain lag which was address by the use of three different techniques: feedback in the first (TIA) stage, cascoding in the second and third stages, and hard-limiting just prior to the final comparator. Because of hard limiting, both the bandwidth and the gain play a role in determining the propagation delay and jitter of the receiver. The need drive the signal to its hard limits also establishes a “dynamic” sensitivity which exceeds the noise-limited “static” sensitivity. Future receiver designs could particularly benefit from increased gain, decreased area, and operation from a single supply.

The receiver was characterized in conjunction with a Vitesse-fabricated, integrated photodetector studied in Chapter 4. Unfortunately, the poor performance of the photodetector did not allow the performance

limits of the receiver to be determined. The receiver was operated at up to 500 Mb/s with a peak-to-peak input signal amplitude corresponding to 60 μA of photocurrent. The performance limits expected from simulations are 1 Gb/s and 20 μA .

Because of the limits imposed by the detector, thorough measurements of the receiver were not completed. For future reference, the detailed characterization of the type of receiver developed here needs to address its digital nature. The receiver output goes through two phases following each input transition. There is an initial time interval over which the output is very uncertain, following which the output (hopefully) settles to the correct logic state. The duration of the first interval is the jitter. Whereas a communication receiver is characterized by the relationship between its input power and output bit-error-rate, a third variable, the jitter, must be added in characterizing the digital receiver. Rather than simply measuring the bit-error-rate based on optimally-timed samples of the receiver output, the bit-error-rate should be measured while varying the sampling point (delay between data and clock) over a full bit-period.¹ This series of measurements can be repeated for a range of input amplitudes and bit rates. From this data, it is possible to determine the minimum input amplitude needed to achieve a given jitter requirement and bit-error-rate (or more correctly, the error probability in the post-jitter interval). A less elaborate measurement, which can be carried out without the use of a bit-error tester, is to directly measure the jitter as a function of input amplitude and bit rate. Assuming that, for a large enough input amplitude, the steady-state (post-jitter interval) error probability is “zero”, then the increase in jitter with decreasing input amplitude and with increasing bit-rate should give great insight into receiver’s inner workings.

1. Measurements of this type have been reported in conjunction with commercial parallel optical interconnect development [43,227].

8

Conclusion

This thesis is the result of a little over six year of work beginning in the fall of 1993. Much has been achieved, and a great deal more remains to be done. Some of the accomplishments and possible directions for future work are reviewed in this chapter.

8.1 Accomplishments

A unified approach has been taken to the development of the Epitaxy-on-Electronics (EoE) optoelectronic integration technology. Initial efforts were focused on growth and fabrication, as needed to produce fully functional high performance optoelectronic integrated circuits (OEICs). This was followed by system- and circuit-level pursuits in order to make the capabilities of the resulting integration process accessible to optical interconnect researchers. This thesis is summarized by the following accomplishments.

EoE integration process

The EoE integration process produces monolithic optoelectronic-VLSI circuits by augmenting conventional electronic and optoelectronic device fabrication steps with the two EoE-specific steps of dielectric growth window (DGW) formation and polycrystalline deposit removal. By making use of commercial electronics, EoE allows the fabrication of complex OEICs without a massive up-front investment in VLSI process development. By growing epitaxial heterostructures within implanted DGWs, EoE overcomes the non-planarity limitation of traditional optoelectronic integration approaches. Furthermore, by using well-established growth and fabrication techniques, EoE lends itself to high-volume, high-yield, and high-reliability manufacturing.

In order to advance the EoE integration technology beyond the proof-of-concept established by Shenoy, *et. al.* [1-5], this thesis has revised and extended each step of the process:

- A new DGW structure and preparation procedure have been implemented. Dry-etch related substrate damage and contamination are eliminated by using the metal-1 layer of the electronics process as an etch stop to control the foundry reactive ion etch through the majority of the dielectric stack. A selective wet etch is used to complete the process.
- The use of atomic hydrogen within the molecular beam epitaxy (MBE) system has been introduced in order to remove the native GaAs surface oxide prior to growth. This is accomplished at or below 470°C whereas the previous procedure thermally adsorbed the oxide at over 580°C, resulting in damage to the electronic interconnects.
- The growth of high-quality optical emitters (lasers and LEDs) in the sub-475°C EoE-compatible temperature range has been demonstrated using the aluminum-free InGaAsP material system. Whereas the growth of AlGaAs-based emitters at the compromised temperature of 530°C previously resulted in poor emitter efficiency and some interconnect degradation, devices with aluminum-free active regions may be integrated without sacrificing the performance of either the optoelectronic or electronic performance.
- For compatibility with the new material system and to improve the robustness of the process, a procedure to remove the polycrystalline deposits which remain on the overglass following growth has been implemented which uses anisotropic dry etching rather than the wet processing.

One indication of the effectiveness of the improved process is the fact that it has recently been put into practice outside of MIT and has been used successfully with minimal involvement by the author¹.

Emitter-based VLSI-complexity OEICs

A research foundry project known as OPTOCHIP was used to demonstrate the EoE process, catalyze optical interconnect research, and identify directions for additional technology development. Aluminum-free LEDs were used in this effort in order to avoid the growth and fabrication complexities of vertical-cavity surface-emitting lasers (VCSELs). For the first time, VLSI-complexity emitter-based OEIC were fabricated. EoE's rich design capabilities were exemplified by a variety of circuits from nine research groups outside of MIT.

1. S. Choi of the University of Massachusetts, Lowell, has carried out the process at Lincoln Laboratory to integrate in-plane surface emitting lasers. An experimental phosphor source was used by A. Postigo at MIT to grow the devices, but adequate growth control was not achieved and the performance of both bulk and integrated devices was poor.

MSM photodetectors

Low-cost integrated photodetectors based on the standard, or minimally-augmented, VLSI process from Vitesse Semiconductor Corporation were investigated as an alternative to existing, commercial devices. These structures are based on the metal-semiconductor-metal (MSM) structure, but are limited by the presence of a p-n junction below the device. Their use has been demonstrated in integrated optical receivers operating at up to 500 Mb/s. An alternate device structures has been proposed which should overcome the current photodetectors' limitations.

Digital optical interconnect design analysis

A system level examination of potential EoE OEIC applications was completed and guidelines for the design of the digital optical interconnects needed for these applications were observed. The key performance parameters of the digital optical interconnect, error probability and jitter, were formulated analytically in terms of the physical parameters of its components. The analysis identifies an optimal receiver bandwidth and predicts the performance achievable by a design.

OE-VLSI laser driver design

A compact, low-power laser driver was designed for use with an index-guided VCSEL in a dense OEIC. The laser driver is based on the active control of a current mirror. The design was found to be intrinsically suitable for use in a 1 Gb/s optical interconnect, but the MESFET drain-lag effect, in conjunction with an anomalously high VCSEL input impedance, prevented this. Modifications to the circuit were described to overcome this constraint. A current-steering laser driver which uses a new buffer stage design to reduce power consumption was also introduced.

OE-VLSI receiver design

A 1 Gb/s receiver dissipating 5 mW was designed for use in a digital VLSI environment. A topology consisting of a differential amplifier preceded by symmetric pairs of single-ended stages was used to achieve low power consumption while maintaining supply noise rejection. The receiver operated at up to ~500 Mb/s, limited by the performance of the integrated photodetector. The receiver was used in a stand-alone 4x4 array OEIC for use by optical interconnect researchers. A second 4x4 array chip was implemented using a more robust, fully differential receiver design. A transconductance/transimpedance cascade is used to eliminate the need for a line driver, thereby reducing noise generation and simplifying power distribution.

8.2 Recommendations

With the goal of extending the performance of high-speed digital systems, the fields of optoelectronic integration and optical interconnect technology are rapidly evolving. This thesis has played a part in this endeavor by developing EoE into a practical, though still research-grade, optoelectronic VLSI technology. Much work remains to realize the full potential of this technology.

Manufacturing demonstration

The EoE process has been designed for compatibility with full-wafer manufacturing. As a result of budget constraints, however, the quality of the most recent EoE demonstrations has been limited by the need to handle individual die during processing. A natural next step in EoE process development would be a full-wafer implementation. With the effects of “die-and-tweezer” processing eliminated, meaningful studies of the material quality in the DGWs and of the integrated device reliability may be undertaken.

VCSEL integration

Integration of VCSELs is needed to allow the full potential of the EoE technology to be utilized. In order to accomplish this, the aluminum-free laser active regions demonstrated in this work may be combined with AlGaAs-based distributed Bragg reflectors grown at an EoE-compatible temperature. Two possible integrated VCSEL structures are shown in Chapter 5 (Figure 5.3).

VCSEL modeling and laser driver design

The two laser driver design described in Section 6.8 may serve as the starting point for additional optical transmitter development. This work should encompass a VCSEL characterization and modeling effort allowing the large signal dynamics of the VCSEL to be represented in circuit simulations.

Photodetectors

In addition to integrated VCSELs, EoE OEICs require high performance integrated photodetectors. A lateral pin detector structure implemented using the standard Vitesse process was proposed in Section 4.3 as a way to accomplish this.

Receivers

The performance limits of the current receiver designs have not been established. Either in conjunction with a suitable detector, or by means of an appropriately designed electrical test structure, the performance

of the current receiver circuits needs to be tested. Future designs may then attempt to increase bandwidth, reduce power, reduce area, reduce jitter, increase sensitivity, and improve supply noise rejection.

Digital optical interconnect design

The laser driver switching time, the high- and low-level optical power levels, and the receiver bandwidth are system-level design parameters which must be selected to meet the goals of a specific application. The analytical framework detailed in Section 5.4 provides a way to relate these parameters, along with the physical parameters which describe the receiver noise, to the error probability and jitter of the optical interconnect. Additional work may be needed to experimentally establish the validity of this model. By relating the link parameters to system-level costs (e.g. power dissipation, area, etc.) and applying additional constraints (e.g. gain available per amplifier stage, minimum signal required for hard-limiting, etc.) the model may then be used as part of an optimization procedure to establish specification for the laser driver and receiver circuits. This type of optimization procedure has been reported, but does not take any real measure of interconnect performance into account [201,202].

OEIC foundry services

The EoE technology development effort has benefited optical interconnect systems research through the OPTOCHIP project and through the receiver array OEICs. These efforts have simultaneously benefited EoE by providing a realistic context for the work. Future involvement in similar projects is like to pay even higher dividends as the effectiveness of the technology and the sophistication of the optical interconnect systems increases. Possible avenues include larger, more functional receiver OEICs, and a second OPTOCHIP offering which includes integrated VCSELs, high-speed photodetectors, and highly capable interface electronics.

The future of optical interconnects and optoelectronic integration technology

Optical interconnects play an important role in many applications. It is not known at this time how extensive their role will be in the future, and it is even less certain what the underlying technologies will look like. Ongoing work must continue to study the needs of new and existing applications, to develop pertinent solutions, and to contribute to an honest and critical assessment all of the possibilities [228].

A

Vitesse GaAs VLSI MESFET Technology

This appendix provides an introduction to the GaAs VLSI electronics process used as the starting point for the EoE-integrated OEICs detailed in this thesis. Section A.1 describes the H-GaAs III process and Section A.2 summarizes the modifications made in the H-GaAs IV process. For an introduction to MESFET circuit design, the reader is referred to [229-232], while some indication of the workings of the Vitesse HSPICE MESFET models are given in [233,234].

A.1 Vitesse H-GaAs III process¹

The Vitesse H-GaAs III E/D MESFET process produces self-aligned enhancement and depletion mode metal-semiconductor field effect transistors (MESFETs) and multiple levels of electrical interconnects. In addition, Schottky barrier diodes and metal-semiconductor-metal (MSM) photodetectors may be created using the standard gate metallization.

The thirty-eight step, thirteen mask process flow is briefly described here, following the Vitesse Foundry Design Guide [235]. The tables and figure are taken from Braun [7]. The process uses four or six inch semi-insulating GaAs substrates. Table A.1 lists the steps, which are diagrammed (up to the first layer of metallization) in Figure A.1, and described below.

Step 0: Following receipt and qualification of the semi-insulating wafers, a blanket p-type (Be) implant is performed. The non-uniform doping profile is around 10^{17} cm^{-3} near the surface extends to a depth of

1. This section is largely a reproduction of a similar appendix in the author's S.M thesis [10] which, in turn, was based on the S.M thesis of E. K. Braun [7]. Additional descriptive material may be found in [2].

Step	Process
1	Field Oxide deposition
2	Mask 1, active area definition
3	Enhancement device implant
4	Mask 2, depletion implant definition
5	Depletion device implant
6	Gate metal deposition
7	Mask 3, gate metal definition
8	LDD implant
9	Spacer oxide deposition
10	Anisotropic spacer etch
11	Source and drain implant
12	Oxide 0 deposition
13	Activation Anneal
14	Mask 4, ohmic contact metallization
15	Ohmic contact sintering
16 (Metal 1)	Oxide 1 deposition
17 (Metal 1)	Mask 5, via 1 definition
18 (Metal 1)	Metal 1 deposition
19 (Metal 1)	Mask 6, Metal 1 definition
20 (Metal 1)	Metal 1 etch
21-25 (Metal 2)	Repeat 16-20 for Metal 2
26-30 (Metal 3)	Repeat 16-20 for Metal 3
31-35 (Metal 4)	Repeat 16-20 for Metal 4
36	Passivation Dielectric deposition
37	Mask 13, passivation contact definition
38	Passivation dielectric etch

Table A.1: Vitesse HGaAs₃ Fabrication Sequence [7]




Figure A.1: HGaAs3 Process Flow Schematic through Metal 1 [7]

$\sim 0.8\text{-}1\ \mu\text{m}$ with a doping of around $10^{16}\ \text{cm}^{-3}$. This implant promotes process uniformity, and allows for junction isolation of adjacent devices. The As-rich semi-insulating substrates are typically n-type due to the presence of EL2 deep-level traps at a concentration of $\sim 10^{16}\ \text{cm}^{-3}$.

Step 1, 2: The substrate wafer is cleaned and capped with a field oxide which is patterned to define the device active areas.

Step 3, 4, 5: Silicon is ion implanted to form the MESFET channels. The dose is controlled to produce the correct threshold voltage for enhancement-mode transistors. The active areas of the depletion-mode devices are then lithographically patterned and additional silicon ions are implant to lower the threshold voltage.

Step 6, 7: The tungsten nitride refractory gate metal is deposited by reactively ion sputtering tungsten in a nitrogen ambient. W_{1-x}N_x with $x=0.1\text{-}0.2$ is produced. The gate metal is then patterned using standard lithographic techniques.

Step 8: A light implant, referred to as the LDD implant, is applied to reduce the series resistance from the source and drain to the channel.

Step 9, 10, 11: A spacer dielectric is deposited and patterned around the gates. A high dose silicon ion implant is used to form the source and drain regions. The spacer dielectric protects the channel region from this implant.

Step 12, 13: Layers of silicon nitride and silicon dioxide are deposited to cap the implant regions. The implants are then activated by annealing for one hour at 800°C .

Step 14, 15: Ohmic contacts are formed from Ni-Ge and capped with a tungsten-nitride barrier layer. The contacts are sintered at 550°C for 30 minutes.

Five steps (two mask layers) are required to form each layer of metal interconnection. The following steps are repeated for each of the four “upper level” metal layers:

Step 16, 17: Silicon dioxide, which forms the intermetal dielectric, is deposited at 380°C . Vias are opened to contact the gate and ohmic metals.

Step 18, 19, 20: The upper level metal consists of an aluminum $\text{Al}_{0.99}\text{Cu}_{0.01}$ core sandwiched between tungsten-nitride claddings (same material as the gates). The maximum temperature the chips are exposed to in this process is 200°C .

Step 36, 37, 38: The process is completed by application and patterning of the passivation dielectric layer. After spin-on-glass (SOG) is hot plate cured at 200°C, vias are etched to the underlying metal 4 layer, exposing the bond pads.

Following steps 1-38, an additional mask step, referred to as the boundary step, is performed to define saw openings in the dielectric stack to allow chips to be cut from the wafer. This step was eliminated in later versions of the process, but can be inserted by special request.

The composition and nominal thickness of each dielectric and metal layer is summarized in Table A.2. For the three metal layer process used in previous designs, the total dielectric/metallization stack thickness was around 3.7 μm . For the four metal layer process, the nominal dielectric thickness is 6.5 μm . This sets

Layer	Composition	Thickness (Angstroms)	Deposition Method	Deposition Time/Temp.
Field oxide	$\text{Si}_3\text{N}_4/\text{SiO}_2$	200/3500	PECVD	20min/380°C
Oxide 0	$\text{Si}_3\text{N}_4/\text{SiO}_2$	375/1600	PECVD	10min/380°C
Gate Metal	WN_x	3500	RIS	10min/<100°C
Ohmic Metal	$\text{Ni-Ge}/\text{WN}_x$	700/1000	RIS	5min/<100°C
Metal 1	$\text{WN}_x/\text{AlCu}_x/\text{WN}_x$	1500/8000/1000	RIS	10min/<200°C
Oxide 1	SiO_2	3000	PECVD	15min/380°C
Metal 2	$\text{WN}_x/\text{AlCu}_x/\text{WN}_x$	1000/12500/1000	RIS	15min/<200°C
Oxide 2	SiO_2	10000	PECVD	50min/380°C
Metal 3	$\text{WN}_x/\text{AlCu}_x/\text{WN}_x$	1500/17000/1000	RIS	20min/<200°C
Oxide 3	SiO_2	15000	PECVD	75min/380°C
Metal 4	$\text{WN}_x/\text{AlCu}_x$	1500/17000	RIS	20min/<200°C
Oxide 4	SiO_2	15000	PECVD	75min/380°C
Passivation	SOG	10000	SOG	<1hour/200°C

PECVD: Plasma Enhanced Chemical Vapor Deposition

RIS: Reactive Ion Sputtering

SOG: Spin on Glass

Table A.2: HGaAs3 Dielectric and Metal Layer Deposition Characteristics [7]

the total thickness of epitaxial material that must be grown in order to attain planarity after the polycrystalline material is removed.

A.2 Vitesse H-GaAs IV process

Aside from evolutionary improvements such as reduced line-widths and shallower MESFET channels, the H-GaAs IV process backgate contacts and isolation structures which are crucial to OEIC design. Details of the technology are given in the foundry design manual [208]. Diagrams of these structures, and a discussion of their importance, are given in Chapter 3 and Chapter 4.

Whereas H-GaAs III allows the blanket p^- layer, which forms a common MESFET backgate node, to float, the H-GaAs IV process provides p contacts to this layer so that its potential may be controlled. The p contact structure consists of overlapped active (mask 1), p -contact (mask 13), and gate metal (mask 3) regions. That is, the MESFET active definition step is used to open the field oxide for the p contact regions. The p -contact layer is used to cover the non- p -contact active areas while a heavy Be implant is performed, and the standard gate metallization is used to form ohmic (tunneling) contacts to the p -implanted material.

H-GaAs IV also provides isolation structures which electrically separate various regions of the p layer. The isolation regions, which take the form of rings surrounding the desired circuit area, are defined by an overlap of active (mask 1) and isolation (mask 54) layers. Since they are not covered by gate metal, the isolation regions receive the source/drain implants. They also receive a deeper n^+ implant from which the other active areas are protected by the isolation mask. The isolation implant is deep enough so that its depletion region pinches off the p layer. The isolation structure is electrically contacted using the standard ohmic contact and should be connected to the positive-most supply.

B

EoE Fabrication

B.1 Degrease Procedures

It is imperative that ICs be thoroughly cleaned prior to being introduced into the MBE environment. A complete solvent cleaning procedure should be followed at the beginning the DGW preparation process and in between every lithography step. The ICs must be inspected (under a microscope) VERY CAREFULLY to ensure that no photoresist remains following the degrease procedure.

The degrease procedure generally consists of a sequence of 15 min. boiling trichloroethylene (TCE), 15 min. acetone (ACE), and 15 min. methanol (METH) followed by a 2-propanol (PROP) rinse immediately before drying with house nitrogen (N_2). As usual, do not allow TCE or ACE to dry on the IC surface. If stubborn photoresist is observed, hot ($>60^\circ C$) or boiling¹ 1-methyl-2-pyrrolidinone (NMP) may be used. This is a good solvent for photoresist which is sometimes found in commercial photoresist stripper. Do not breath NMP fumes or allow NMP to contact skin.² NMP (hot or cold) must be removed from the IC using ACE. The recommended sequence is (15 min. each) hot-NMP, ACE, boiling-TCE, ACE³, METH, PROP, N_2 dry.

Do NOT use ultrasonic agitation in the degrease procedure or anywhere else in the EoE process. Vit-esse has indicated that ultrasonic vibration is believed to create defects in GaAs.

-
1. This is the author's preferred mode. Use caution and consult with appropriate safety personnel.
 2. NMP will may you feel dizzy and "sick".
 3. The author also heats the ACE. See note for NMP.
-

B.2 Photoresist

Relatively thick photoresist are used in the EoE process in order to ensure proper coverage over the IC surface non-planarities and the extreme DGW sidewalls.


The positive photoresist used was Shipley SC-1827 [126]. The following procedure results in $\sim 3 \mu\text{m}$ resist film (on a flat surface).

- 1) Degrease sample
- 2) Pre-bake sample on hot plate at 120°C to remove moisture. Leave on hot plate until ready to apply resist.
- 3) Remove sample from hot plate. Blow with ionized N_2 . Allow time for sample to cool. Place on spin chuck. Test spin to make sure sample does not fly off. Blow with ionized N_2 while spinning.
- 4) Apply one or more drops of Mucroposit Primer (HMDS) to surface. Let sit for 10 seconds. Spin at 3000 RPM for 40 s.
- 5) Promptly coat surface with photoresist. A glass or disposable pipette is recommended for this purpose. Spin at 3000 RPM for 40 s.
- 6) Remove photoresist build-up from sample perimeter and corners (edge-bead). A razor blade tightly covered by clean-room wipe makes a good edge-bead removal tool. Next., use a lint-free foam swab, moistened with ACE, to remove any resist which may appear on the back of the substrate. Be careful not to drop the sample.
- 7) Place sample on 120°C hot plate for 90 sec.
- 8) Repeat for each sample, then proceed to exposure.
- 9) Develop in MF-319 developer. Use plenty of developer and agitate sample continuously. Rinse in deionized water for 20 s.
- 10) Hard bake in oven at 130°C for 15 min.

The negative photoresist used in this work (for lift-off metallization steps) was Futurrex NR8-3000. Use the same procedure as above, EXCEPT: soft bake for 60 sec. (not 90 sec.), develop in RD-2 developer, and don't hard bake (since this is used for lift-off only). Again, the following procedure results in $\sim 3 \mu\text{m}$ films. NR-8 generally does not work well as an etch mask. To use as an etch mask, it must be hard baked. It still does is not very resistant to typical etches, but may be hard to remove with a solvent!

B.3 Dielectric Growth Window Preparation

The EoE process begins with the dielectric growth window (DGW) preparation procedure. The procedure is summarized in Figure 2.12. The initial structure received from Vitesse is shown below




Following a thorough degrease, sputter $0.6 \mu\text{m}$ of SiO_2 over the IC. This serves two purposes. The first is to cover microscopic openings in Vitesse's overglass layer which sometimes appear along metal-4 edges. (These may result from the extended RIE etches used at Vitesse to form the DGW structures.) Second, to provide additional strength to the overglass in order to eliminate cracking which sometimes appears following growth.

The SiO_2 must next be removed from the metal 1 regions in the DGWs. Use the dark-field WELLCLEAN mask, which is inset from the DGW border by $3 \mu\text{m}$, along with a positive resist. Etch the SiO_2 with buffered HF (BOE). *ED?*

Next, to remove the metal 1 layer, use alternating steps of hydrogen peroxide (H_2O_2) and HCl. The metal 1 layer consists of Cu-doped Al surrounded by WN_x cladding layers. The WN_x is removed by the H_2O_2 while the Al is removed by the HCl. Soak the sample in H_2O_2 for 10-15 min., rinse in DI water, then soak in HCl for 10-15 min., rinse in DI water, then soak again in H_2O_2 for 10-15 min., etc. Continue to repeat this cycle until all metal is gone and a smooth, shiny, uniformly colored surface (often the color is purple) is observed. NEVER mix the peroxide and HCl since this will quickly etch the GaAs.


Once the metal 1 layer is removed, the structure appears as shown below



Next, the $\sim 0.5 \mu\text{m}$ dielectric over the GaAs is removed using BOE. Because of Vitesse's high temperature anneals, the Si_3N_4 film immediately above the GaAs may etch very slowly. Again use the dark-field WELLCLEAN mask with the positive resist. The etch time is limited by two factors. First is lateral undercutting. The second is deterioration of the resist. Typically, the sample must be patterned and etched four times, ten minutes each, in order to complete the etch (this is using a BOE composition of ammonium fluoride 7 : HF 1). This process may proceed as follows: The first etch removes the SiO_2 above the Si_3N_4 . When the SiO_2 is removed (but before the Si_3N_4 has been etched for a long time) the surface is smooth; often it is a deep blue color. The second etch partially removes the Si_3N_4 . The etch is non-uniform and the resulting surface is splotchy and multi-colored. The third etch clears most or all of the Si_3N_4 . The surface is smooth and grey. It looks like GaAs, but this is very difficult to accurately judge since a very thin dielectric layer may remain. The fourth etch (or more correctly, the etch which follows the observation of a smooth gray surface) removes the mysterious final surface layer.

One technique to test for complete dielectric removal is to electrically probe the surface. Use two clean probe tips and contact the surface very lightly. A Schottky contact should be observed if the dielectric has been removed. If dielectric remains, a breakdown voltage of over 10 V is likely. Naturally, this test should be conducted in a non-critical DGW region since it creates surface damage. The sample should be well degreased after this test has been conducted.

Hopefully, the DGW now looks something like this:



Degrease and inspect the sample carefully before introducing it into the MBE environment.

B.4 Growth Preparation

Typically, several ICs are grown along with a piece of epi-ready GaAs wafer. The use of an In-free block is recommended rather than a Mo block. Prior to mounting, dip all of the ICs and the wafer piece in BOE for a few seconds. Rinse the pieces in DI water then soak them in clean DI water for 15 min. This step is used to form a fresh, clean surface oxide. The GaAs wafer piece should be prepared with the same surface oxide as the ICs so that it may be accurately used to judge when the oxide removal is complete. Thoroughly dry the samples with N_2 , being careful not to leave water spots on the surface. (The water must flow off the samples in a uniform sheet when the N_2 is applied.) Mount the ICs and wafer piece as required.

B.5 Polycrystalline Deposit Removal

Following growth, dismount the ICs and remove the In from the back of each sample. Sputter $\sim 0.4 \mu\text{m}$ of SiO_2 over the sample. This will serve as an etch mask for the dry etch of the polycrystalline deposits. Use the clear-field POLYSTRIP mask, which is outset from the DGW border by $5 \mu\text{m}$, and positive photoresist to pattern the SiO_2 with BOE. Complete the dry etch. This should be selective and anisotropic, so timing is not critical.

Ridges of SiO₂-covered deposits will not border each DGW (see Figure 2.21). To remove these, place the sample in propanol and rub the surface with a cotton-tipped swab. This will break off the ridges--the poly has very poor adhesion to the overglass surface.

To remove the SiO₂ that remains on the epitaxial material in the DGWs, use the dark-field WELLCLEAN mask (3 μm inset from DGW border), positive photoresist, and BOE.

Proceed to device fabrication.

B.6 LED Fabrication

This section details the integrated LED fabrication process, including interconnection of the LED top contact to the on-chip electronics. The LED structure is shown in Figure 3.1. This process may be adapted to other devices.

Ohmic contact

Use the clear-field LED-MET-1 mask with the negative resist. Develop thoroughly to ensure that no resist remains of the surface. If a *clean* oxygen plasma or ozone source is available, use it to de-scum (ash) the surface.

Just prior to loading into evaporator, dip sample in BOE for 5 sec. Rinse in DI water. Dip in GaAs etchant (25 H₂O : 1 H₂O₂ : 1 H₃PO₄) for 3 sec. Rinse in DI water. Dip in BOE for 5 sec. again. Rinse in DI water. Dry with N₂. Mount and promptly load into chamber and pump down.

Evaporate (e-beam) 30 nm Ti, 150 nm Au, 50 nm Ni.

Lift-off: Soak in ACE to remove resist. Rub surface with cotton swab if needed.

Degrease.

Anneal the ohmic contact. Rapid thermal anneal (RTA) at 420°C for 20 s was used.

Mesa etch

Use the dark-field LED-MESA mask with positive resist.

Both the GaAs contact layer and the upper InGaP layer must be etched. The etchant for the GaAs is 1 H₂O₂ : 1 H₃PO₄ : 25 H₂O. Apply this for 90 sec. to remove the GaAs contact layer. This etch does not attack InGaP. The etchant for InGaP is 1 HCl : 1 H₃PO₄ : 1 H₂O. This etch takes around 20-35 minutes to remove the 0.6 μm InGaP layer above the GaAs etch stop just above the LED core. The InGaP etch does

not attack GaAs. In high quality material, the etch results are good, producing vertical, smooth sidewalls and a smooth bottom defined by the 4 nm GaAs etch stop. The presence of defects, possibly resulting from a degree of lattice mismatch, results in rough etching results, poorly defined sidewalls, and the likely breach of the etch stop.

Passivation and Via etch

Sputter 0.3 μm SiO_2 over the sample.

To form vias to the LED ohmic contacts, use the dark-field LED-VIA-1 mask with the positive resist. Etch the SiO_2 with BOE as needed (~5 min.). Alternately, use a dry etch (see next step) for better control of the via pattern.

IC passivation cut

Openings to the IC's metal 4 layer must be completed with a dry etch. Use the dark-field LED-PASS mask with the positive resist. RIE etch down to the metal 4. The PlasmaQuest RIE was used with 45 sccm CHF_3 (95%), 2.5 sccm O_2 (5%), 50 mTorr chamber pressure, 210 W incident power (180-190 V self-bias) unheated (~30°C). This etch took around 4 hours to reach the metal 4.

Interconnect metallization

Sputter ~1 μm of Al. Use the clear-field LED-MET-5 mask with the positive resist and PAN etch (77% H_3PO_4 : 20% acetic acid : 3 % nitric acid). Heat PAN etch to 46°C. Etch for 2-2.5 min. and rinse well in DI water.

C

Circuit Characterization Setups

This appendix provides the details of the board-level circuit implementations used to characterize the modified MSM photodetectors of Chapter 4, the laser driver of Chapter 6, and the continuous time optical receiver of Chapter 7. All of these measurements, as well as the clocked optical receiver of Appendix I, are based on device and circuit test structures found on the MIT-OEIC-7 test chip. The layout of this 4.6 mm x 4.1 mm IC is shown in Figure C.1.

C.1 Setup for Low Power Laser Driver Characterization

This section describes the on-chip circuitry and board designs used to test the laser driver presented in Chapter 6.

Laser driver test block

The layout of the laser driver test block, which appears as part of the MIT-OEIC-7 test chip shown in Figure C.1, is shown in Figure C.2. This die contains the laser driver shown in Figure 6.3 and Figure 6.6, referred to here as the VCSEL driver, as well as a similar laser driver scaled to support the higher current requirements of an in-plane laser. The primary laser driver test block, located at the bottom of the die, includes on-chip clock and data generation and is detailed in this section. The definitions of the pad names indicated in Figure C.2 are given in Table C.1. This test circuit allows either an external or an internally generated clock to be used with an internal pseudorandom bit sequence (PRBS) generator. Alternately, the

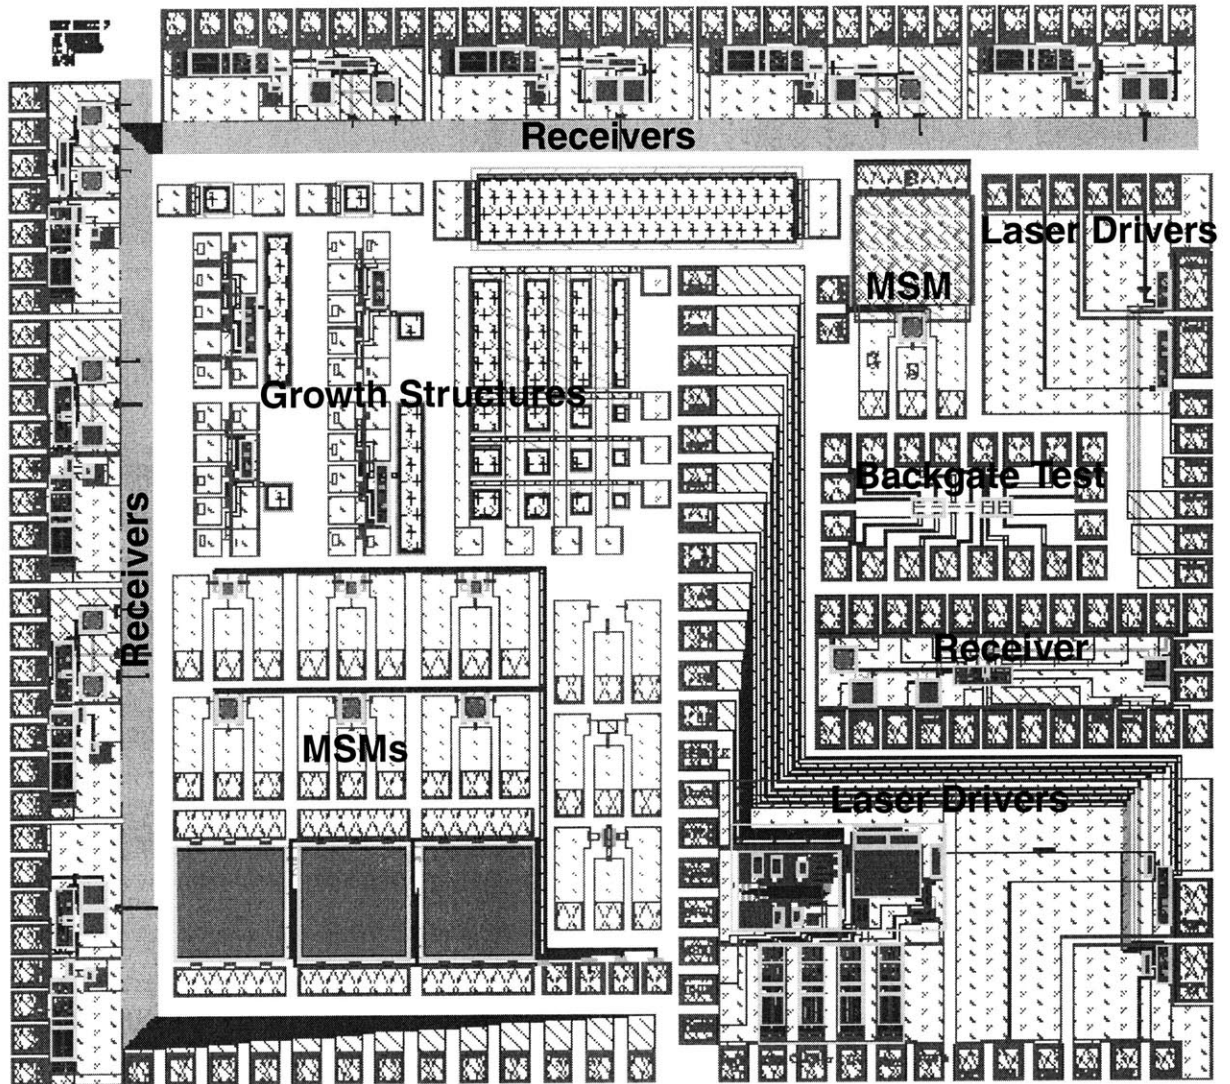


Figure C.1: Layout MIT-OEIC-7

This test chip, fabricated in Vitesse's H-GaAs IV process, contains the modified MSM photodetectors, laser driver, and receivers examined in Chapter 4, Chapter 6, Chapter 7, and Appendix I. It measures 4.6 mm x 4.1 mm.

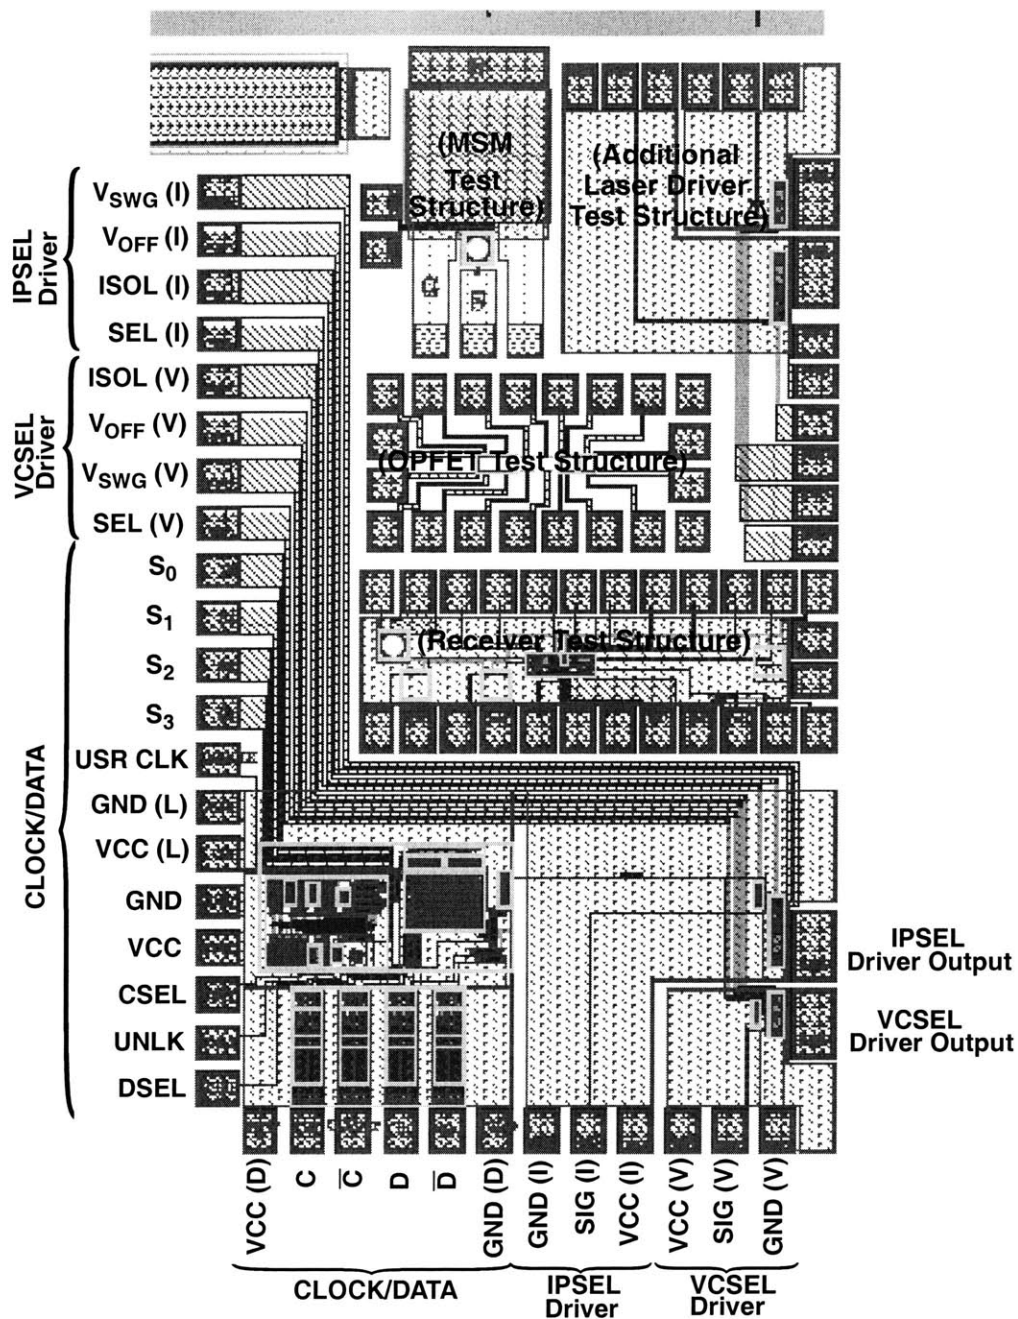


Figure C.2: Layout of die cut from MIT-OEIC-7 for laser driver characterization

The MIT-OEIC-7 test chip shown in Figure C.1 maybe cut to isolate the die shown here. This die contains the laser driver shown in Figure 6.3 and Figure 6.6, referred to here as the VCSEL driver, as well as a similar laser driver scaled to support the higher current requirements of an in-plane laser. The primary laser driver test block, located at the bottom of the die, includes on-chip clock and data generation and is detailed in this section. An additional test block at the upper right corner of the die gives direct access to the laser driver. This die also includes a number of test structures which are not related to the laser drivers.

Pad	Description
VCC (V), GND (V)	+2 V supply and ground for VCSEL driver
V _{SWG} (V)	Sets modulation reference current for VCSEL driver
V _{OFF} (V)	Sets pre-bias reference current for VCSEL driver
SEL (V)	Data select for VCSEL Driver (SEL=HI selects SIG (V) as input; SEL=LO selects internal PRBS data)
SIG (V)	External data input to VCSEL Driver
ISOL (V)	Connection to isolation structure surrounding VCSEL driver
VCC (I), GND (I)	+2 V supply and ground for IPSEL driver
V _{SWG} (I)	Sets modulation reference current for IPSEL driver
V _{OFF} (I)	Sets pre-bias reference current for IPSEL driver
SEL (I)	Data select for IPSEL Driver (SEL=HI selects SIG (I) as input; SEL=LO selects internal PRBS data)
SIG (I)	External data input to IPSEL Driver
ISOL (I)	Connection to isolation structure surrounding IPSEL driver
VCC, GND	+2 V supply and ground for clock and data generators
VCC (L), GND (L)	+2 V supply and ground for line drivers
CSEL	Select clock source (CSEL=LO selects internal clock; CSEL=HI selects USRCLK)
USRCLK	External clock input
S ₀ , S ₁ , S ₂ , S ₃	Four-bit value sets internal clock frequency; see Table C.2
UNLK	Use to force PRBS generator out of locked state (UNLK=HI flushes PRBS shift register with 0's)
DSEL	Selects polarity of PRBS data sent to laser drivers
VCC (D), GND (D)	+2 V supply and ground for ECL output drivers
C, \bar{C}	Clock outputs (ECL logic levels)
D, \bar{D}	PRBS data outputs (ECL logic levels)

Table C.1: Definition of pad names used in Figure C.2

VCSEL driver may be driven from an off-chip data source. All logic inputs to the laser test block must use the DCFL low and high levels of ~ 0 V and ~ 0.6 V.

A block diagram of the contents of the laser driver test block is provided in Figure C.3. More detailed layout views are provided in Figure C.4, Figure C.5, and Figure C.6. As shown in the block diagram, CSEL selects between the internal clock and the external user clock. The frequency of the internal clock is set by inputs S_0 - S_3 . The selected clock source is output from the test block as an ECL signal and is used internally to operate the PRBS generator. The output of the PRBS generator is also provided in the form of ECL signal as well as being delivered to the laser drivers.

A schematic of the line driver circuit indicated in Figure C.3 is shown in Figure C.7. The E/D300-302 make up a Vitesse patented “squirt buffer” [208]. The squirt buffer is preceded by a four-stage telescoping buffer.

Note that, as shown in Figure 6.9, there is a 185 ps timing asymmetry between the rising and falling edges of the data outputs. The cause of the asymmetry is not certain, but may be related to the ECL driver. Finally, as indicated in Figure C.6, either the PRBS data signal or an external signal may be selected by the user as an input to the laser driver.

Internal clock generation

The internal clock consists of five different ring oscillators of different lengths. Eleven toggle-connected D-flip-flops operate on these five sources. A total of sixteen frequencies are thus produced, and one of these is selected by a sixteen-to-one multiplexer based on the four-bit binary value applied to S_0 - S_3 . Table C.2 lists the nominal frequency values and the corresponding input settings.

The jitter in the internally generated clock was found to be too great to allow its use in characterizing the laser driver. The jitter is likely to be due to noise coupled into the ring oscillators through the supply and the substrate. Future designs may be able to reduce the jitter by 1) providing a dedicated supply for the ring oscillators (kept separate both on and off chip), and 2) operating only one ring oscillator at a time (e.g. using a NOT gate within the ring to allow the each oscillator to be disabled based on a which output has been selected).

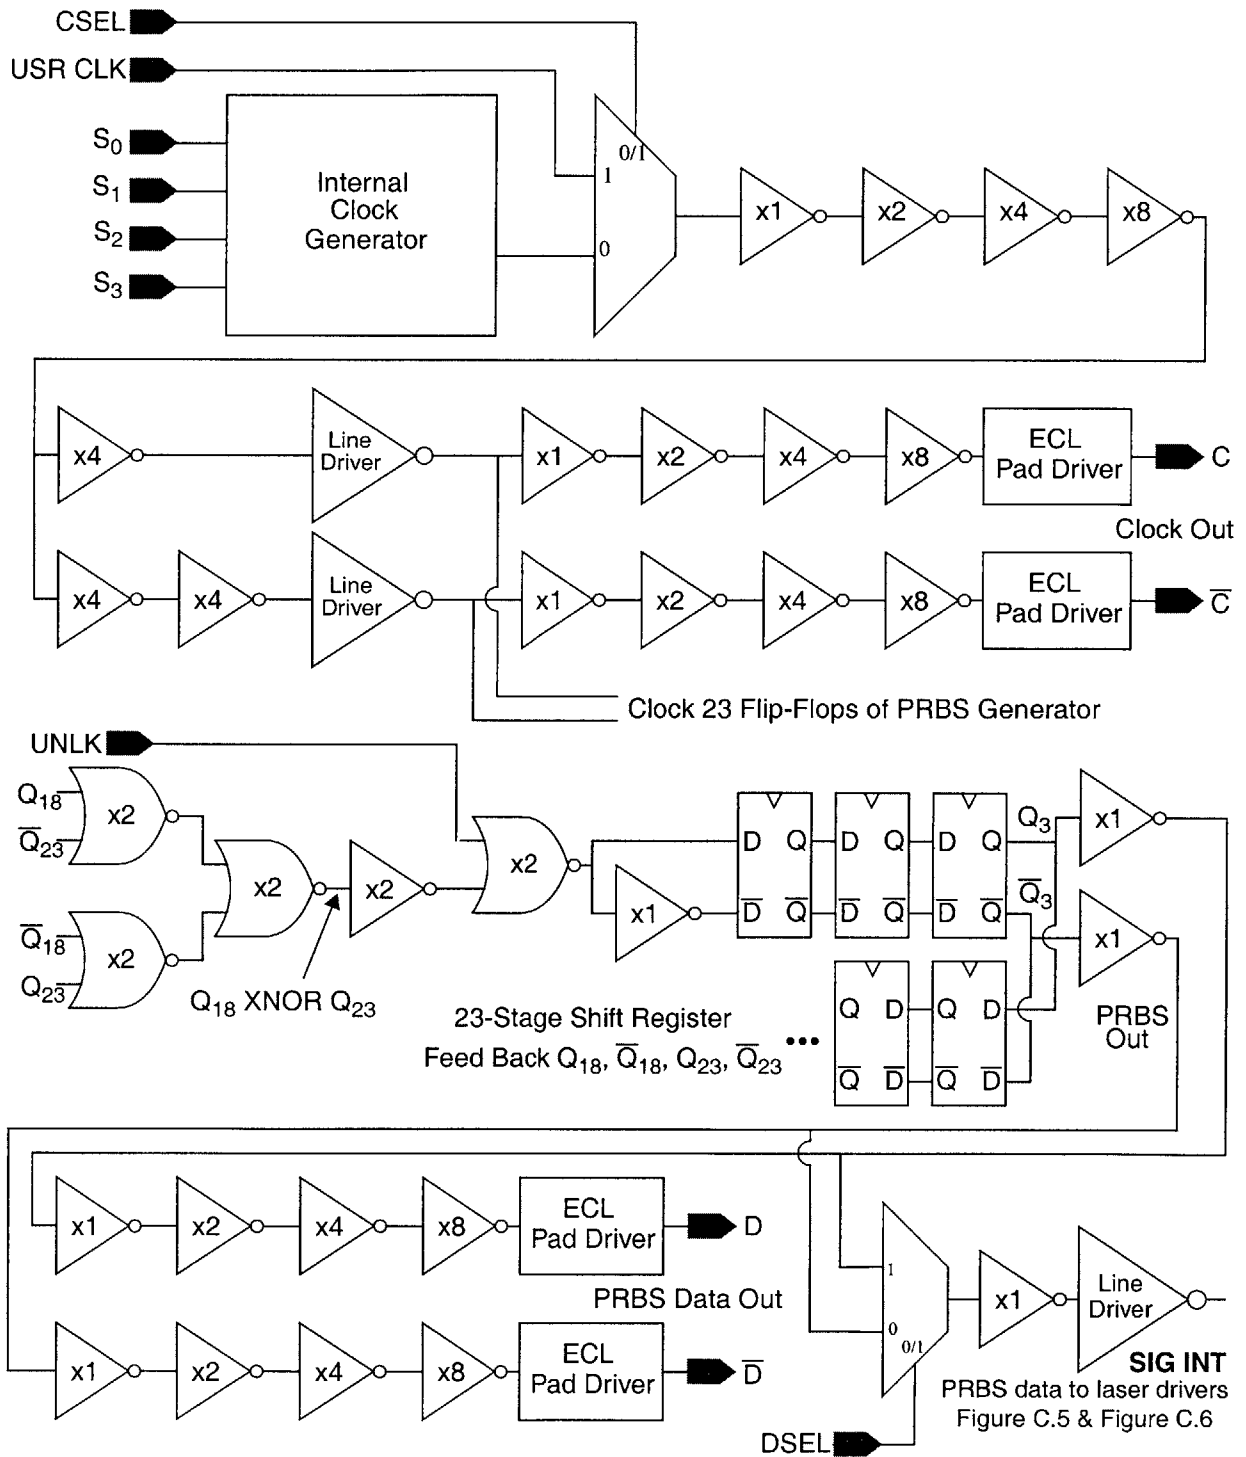


Figure C.3: Block representation of clock and data generators

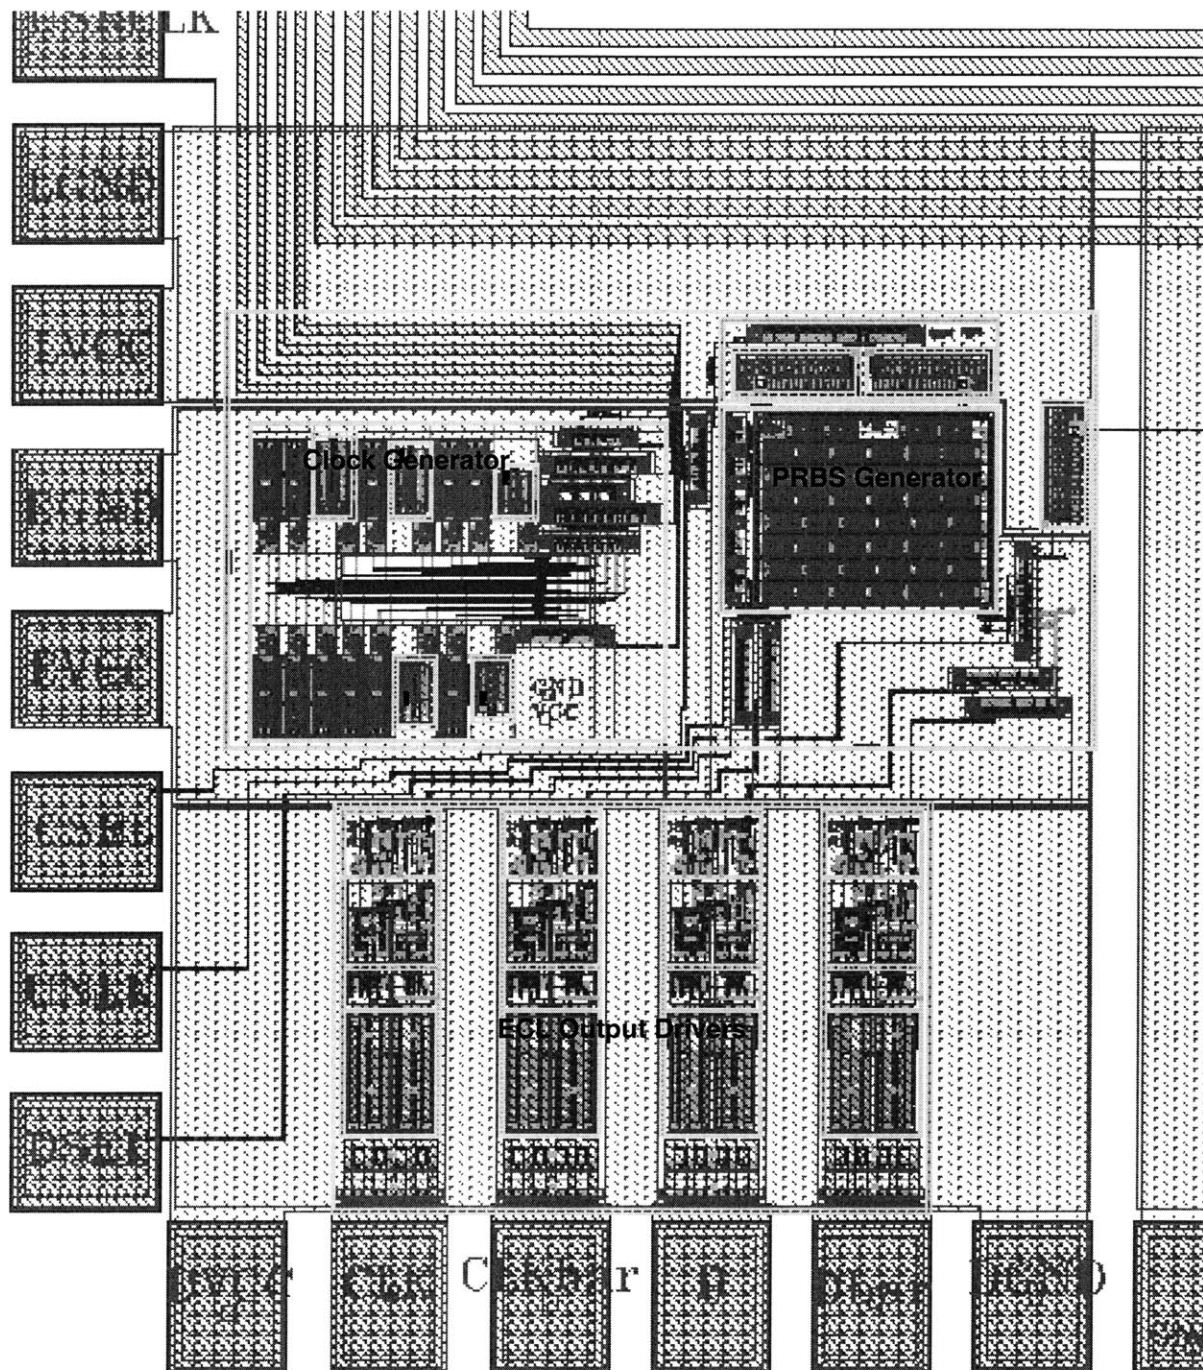


Figure C.4: Clock and data generator portion of laser driver test block

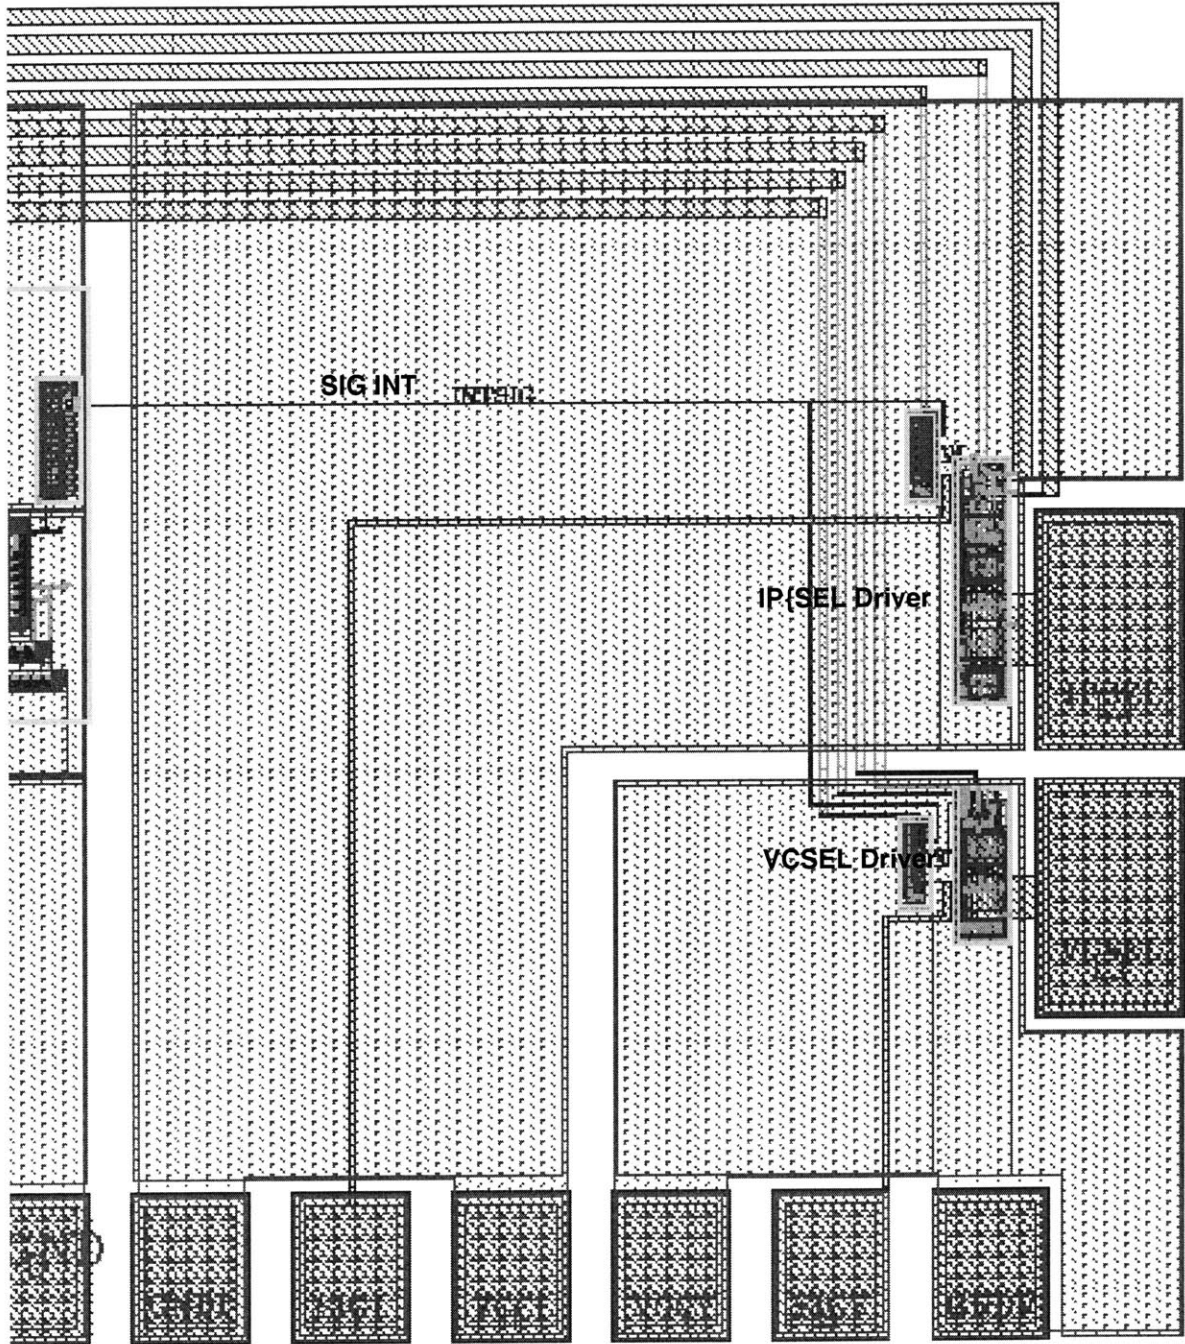


Figure C.5: Laser driver portion of laser driver test block

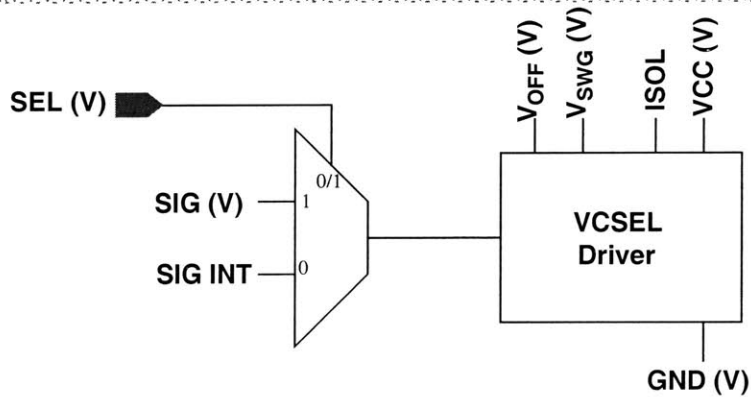
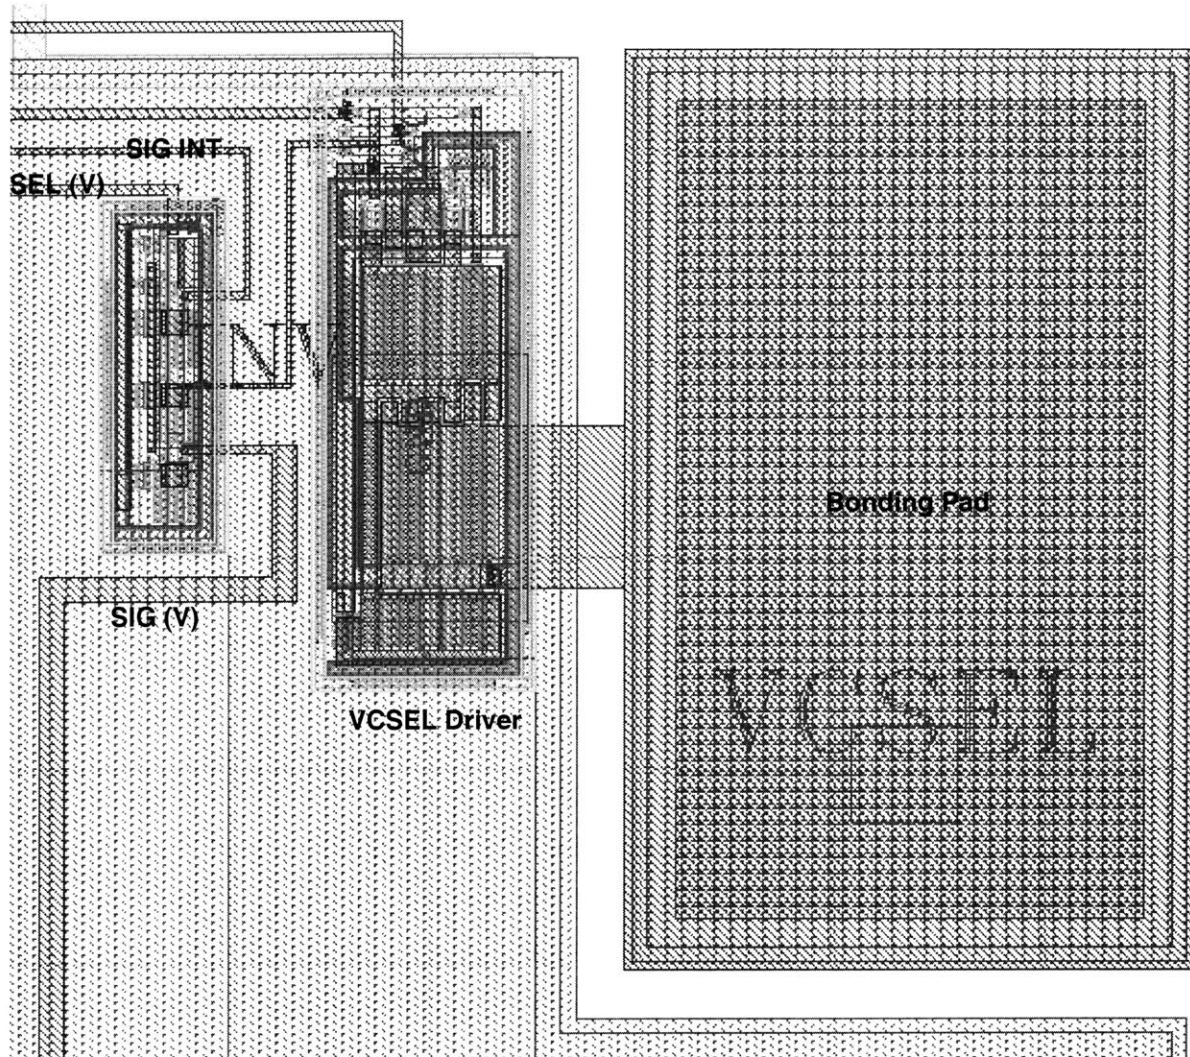


Figure C.6: Close up of laser driver test block in the vicinity of the VCSEL driver




Figure C.7: Schematic of line driver circuit

The FETs E/D300-302 make up a Vitesse-patented squirt buffer. The squirt buffer is well suited to driving large capacitive loads. It is preceded by a four-stage telescoping buffer.

Pseudorandom data generation

The PRBS generator uses a 23-bit shift register with the XNOR of the 18th and 23rd bits fed back into the first bit. The design is based on the Table C.3, which lists the appropriate taps for PRBS generators of various lengths [236]. For a given generator length, n , the bit sequence will run for $2^n - 1$ bits before repeating. Furthermore, the contents of the n flip-flops will take on all possible n -bit binary values except for one. For the XNOR feedback design, the excluded state is the all-ones state. The input UNLK is provided in order to force the generator out of this state by flushing it with zeros.

Module construction

Two circuit boards were implemented to support electrical testing the VCSEL driver as well as testing the VCSEL driver in conjunction with a VCSEL. The 6 inch square boards used standard 62 mil (1/16 inch) FR4 substrates and two layers of copper with 8 mil minimum lines and spaces. They were fabricated by Alberta Printed Circuits [237]. SMA connectors are used for all high-speed signals along with 100 mil, 50 Ω transmission line traces (with the exception of the USR CLK input with uses a narrower trace). As

S_3	S_2	S_1	S_0	Clock Frequency (MHz)
0	0	0	0	25
0	0	0	1	50
0	0	1	0	100
0	0	1	1	150
0	1	0	0	200
0	1	0	1	250
0	1	1	0	300
0	1	1	1	350
1	0	0	0	400
1	0	0	1	450
1	0	1	0	500
1	0	1	1	600
1	1	0	0	700
1	1	0	1	800
1	1	1	0	900
1	1	1	1	1000

Table C.2: Internal clock frequency settings

Inputs S_3 , S_2 , S_1 , and S_0 form a four-bit binary number which selects the desired clock frequency. The inputs use DCFL (direct-coupled FET logic) levels: Low=0 V, High=0.6 V. The clock frequencies indicated are design values. On the ICs tested, the actual frequencies were found to be roughly 20% higher.

returned from the manufacturer, the copper on the circuit boards is tin-plated. To allow wire bonding, the tin is sanded off in the vicinity of the die attachment site shortly before bonding.

The circuit board for testing the VCSEL driver along with a VCSEL is shown in Figure C.8. A bonding diagram showing how the die of Figure C.2 and a VCSEL chip are to be packages onto this board is given in Figure C.9. The VCSEL is fabricated on a p-type substrate and is supplied through the copper region labeled V_{LSR} . Figure 6.12 diagrams the signal path implemented by this configuration.

The circuit board for electrical testing of the VCSEL driver is shown in Figure C.10 and a bonding diagram may be found in Figure C.11. Figure 6.10 diagrams the signal path implemented by this configuration.

n	XNOR from	n	XNOR from	n	XNOR from	n	XNOR from
3	3,2	45	45,44,42,41	87	87,74	129	129,124
4	4,3	46	46,45,26,25	88	88,87,17,16	130	130,127
5	5,3	47	47,42	89	89,51	131	131,130,84,83
6	6,5	48	48,47,21,20	90	90,89,72,71	132	132,103
7	7,6	49	49,40	91	91,90,8,7	133	133,132,82,81
8	8,6,5,4	50	50,49,24,23	92	92,91,80,79	134	134,77
9	9,5	51	51,50,36,35	93	93,91	135	135,124
10	10,7	52	52,49	94	94,73	136	136,135,11,10
11	11,9	53	53,52,38,37	95	95,84	137	137,116
12	12,6,4,1	54	54,53,18,17	96	96,94,49,47	138	138,137,131,130
13	13,4,3,1	55	55,31	97	97,91	139	139,136,134,131
14	14,5,3,1	56	56,55,35,34	98	98,87	140	140,111
15	15,14	57	57,50	99	99,97,54,52	141	141,140,110,109
16	16,15,13,4	58	58,39	100	100,63	142	142,121
17	17,14	59	59,58,38,37	101	101,100,95,94	143	143,142,123,122
18	18,11	60	60,59	102	102,101,36,35	144	144,143,75,74
19	19,6,2,1	61	61,60,46,45	103	103,94	145	145,93
20	20,17	62	62,61,6,5	104	104,103,94,93	146	146,145,87,86
21	21,19	63	63,62	105	105,89	147	147,146,110,109
22	22,21	64	64,63,61,60	106	106,91	148	148,121
23	23,18	65	65,47	107	107,105,44,42	149	149,148,40,39
24	24,23,22,17	66	66,65,57,56	108	108,77	150	150,97
25	25,22	67	67,66,58,57	109	109,108,103,102	151	151,148
26	26,6,2,1	68	68,59	110	110,109,98,97	152	152,151,87,86
27	27,5,2,1	69	69,67,42,40	111	111,101	153	153,152
28	28,25	70	70,69,55,54	112	112,110,69,67	154	154,152,27,25
29	29,27	71	71,65	113	113,104	155	155,154,124,123
30	30,6,4,1	72	72,66,25,19	114	114,113,33,32	156	156,155,41,40
31	31,28	73	73,48	115	115,114,101,100	157	157,156,131,130
32	32,22,2,1	74	74,73,59,58	116	116,115,46,45	158	158,157,132,131
33	33,20	75	75,74,65,64	117	117,115,99,97	159	159,128
34	34,27,2,1	76	76,75,41,40	118	118,85	160	160,159,142,141
35	35,33	77	77,76,47,46	119	119,111	161	161,143
36	36,25	78	78,77,59,58	120	120,113,9,2	162	162,161,75,74
37	37,5,4,3,2,1	79	79,70	121	121,103	163	163,162,104,103
38	38,6,5,1	80	80,79,43,42	122	122,121,63,62	164	164,163,151,150
39	39,35	81	81,77	123	123,121	165	165,164,135,134
40	40,38,21,19	82	82,79,47,44	124	124,87	166	166,165,128,127
41	41,38	83	83,82,38,37	125	125,124,18,17	167	167,161
42	42,41,20,19	84	84,71	126	126,125,90,89	168	168,166,153,151
43	43,42,38,37	85	85,84,58,57	127	127,126		
44	44,43,18,17	86	86,85,74,73	128	128,126,101,99		

Table C.3: Taps for maximum-length PRBS generators

A 2^n-1 bit pseudorandom bit sequence (PRBS) generator may be constructed from a n -bit shift register by forming the exclusive-nor (XNOR) of the indicated register bits (bits are numbered starting with 1) and feeding this value back into the first bit of the shift register. This table was taken from [236]. The all-1 state locks up this state machine.

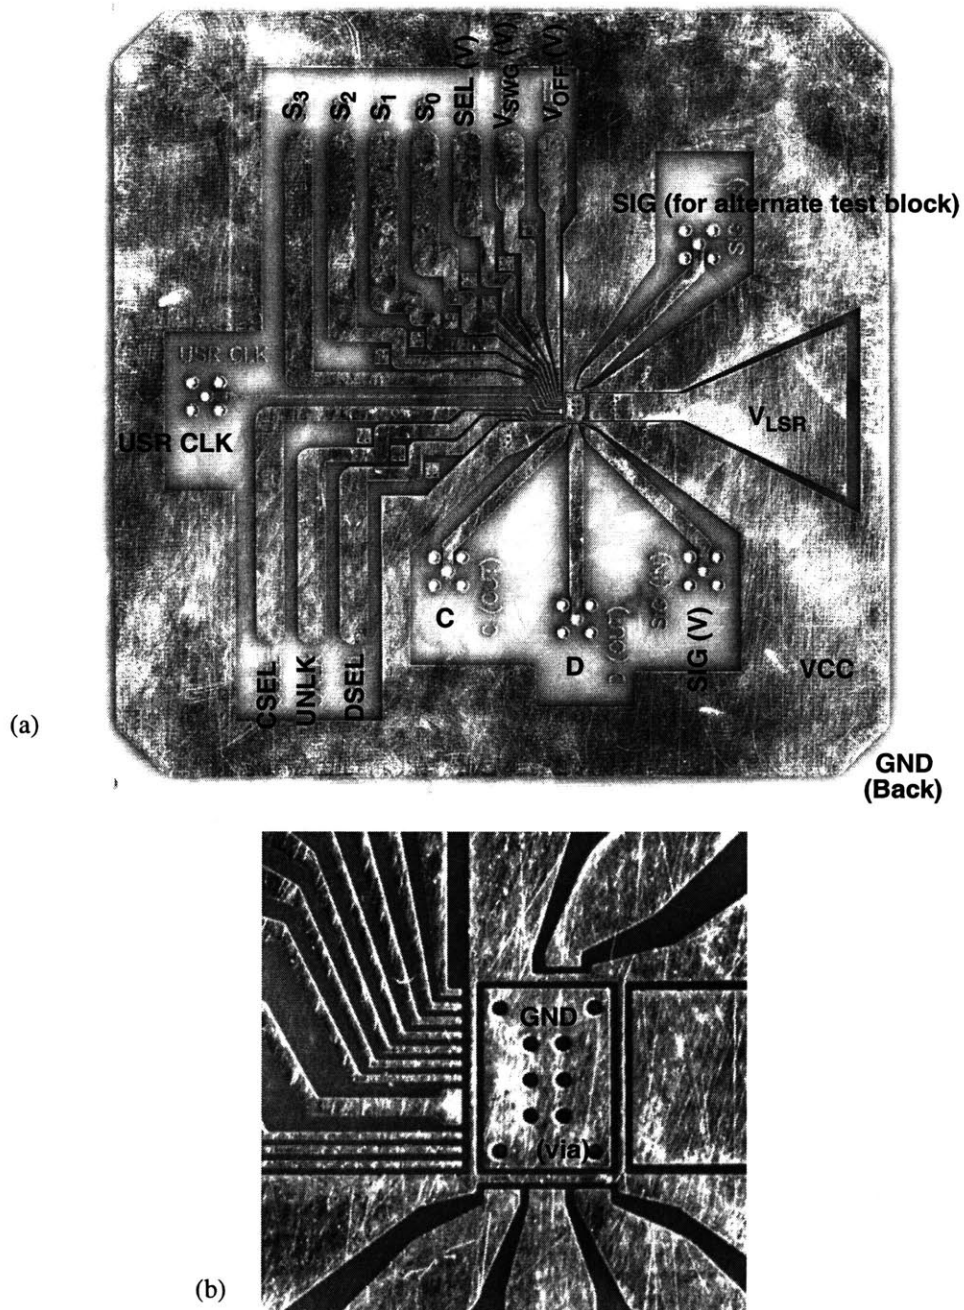


Figure C.8: Test board for laser driver + VCSEL characterization

(a) is the six inch square board (top side) and (b) is a close up of the site for attaching the laser test block IC. The back is a dedicated ground plane.




Figure C.9: Bonding diagram for laser driver + VCSEL characterization

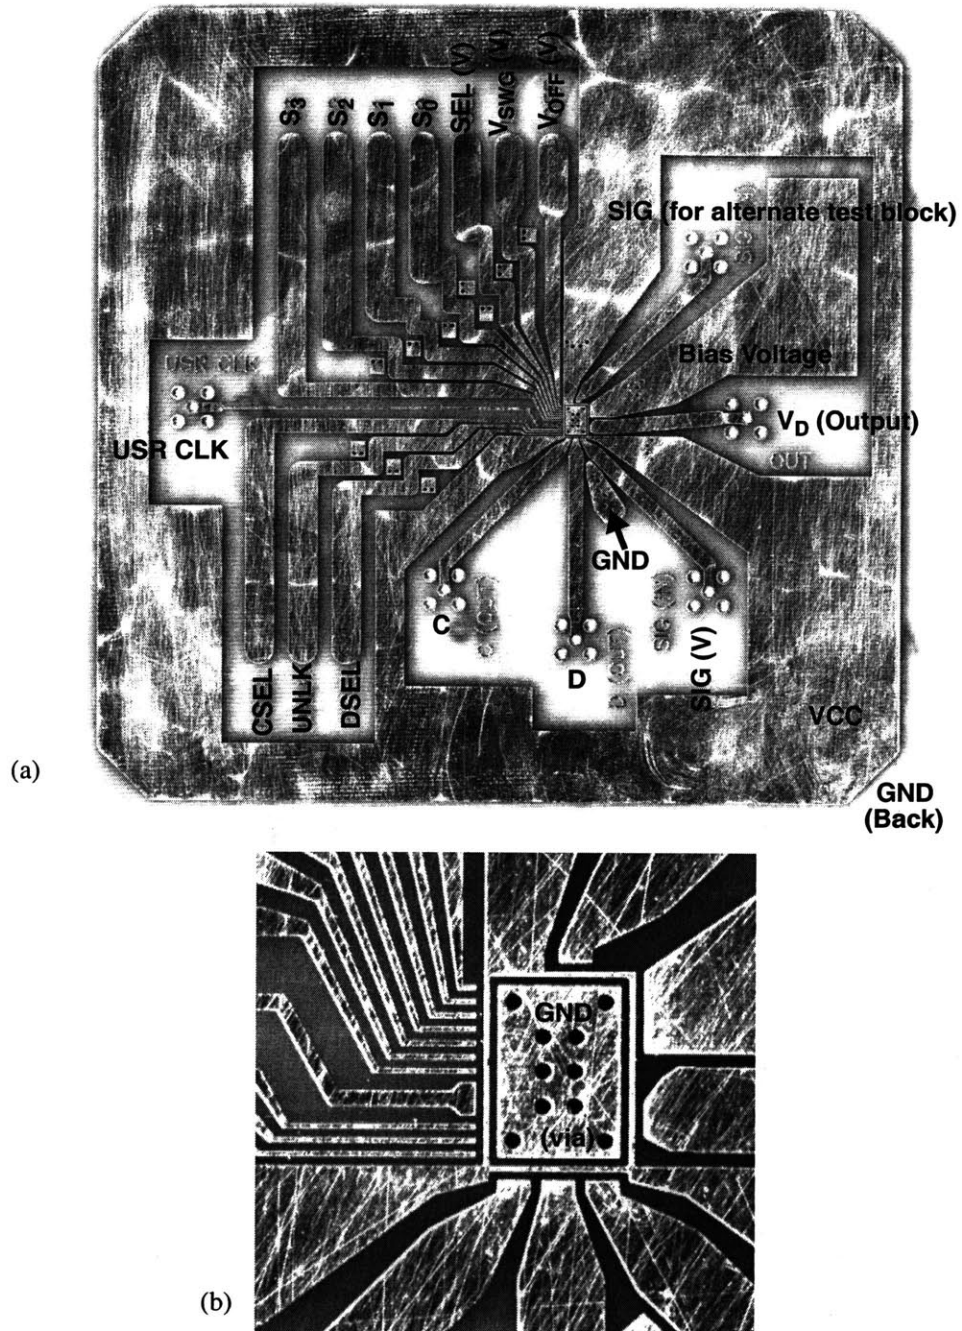


Figure C.10: Test board for laser driver electrical characterization

(a) is the six inch square board (top side) and (b) is a close up of the site for attaching the laser test block IC. The back is a dedicated ground plane.




Figure C.11: Bonding diagram for laser driver electrical characterization

C.2 ECL-Driven Laser Module

A module was constructed to allow a VCSEL to be modulated by an ECL signal. The VCSEL used in this module is of the same type detailed in Section 6.3. The circuit board, of the same construction as those described above, was intended for use in the characterization of the clocked optical receiver described in Appendix I and was retrofitted for use in the laser module. The board is shown in Figure C.12. The electrical schematic of the module, and the physical implementation using the board in Figure C.12 are shown in Figure C.13.

In this module, the ECL signal is AC-coupled onto the board through the external bias tee. An in-line BNC $50\ \Omega$ terminator placed on the DC port of the bias tee provides the DC current path for the ECL output. The $50\ \Omega$ AC termination is provided by the $50\ \Omega$ surface mount resistor placed at the end of the on-board transmission line immediately prior to the laser. This resistor is AC-grounded through bypass capacitors on the back side of the board. An external current source is used to provide the VCSEL bias current. The performance characteristics of this module are shown in Section 4.2.

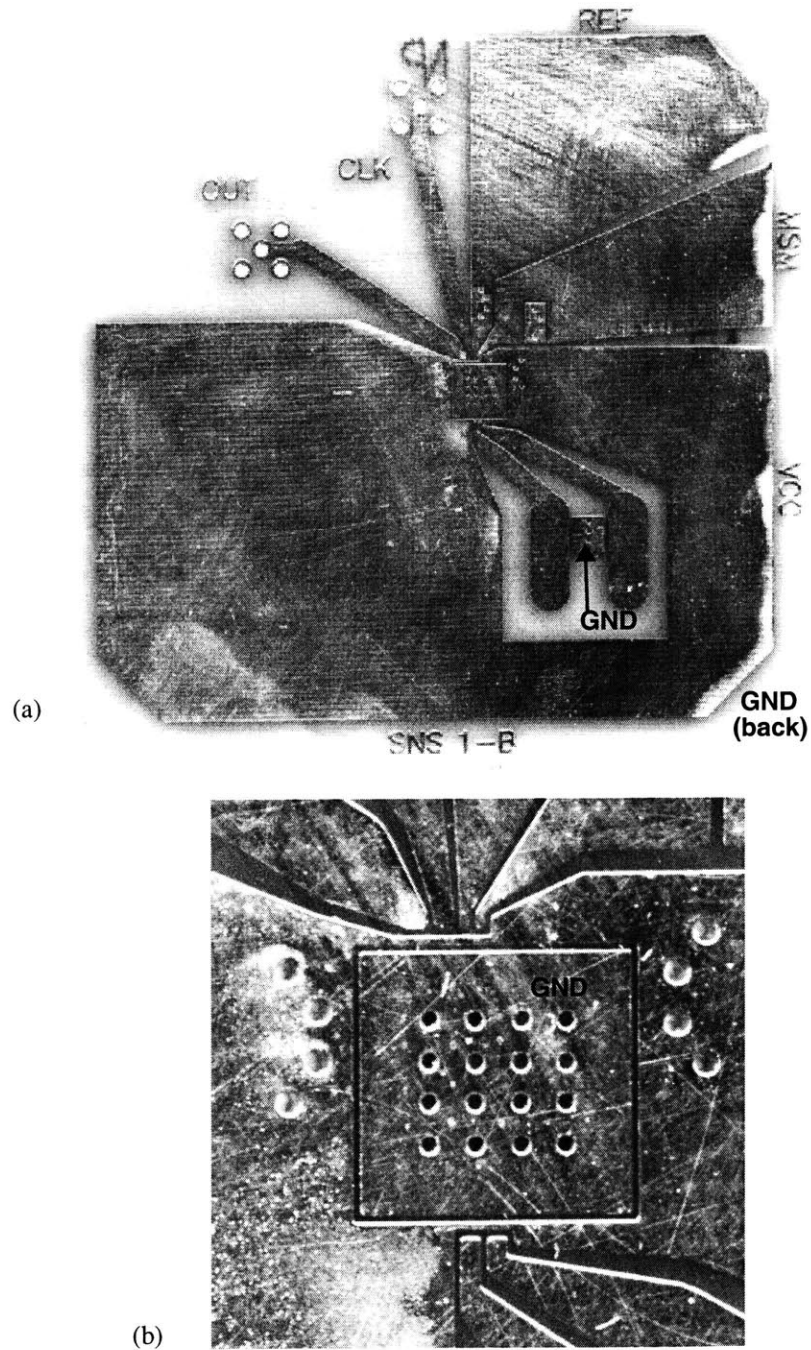


Figure C.12: Laser module test board

(a) is the four inch square board (top side) and (b) is a close up of the site for attaching the VCSEL die. This board was designed for testing the clocked receiver of Appendix I as was retrofitted for use in the laser module.


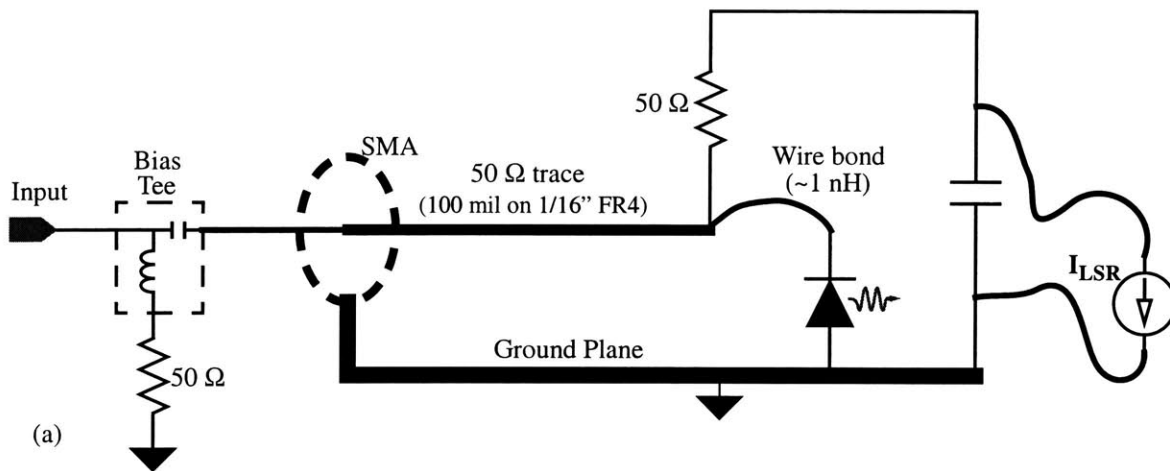


Figure C.13: Laser module construction and schematic

(a) is the schematic of the laser module circuit and (b) shows the physical construction.

C.3 Setup for Receiver Characterization

A two-layer circuit board, similar in construction to those used for the laser driver measurements, was used to characterize the low power optical receiver of Chapter 7. This section describes the MIT-OEIC-7 test cell and the circuit board used to package it.

The MSM27 photodetector detailed in Chapter 4 was coupled with the low power receiver in the test cell shown in Figure C.14. The contents of this test cell is described in Figure C.15. A 200 k Ω source/drain implant resistor is connected to the reference input of the receiver so that the REF input voltage may be used to set the reference current. The p-contact beneath the MSM27 device is routed to the PCNT pad along the bottom of the die. The isolation rings surrounding the MSM27, the reference resistor, and the receiver are routed to the ISOL pad next to the PCNT pad. The receiver output buffered and applied to both a Vitesse ECL output driver and a small pad which may be probed using a high-speed active probe.

The four-inch square circuit board that was used is shown in Figure C.16 and a bonding diagram showing how MIT-OEIC-7 is mounted onto this board is shown in Figure C.17.

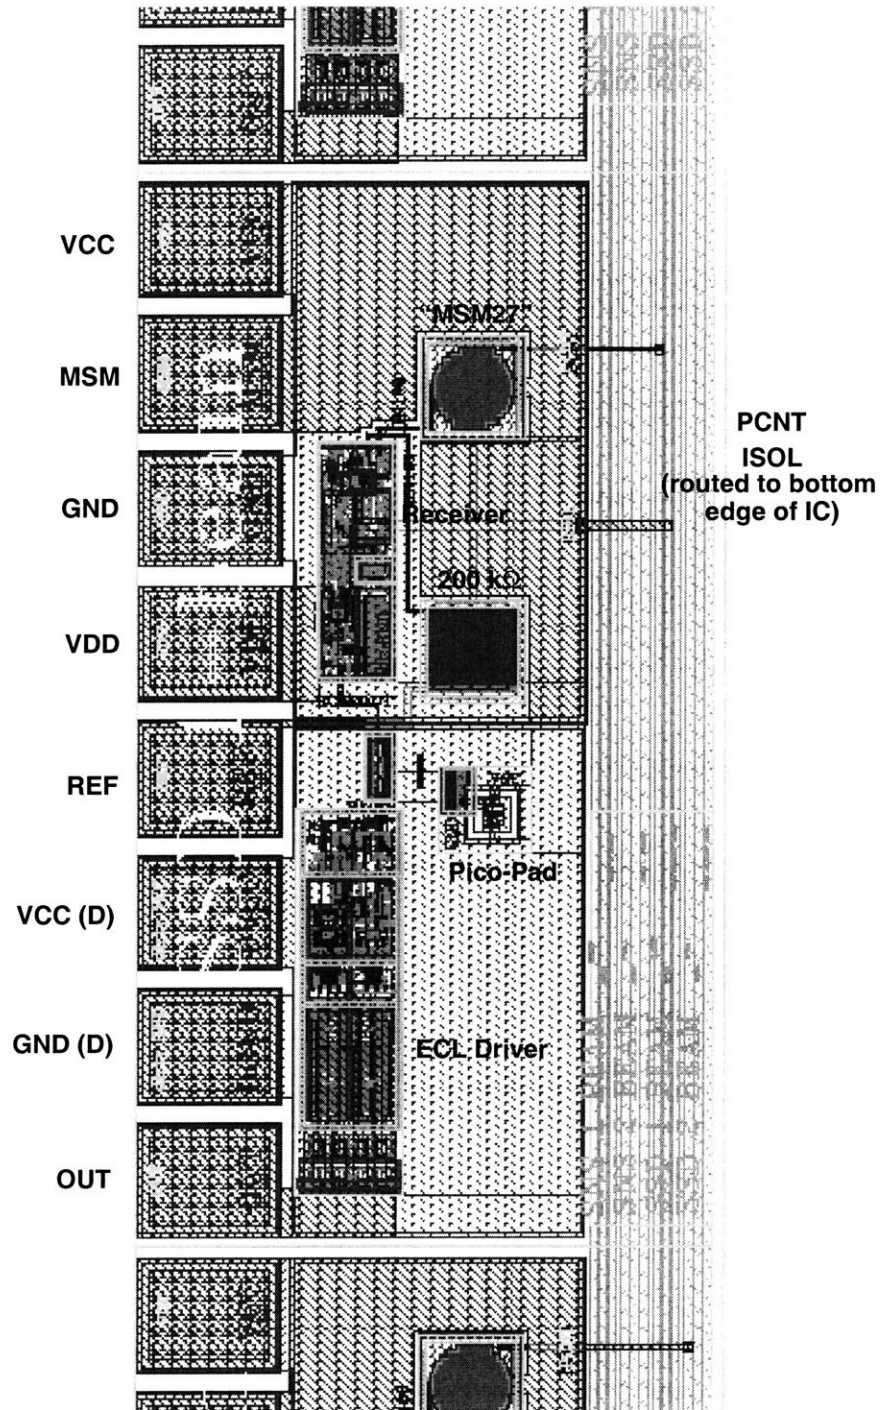


Figure C.14: Layout of low power optical receiver test structure on MIT-OEIC-7

Test structure used to characterize the low power optical receiver. This cell is located along the left edge of the MIT-OEIC-7 die.

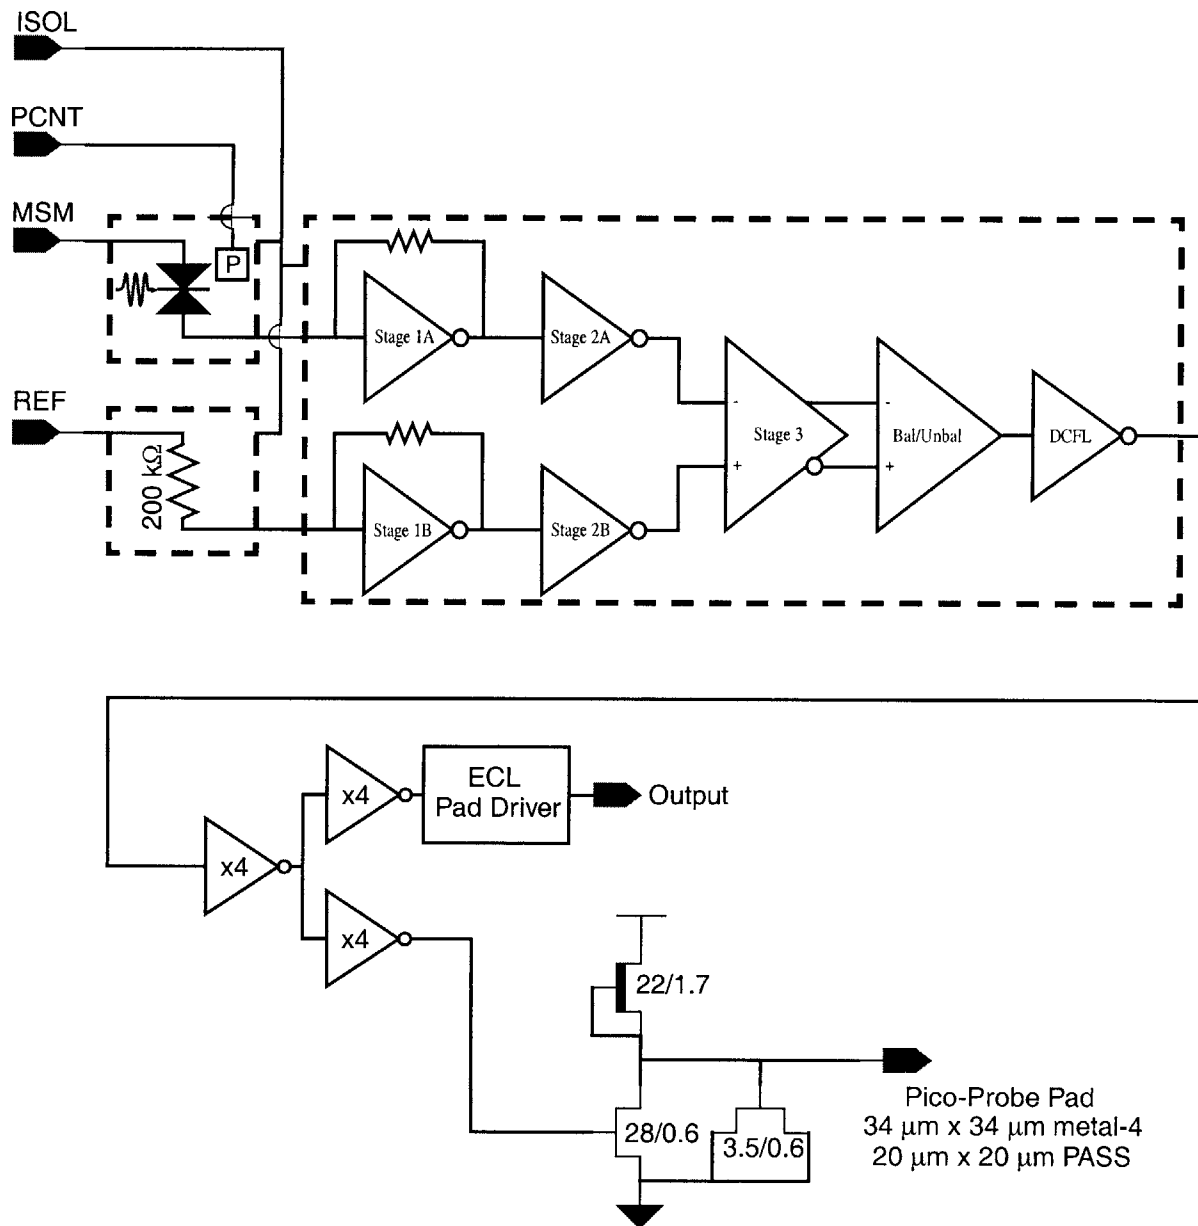


Figure C.15: Block diagram of low power optical receiver test structure on MIT-OEIC-7

The MSM27 photodetector is coupled with the low power optical receiver in this test cell. A 200 k Ω source/drain implant resistor is used to set the reference current. The receiver output is buffered and applied to an ECL output driver as well as a large DCFL inverter which drives a small pad suitable for use with an active probe ("Pico-Probe"). The 3.5/0.6 EFET connected to this pad is used to clamp the output to a valid, high speed, drain-lag-free DCFL logic signal.





Figure C.16: Test board for low power receiver characterization

(a) is the four inch square board (top side) and (b) is a close up of the site for attaching the MIT-OEIC-7 die which contains the receiver test circuit. The back is a dedicated ground plane.




Figure C.17: Bonding diagram for low power receiver characterization

D

Jitter

This appendix describes some of the mechanisms by which jitter can appear in a digital optical link, and derives simple formulae for the magnitude of this jitter. The basic sources of jitter which are considered are intersymbol interference and additive noise.

D.1 Intersymbol Interference, Eye Closure, and Jitter

A signal in either a communication channel or a digital interconnect may be modeled as a sum of delayed pulses modulated by the data stream

$$s(t) = \sum_{k = \dots, -1, 0, 1, \dots} A_k p(t - kT) \quad (\text{D.1})$$

where A_k are the data symbols (bits in the digital interconnect) and $p(t)$ is pulse shape. Figure D.1 shows examples of pulse shapes and data signals which may be encountered in communication systems and digital interconnects. The data signals are shown as eye diagrams formed by overlaying waveforms corresponding to different, random bit sequences. The bandwidth required for $p(t)$ in a communication system can be as small as half the bit rate, while a digital link typically requires a bandwidth of three times the bit rate.

In the case of a communication system, the data is recovered by sampling $s(t)$ at bit-period intervals, $t=kT$. In the receiver, the sample timing is based on a clock which is recovered from $s(t)$ itself. The pulse, $p(t)$ is required to be finite at $t=0$ and to equal zero at the other sample points. In this way, the sample $s(kT)$

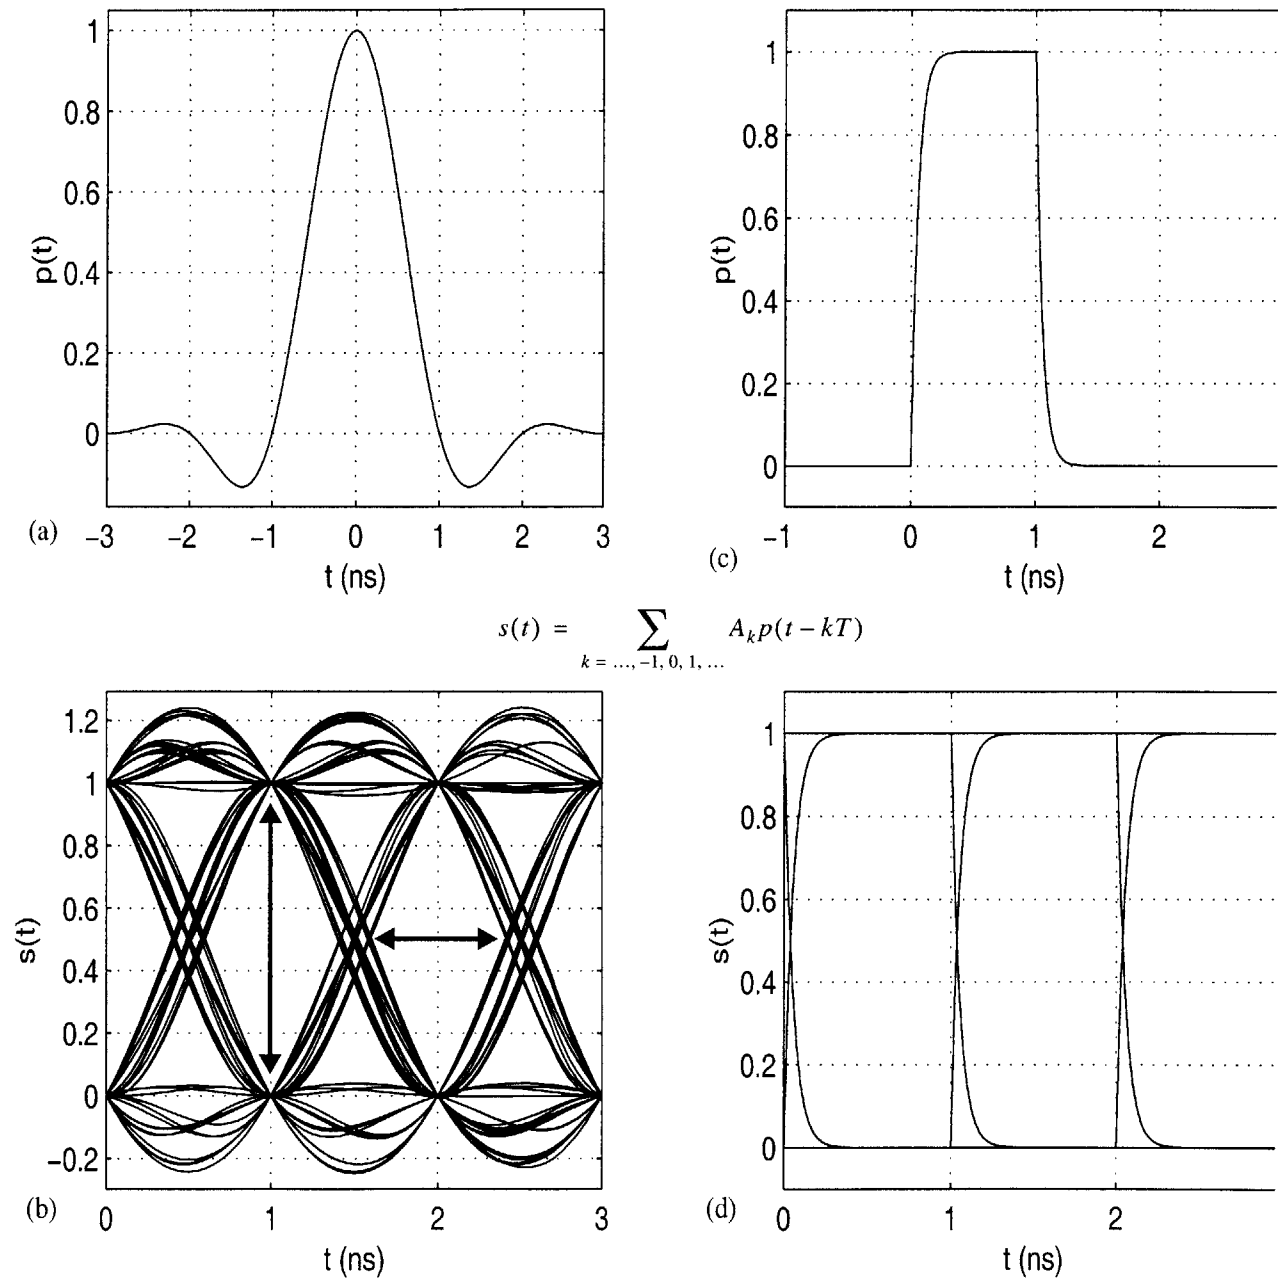


Figure D.1: Pulse shapes and data signals in communication links and digital interconnects

(a) is a typical pulse shape for a communication system (this is a raised cosine with 50% excess bandwidth [192]) (b) is a simulated eye pattern based on this pulse. The arrows indicate the vertical and horizontal "eye openings". (c) is a pulse shape typical of a digital interconnect. It has an RC-limited bandwidth of 3 GHz. (d) is a digital interconnect eye pattern.

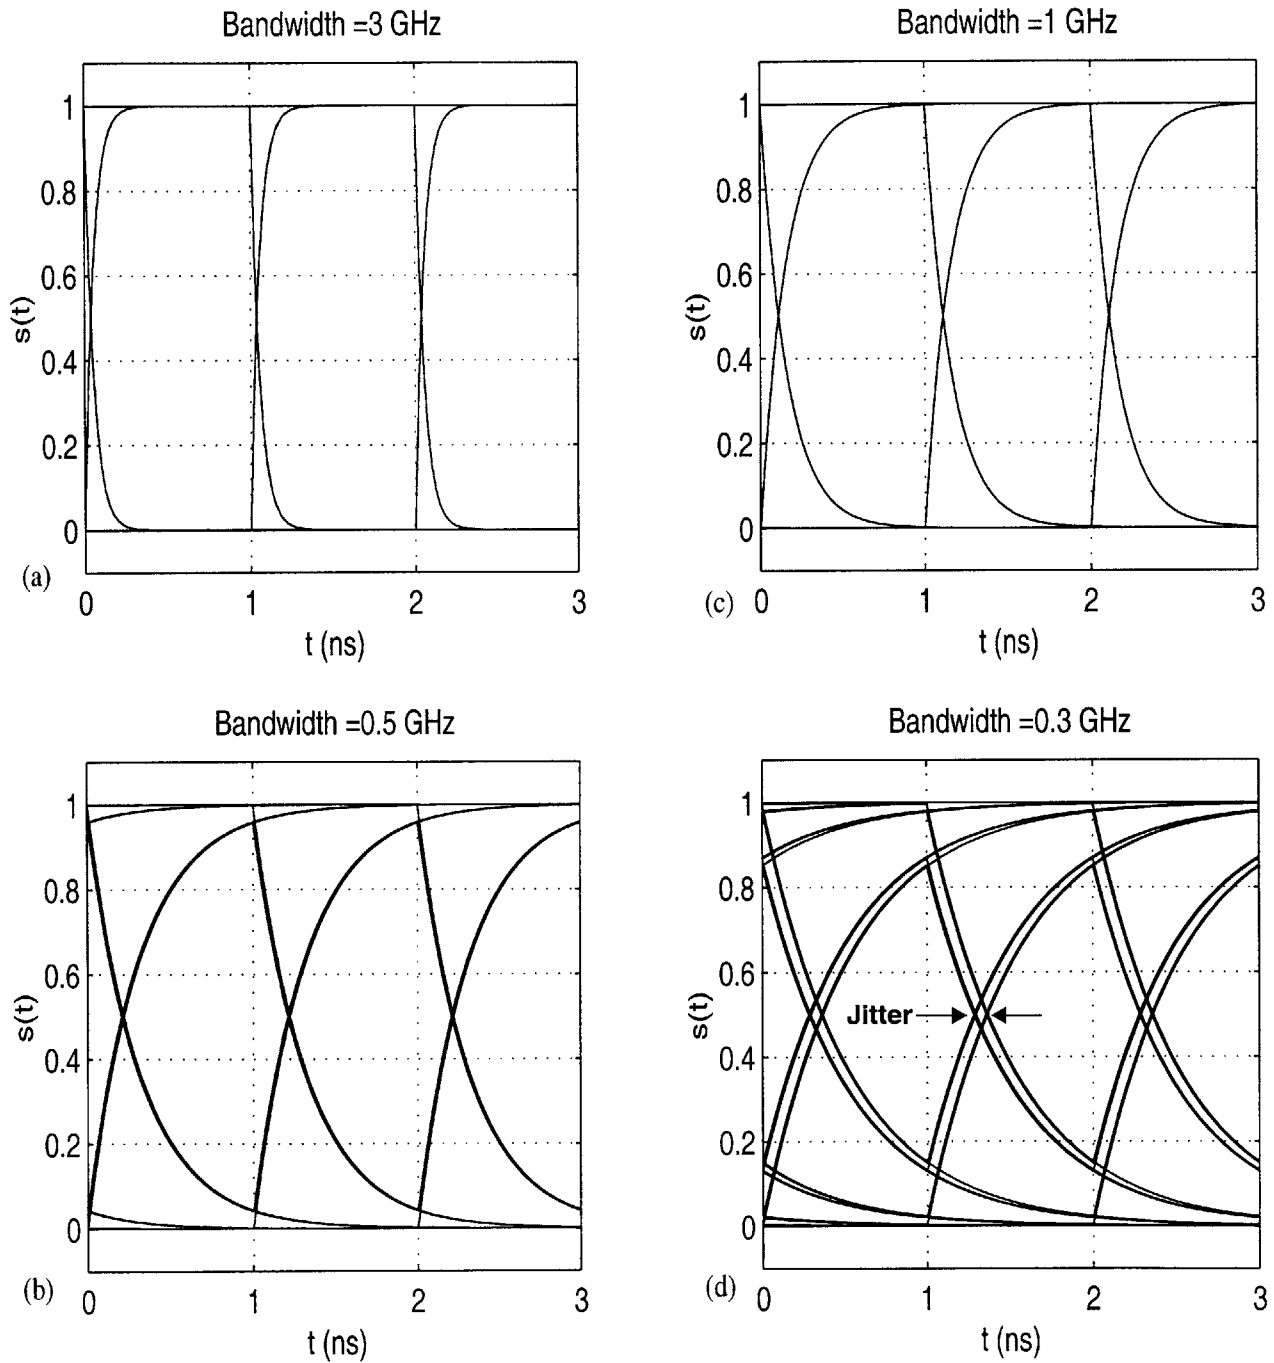


Figure D.2: ISI in a digital interconnect

As the bandwidth of the digital interconnect is reduced, the transitions take longer to settle. If they do not settle within one bit period, the eye opening just prior to the transitions is reduced. This translates into jitter.

contains information only about A_k . If the pulse is distorted so that so that this condition is not met, $s(t)$ is said to suffer from “intersymbol interference”, or ISI [192]. ISI is manifest in the eye diagram as “eye closure”, a reduction in the vertical and/or horizontal openings (see Figure D.1(b)). Vertical eye closure has the same effect as a reduction in signal-to-noise ration (SNR), while horizontal eye closure reduces the timing margins allowed in sampling $s(t)$.

In a digital interconnect, the situation is somewhat different. In general, $s(t)$ is not sampled. It may represent either a data or a clock signal. Ideally, $s(t)$ should make sharp transition between the specified noise-margins and remain there for there until the next transition. The effects of ISI on a digital signal are shown in Figure D.2. The time needed for the signal to transition between any given set of noise margins increases as the bandwidth of $p(t)$ is decreased. Also, the eye opening *immediately prior to the transition* is seen to effect the transition time. This is the mechanism by which ISI translates into jitter.¹

To quantify the ISI induced jitter in terms of the eye closure, consider the simplified waveform of Figure D.3. Suppose the signal has an amplitude of V_A and makes transitions in a time t_r . Near the center of the transition, and with the time axis shifted to this point, the waveform may be approximated as

$$s(t) = V_A \frac{t}{t_r} \quad (\text{D.2})$$

If the signal is initially displaced vertically, due to ISI, by an amount V_Δ , then the new waveform may be written

$$s'(t) = V_A \frac{t}{t_r} + V_\Delta \quad (\text{D.3})$$

The waveform now reaches the decision level, zero in this case, at time $t = -T_\Delta$, that is $s'(-T_\Delta) = 0$. Solving for T_Δ gives.

$$T_\Delta = \frac{V_\Delta}{V_A} t_r \quad (\text{D.4})$$

A similar calculation applies to both the rising and falling edges of the waveform. The total jitter would then be

1. There does not appear to be a great deal of existing research on ISI in digital systems. One exception is [238] which demonstrates the effect using simulations and suggest the use of clamping as a remedy.




Figure D.3: Simplified waveforms used to calculate ISI induced jitter

A signal of amplitude V_A makes a transition in time t_r . It is vertically offset, due to ISI, by an amount V_{Δ} .

$$T_{\Delta} = \frac{V_{\Delta, upper} + V_{\Delta, lower}}{V_A} t_r \quad (D.5)$$

The total eye-closure due to ISI, C_{ISI} , is recognized as the ratio

$$C_{ISI} = \frac{V_{\Delta, upper} + V_{\Delta, lower}}{V_A} \quad (D.6)$$

The peak-to-peak ISI induced jitter is thus seen to be given by the product of the transition time of the signal and the fractional eye closure measured immediately prior to the transition.

$$T_{\Delta} = C_{ISI} \cdot t_r \quad (D.7)$$

D.2 Jitter at the Output of a Comparator

The laser driver at the source end of a digital optical link may be thought of as a comparator. The logic signal at its input is compared to a decision level and the laser is set to the HI or LO state accordingly. Likewise, the receiver may be idealized as a comparator preceded by a linear filter which accounts for the frequency response of the initial, linear gain states. The linear components can create ISI, and ISI induced jitter, as shown above. This section is concerned with jitter that appears at the output of a comparator. The output jitter will include contributions from the jitter of the input signal, additive noise on the input signal, and the comparator's own timing variations.¹

Consider again the waveform of (D.2) appearing at the input of the comparator. Suppose that the input signal appears with a delay T_i and is corrupted by additive noise $n(t)$. The input signal may be written

$$s_i(t) = V_A \frac{t - T_i}{t_r} + n(t) \quad (\text{D.8})$$

When $s_i(t)$ reaches the decision level of the comparator, zero in this case, the output makes its transition following the comparator internal delay, T_d . Call the decision point T_1 so that $s_i(T_1) = 0$ may be solved to give:

$$T_1 = T_i - \frac{n(T_1)}{V_A} t_r \quad (\text{D.9})$$

For this expression to be well defined, it must be assumed that $s_i(t)$ crosses the decision level only once. This requirement is met if the maximum slope of the noise is less than the slope of the signal, a constraint which may be expressed as [242]

$$2\pi \cdot \frac{(\text{RMS Noise Voltage}) \cdot (\text{Noise Bandwidth})}{(\text{Signal Slope})} < 1 \quad (\text{D.10})$$

Adding the comparator delay, T_d , to (D.9), the time at which the comparator output changes may be written

$$T_o = T_i - \frac{n(T_1)}{V_A} t_r + T_d \quad (\text{D.11})$$

1. The problem solved here is a special case of the classic first-crossing time problem [239-241]. The solution presented here agrees with the empirical result used in [242] to analyze jitter in relaxation oscillators and in [243] to examine jitter in CMOS ring oscillators. The result is valid provided the slope of the signal is greater than the maximum slope of the noise. A number of recent studies of jitter in ring oscillators have also been reported [244-246].

Now, T_i , $n(T_I)$, and T_d can be taken to be random variables with variances σ_i^2 , σ_n^2 , and σ_d^2 respectively. Assuming the three are mutually independent, the variance of the output timing simply becomes

$$\sigma_o^2 = \sigma_i^2 + \frac{\sigma_n^2}{V_A^2} t_r^2 + \sigma_d^2 \quad (\text{D.12})$$

Or, recognizing V_A^2/σ_n^2 as the signal-to-noise ration, SNR ,

$$\sigma_o^2 = \sigma_i^2 + \frac{t_r^2}{SNR} + \sigma_d^2 \quad (\text{D.13})$$

To find the peak-to-peak jitter, rather than its variance, suppose that T_i , $n(T_I)$, and T_d have peak-to-peak variations of ΔT_i , ΔV , and ΔT_d , respectively. Then, it follows immediately from (D.11) that the worst case variation in T_o is

$$\Delta T_o = \Delta T_i + \frac{\Delta V}{V_A} t_r + \Delta T_d \quad (\text{D.14})$$

Also, note that V_A/t_r is the time derivative of $s(t)$ at the decision point. The last expression can thus be recast as

$$\Delta T_o = \Delta T_i + \Delta V \left(\frac{d}{dt} s(t) \Big|_{s(t) = \text{decision-level}} \right)^{-1} + \Delta T_d \quad (\text{D.15})$$

The latter form is more general in that it applies directly to waveforms which are not well approximated by the linear ramp transitions assumed in (D.8).

E

MESFET Drain-Lag

A recurring problem in the circuit designs presented in this thesis has been the frequency dependence of the effective MESFET output conductance. This phenomenon is often referred to as “drain-lag” in the literature, and its origins and some design techniques used to control it will be described in this appendix.

Figure E.1 shows the simulated drain current frequency response of Vitesse EFET. Not unexpectedly, the drain current response rises above 1 GHz as a result of the gate-to-drain capacitance. However, there is also a low-frequency rise in drain current response starting at around 1 MHz and leveling off at around 15 MHz. If this EFET is used in a basic common-source amplifier, the apparent rise in output conductance translates into a drop in gain, as seen in the frequency response shown in Figure E.2. In the time domain, this behavior translates into a fast initial response followed by an extended response tail. The step response in Figure E.3 illustrates this behavior.

E.1 Origin and modeling of drain-lag

The “low-frequency anomaly” in GaAs MESFET output conductance has been studied by a number of authors [229,247-251]. The majority of this work has dealt with either ion-implanted or epitaxial, recessed-gate MESFETs implemented directly on semi-insulating GaAs substrates. The output conductance was found to rise at frequencies below 100 kHz, and the effect was attributed to deep-level traps at the channel-to-substrate interface and at the device surface. However, no specific deep-level has been identified to account for the observed frequency and temperature dependencies. A number of models have been




Figure E.1: Frequency response of MESFET output conductance

A $10 \mu\text{m}$ wide, $0.5 \mu\text{m}$ long EFET was simulated as shown in the inset. At the bias point indicated, this device has a transconductance of $g_m = 1.5$ mS, backgate transconductance of $g_{mbs} = 196 \mu\text{S}$, drain conductance of $g_{ds} = 15.4 \mu\text{S}$, and gate-to-source and gate-to-drain capacitance of $C_{gs} = 6.6$ fF and $C_{gd} = 2.2$ fF, respectively. Although the backgate terminal is shown grounded, the device model assumes a finite resistance in this connection. In conjunction with the drain-to-backgate capacitance, this allows drain voltage variations to drive the backgate transconductance, giving rise to the effective increase in output conductance at around 1 MHz. The second apparent rise at 1 GHz is due to the gate-to-drain capacitance.




Figure E.2: Frequency response of common-source amplifier

The EFET of Figure E.1 was used in a common-source amplifier as shown in the inset. Because of the increase in the output conductance at 1 MHz, the gain of the amplifier drops at this frequency. It remains flat out the normally-predicted bandwidth of the amplifier.




Figure E.3: Step response of common-source amplifier


The common-source amplifier of Figure E.2 was used in a transient simulation. The applied gate-to-source voltage is shown in the upper trace and the output waveform is shown in the lower trace. In initial rapid rise corresponding to the high frequency roll-off of the amplifier is followed by a long response tail due to the low-frequency roll-off.

proposed for use in circuit simulation and have involved various combinations of non-physical resistors, capacitors, and controlled sources. A notable exception appears in [250,251] which reproduces the drain-lag effect by recognizing the presence of capacitive drain-to-backgate coupling and the finite resistance of the semi-insulating substrate. In this model, no specific mention is made of deep levels, and the drain-lag effect appears at a low frequency because of the large resistance of the semi-insulating substrate.

In the case of implanted, buried-channel MESFETs fabricated in a p-doped GaAs layer, similar to those in the Vitesse process, output conductance variations do not appear below 100 kHz as they do when the MESFET is fabricated directly in semi-insulating material [252]. This is consistent with the Vitesse devices. Unfortunately, the devices in [252] were not characterized above 100 kHz. A rather elaborate model involving a parasitic npn transistor was proposed in [253] to reconcile this observation with the deep-level theories.

It is not clear whether or not deep levels play a role in the Vitesse MESFETs. As in [250,251], the inclusion of drain-to-backgate capacitance and finite backgate resistance may be used to model the frequency dependence of the output conductance. This is the general approach taken in Vitesse's HSPICE models, though the details are of a proprietary nature. Figure E.4 outlines a simple implementation of a drain-lag model. A pole-zero pair appear in the drain conductance accounting for the observed behavior. According to this model, the output conductance assumes its intrinsic value, g_{ds} , at low frequencies and should start to rise at around the zero frequency of $1/(2\pi(2+g_{mbs}/g_{ds})R_B C_{DS})$ where R_B is the resistance to the backgate contact and C_{DS} is the drain-to-backgate junction capacitance. The drain conductance increases by one-half the value of the backgate transconductance then levels off due to the presence of the pole. The ratio of the pole and zero frequencies is $(1+g_{mbs}/2g_{ds})$. The initial and final conductance values and the width of the transition region predicted in this way are consistent with Figure E.1.

The predicted zero frequency tends to be somewhat high compared to the Vitesse HSPICE model results in Figure E.1. C_{DB} for the EFET of Figure E.1 is ~ 1 fF, and the backgate resistance is ~ 50 k Ω , though these values are not known with great certainty. The resulting zero would be at ~ 130 MHz. This estimate is more consistent with the experimental laser driver results than the Vitesse HSPICE model. The Vitesse HSPICE model overestimates the effects of drain-lag [211].



(b)
$$\frac{V_{BS}}{V_{DS}} = \frac{sR_B C_{DB}}{1 + sR_B(C_{DB} + C_{SB})} \approx \frac{s\tau}{1 + s(2\tau)} \quad \tau = R_B C_{DB} \quad (C_{DB} \approx C_{SB})$$

(c)
$$g_{out} = g_{ds} + g_{mbs} \frac{V_{BS}}{V_{DS}}$$

(d)
$$g_{out} \approx g_{ds} \left\{ \frac{1 + s \left[\left(2 + \frac{g_{mbs}}{g_{ds}} \right) \tau \right]}{1 + s[2\tau]} \right\}$$

Figure E.4: A simple model for drain-lag

(a) Shows a cross-section of a MESFET indicating the source- and drain-to-backgate capacitances and the resistance of the backgate connection through the blanket p⁻ implant layer. With the source and the backgate contact grounded, (b) shows the response of the backgate voltage, V_{BS}, to variations in the drain voltage, V_{DS}. Then, (c) points out that the output conductance has contributions from both the intrinsic drain-to-source conductance, and the backgate transconductance which is driven by the drain voltage through the drain-to-backgate coupling. Finally, (b) is substituted into (c) to give the frequency dependent output conductance in (d)

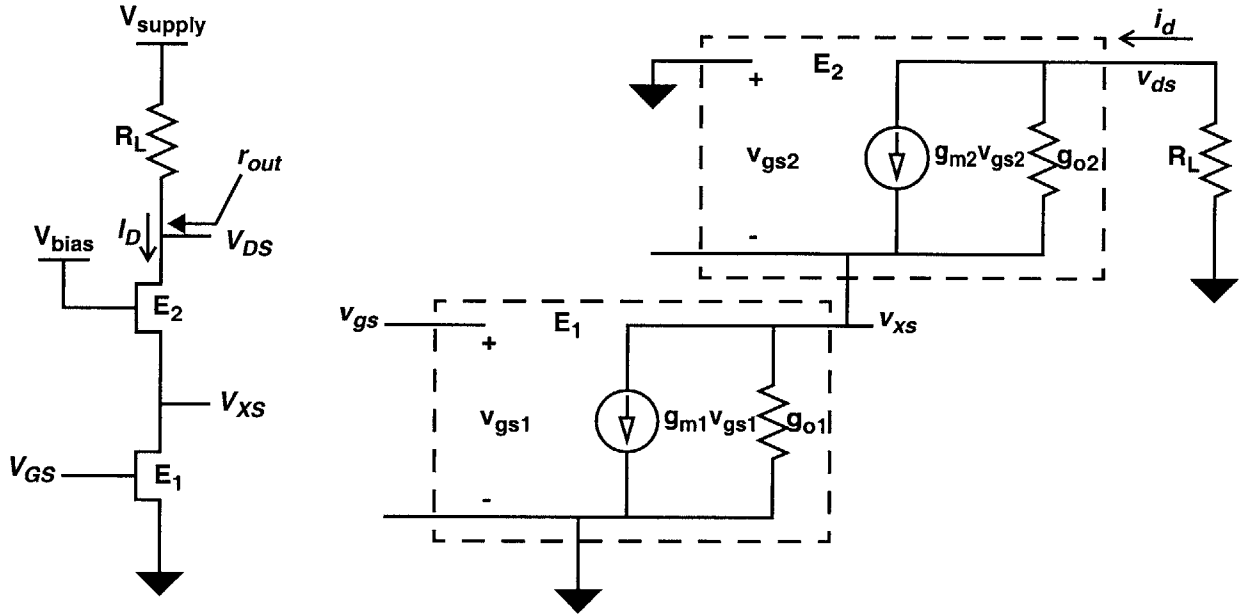
E.2 Controlling Drain-lag

Regardless of the details of its origin, circuit techniques are needed to control the effects of drain lag. Prior to listing three techniques used in this work, it is worth noting that in some applications, it should be possible to sidestep drain-lag by limiting the frequency range of operation to above the onset of drain-lag. For instance, the data in the optical link could be coded so that the lowest frequency component in the transmitted signal is above the onset of the drain-lag effect. An example of this would be a Manchester code. The use of such coded links was not considered in this work which pursued the most general case of a digital optical interconnect.

Controlling drain-lag using a cascode

One approach to controlling the effects of drain lag on a transistor is to reduce its drain voltage variation by using a cascode circuit topology. An example of this in the case of a common-source amplifier is shown in Figure E.5. This type of cascoded amplifier is commonly used in CMOS design to reduce the output conductance of the gain transistor (in order to get higher gain) and to eliminate the input-output coupling capacitor which would otherwise be Miller-multiplied. The presence of the common-gate stage, E_2 , reduces the amplitude of the drain voltage on E_1 by a factor of $[(g_{out}+g_{o1})/g_{m2}]$ relative to what it would be without the cascode. This should reduce drain lag since it is driven by this same drain voltage.

An alternate cascode topology has also been reported which, in some cases, is more effective at controlling drain lag [254,255]. Rather than setting the gate of the cascode transistor to a separate bias voltage, it is connected directly to the gate of the gain transistor (suggesting the name “active cascode”). The gain transistor may be operated in saturation if the threshold voltage of the cascode transistor is lower than that of the gain transistor. In the Vitesse process, this is accomplished by using a DFET for the cascode transistor and an EFET for the gain transistor. This topology is shown in Figure E.6. From the perspective of the drain of the composite transistor, this topology is a small-signal equivalent to the original one, and the drain conductance is reduced by the same factor. But, the cascode transistor is now driven by the input. This feature may be used to further reduce variations in the drain voltage of the gain transistor. Relative to an un-cascoded common-source amplifier, the drain voltage amplitude is now reduced by a factor of $[(g_{out}+g_{o1})/g_{m2} - g_{out}/g_{m1}]$. In principle, the amplitude of this signal may be reduced to zero by sizing the



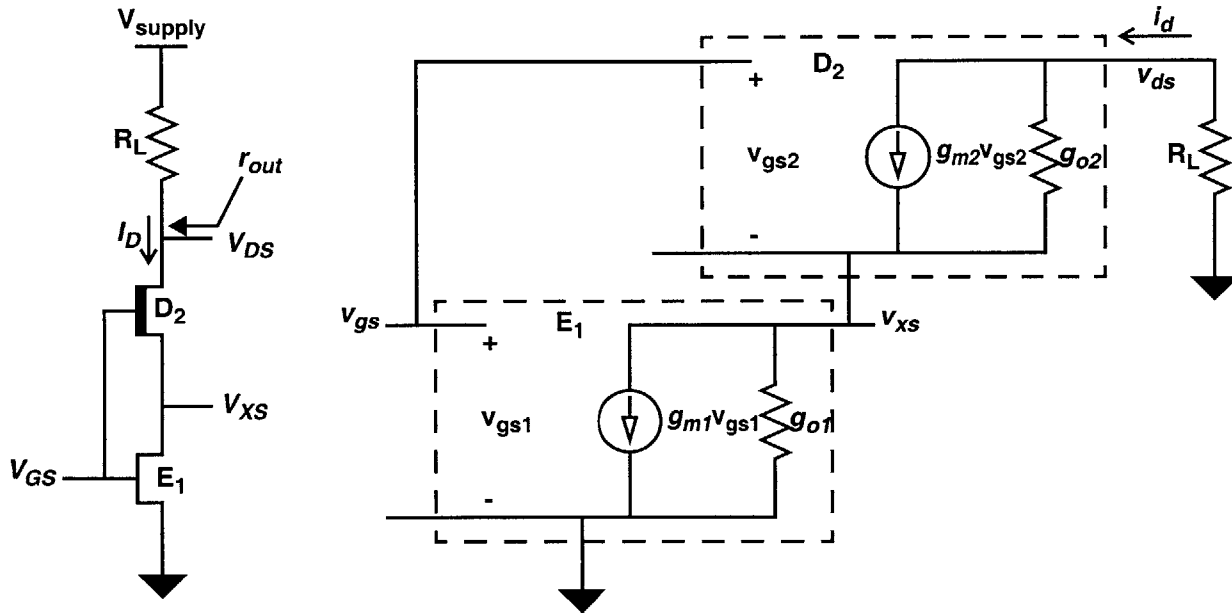
$$i_d \approx -g_{m1} \left[1 + \frac{g_{o1}}{g_{m2}} \right] v_{gs} + \frac{g_{o1}g_{o2}}{g_{m2}} v_{ds}$$

$$v_{ds} \approx -g_{m1} r_{out} v_{gs} \quad r_{out} = \frac{1}{g_{out}} \approx R_L \parallel \frac{g_{m2}}{g_{o1}g_{o2}}$$

$$v_{xs} \approx -\frac{g_{m1}}{g_{m2}} (1 + g_{o2} r_o) v_{gs} = -\left(\frac{g_{out} + g_{o1}}{g_{m2}} \right) g_m r_{out} v_{gs}$$

Figure E.5: A common-source amplifier using a cascode

In this common-source amplifier, gain transistor E_1 is cascoded by E_2 . This type of cascode structure is common in CMOS design. Compared to simply using E_1 without the cascode, the transistor output conductance is decreased from g_{o1} to $g_{o1}g_{o2}/g_{m2}$ while the transconductance, g_{m1} , is not significantly altered. The Miller-multiplied capacitance coupling the output and input is also eliminated. Drain lag is improved since the amplitude of the E_1 drain voltage, v_{xs} , is reduced from the un-cascoded value by the factor $(g_{out} + g_{o1})/g_{m2}$.



$$i_d \approx -g_{m1} \left[1 + \left(1 - \frac{g_{m2}}{g_{m1}} \right) \frac{g_{o2}}{g_{m2}} \right] v_{gs} + \frac{g_{o1}g_{o2}}{g_{m2}} v_{ds}$$

$$v_{ds} \approx -g_{m1} r_{out} v_{gs} \quad r_{out} = \frac{1}{g_{out}} \approx R_L \parallel \frac{g_{m2}}{g_{o1}g_{o2}}$$

$$v_{xs} \approx -\frac{g_{m1}}{g_{m2}} \left(1 + g_{o2} r_{out} - \frac{g_{m2}}{g_{m1}} \right) v_{gs} = -\left(\frac{g_{out} + g_{o1}}{g_{m2}} - \frac{g_{out}}{g_{m1}} \right) g_{m1} r_{out} v_{gs}$$

Figure E.6: A common-source amplifier using an “active” cascode

Like the amplifier in Figure E.5, this stage uses a cascode to reduce the transistor output conductance from g_{o1} to $g_{o1}g_{o2}/g_{m2}$ without significantly altering the transconductance, g_{m1} . Using a DFET for the cascode allows the additional bias voltage to be eliminated, although the Miller-multiplied feedback capacitance is again encountered. This topology is advantageous in controlling drain-lag, however, because the amplitude of the E_1 drain voltage, v_{xs} , can be approximately zeroed out by setting $g_{m2} \sim (1 + g_{o2}r_o)g_{m1}$.

transistors to set $g_{m2}=(1+g_{o2}r_o)g_{m1}$. In practice, it is difficult to achieve perfect cancellation, but this topology has generally been found in simulations to be more effective at reducing drain lag than the standard cascode.

It is important to recognize that the voltage on the drain of the cascode transistor (E_2 or D_2) is varying even though the drain the drive transistor (E_1) is not. If the drive and cascode transistors share a common backgate node, the signal on the cascode drain can still drive the backgate and cause drain lag. The magnitude of the drain lag may be smaller than the non-cascoded circuit, but to realize the full benefit of the cascode topology, the drive and cascode transistors can be placed in separate isolation rings.

Controlling drain-lag using feedback

If negative feedback, with a feedback factor f , is applied around an amplifier with transfer function $A(s)$, the transfer function of the composite system is $A(s)/(1+A(s)f)$ [256]. As long as $|A(s)| \gg |1/f|$, the composite system response is approximately $1/f$. If $A(s)$ represents a MESFET gain stage effected by drain lag, using it in conjunction with negative feedback can result in an amplifier with a broad, flat transfer function.

When implementing a transimpedance amplifier (TIA), negative feedback appears naturally and corrects the drain-lag effects. In principle, voltage amplifiers could also be constructed which include feedback, but this practice was avoided in the designs pursued in this thesis since it would require buffering between stages in order to drive the input resistance of the feedback-based voltage amplifier.

Controlling drain-lag by clamping

Drain lag would appear on the rising edge of DCFL logic signals if not for the fact that the output of each gate is clamped by the Schottky diode at the input of the of the next gate. This is demonstrated in Figure E.7. Hard-limiting can be used in this way to control drain lag in other contexts. EFET and DFET Schottky gates may be used to attain slightly different voltage drops (EFET voltage drop is a little higher). An EFET or DFET with its drain and gate connected may also be used to form a rectifier with a smaller voltage drop. Signals may be clamped relative to ground or relative to a supply. Differential signals may also be clamped relative to each other by using a parallel combination of opposing diodes.

It is worth noting that although no extended response tail appears on the clamped signal, the backgate node is still being driven by the drain voltage and will evolve on a long time scale. This will appear as a




Figure E.7: Clamping controls drain lag in DCFL logic

DCFL (direct-coupled FET logic) circuits experience drain-lag in the same way that analog MESFET circuits do. As long as the output of each gate is connected to the input of another DCFL gate, the Schottky diode at the input (i.e. the gate of the EFET) clamps the output at ~ 0.6 V and does not allow the extended drain-lag response tail to be manifest.

pattern-dependent variation in the MESFET threshold voltage. In practice, this does not appear to be a large effect.

F

First-Generation Optical Receiver

Figure F.1 is the MIT-OEIC-6 test chip. This OEIC was released after the OPTOCHIP project and contains the first set of transmitter and receiver interface electronics to be designed as part of this thesis.




Figure F.1: MIT-OEIC-6

This test chip contains the first set of transmitter and receiver interface electronics developed in this thesis. It implemented in the Vitesse H-GaAs III process. The central 2 mm x 2 mm region of the 4 mm x 4 mm chip contains growth-related cells while the peripheral regions is dedicated to electronic test structures.

F.1 The First-Generation Receiver Circuit

A block representation of the first-generation receiver circuit is given in Figure F.2. In this circuit, a fully-differential design is used for supply noise immunity. This receiver, designed in the H-GaAs III process and operating from a single 2 V supply, is entirely self contained. That is, no external input is required to set the decision level. This approach simplified packaging and testing, but does not allow enough flexibility for experimentation. While it may ultimately be desirable to return to this type of design, the receivers which came after this used an external reference.

The first stage transimpedance amplifier (TIA) is “self-referenced” by feeding back the common-mode output voltage to the positive input terminal. The resulting outputs are differential, but not logically complementary. That is, with no optical input applied, the outputs of the TIA are equal. As the photocurrent is increased, the TIA outputs grow apart.

To examine the stability of this feedback arrangement, define the common-mode signal as the average of the signal pair (either input or output) and the differential-mode signal as the difference of the signal pair. The common-mode output is related to the common-mode input by $V_{OC} = -(A_C)(V_{IC})$ where A_C defined to be the common-mode gain. Similarly, the differential-mode signal are related by

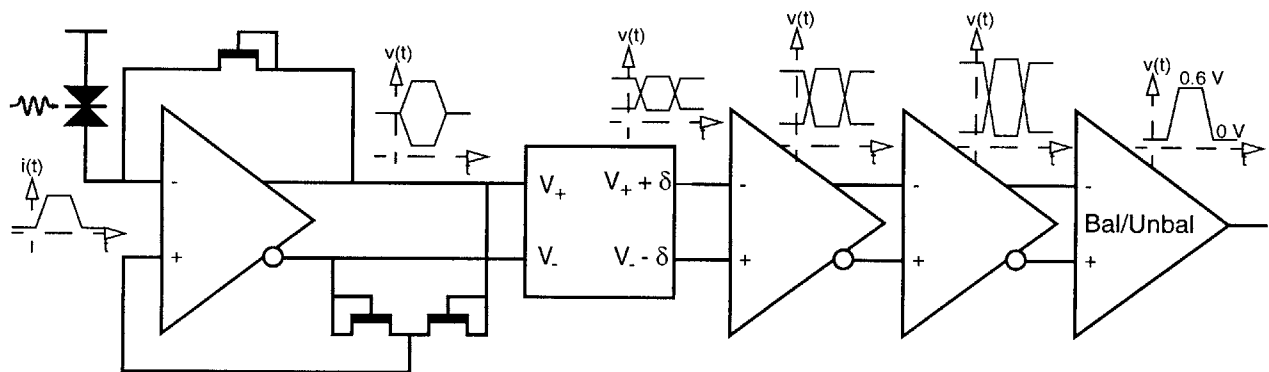


Figure F.2: First generation optical receiver: block diagram

A fully-differential design is used for supply noise immunity. This receiver, designed in the H-GaAs III process and operating from a single 2 V supply, is entirely self contained. The first stage transimpedance amplifier is “self-referenced” by feeding back the common-mode output voltage to the positive input terminal. The resulting outputs are differential, but not logically complementary. An offset stage shifts each of the two signal relative to each other in order to produce a logically complementary pair. This differential signal is then further amplified before combined into a single-ended direct-coupled FET logic (DCFL) logic signal.

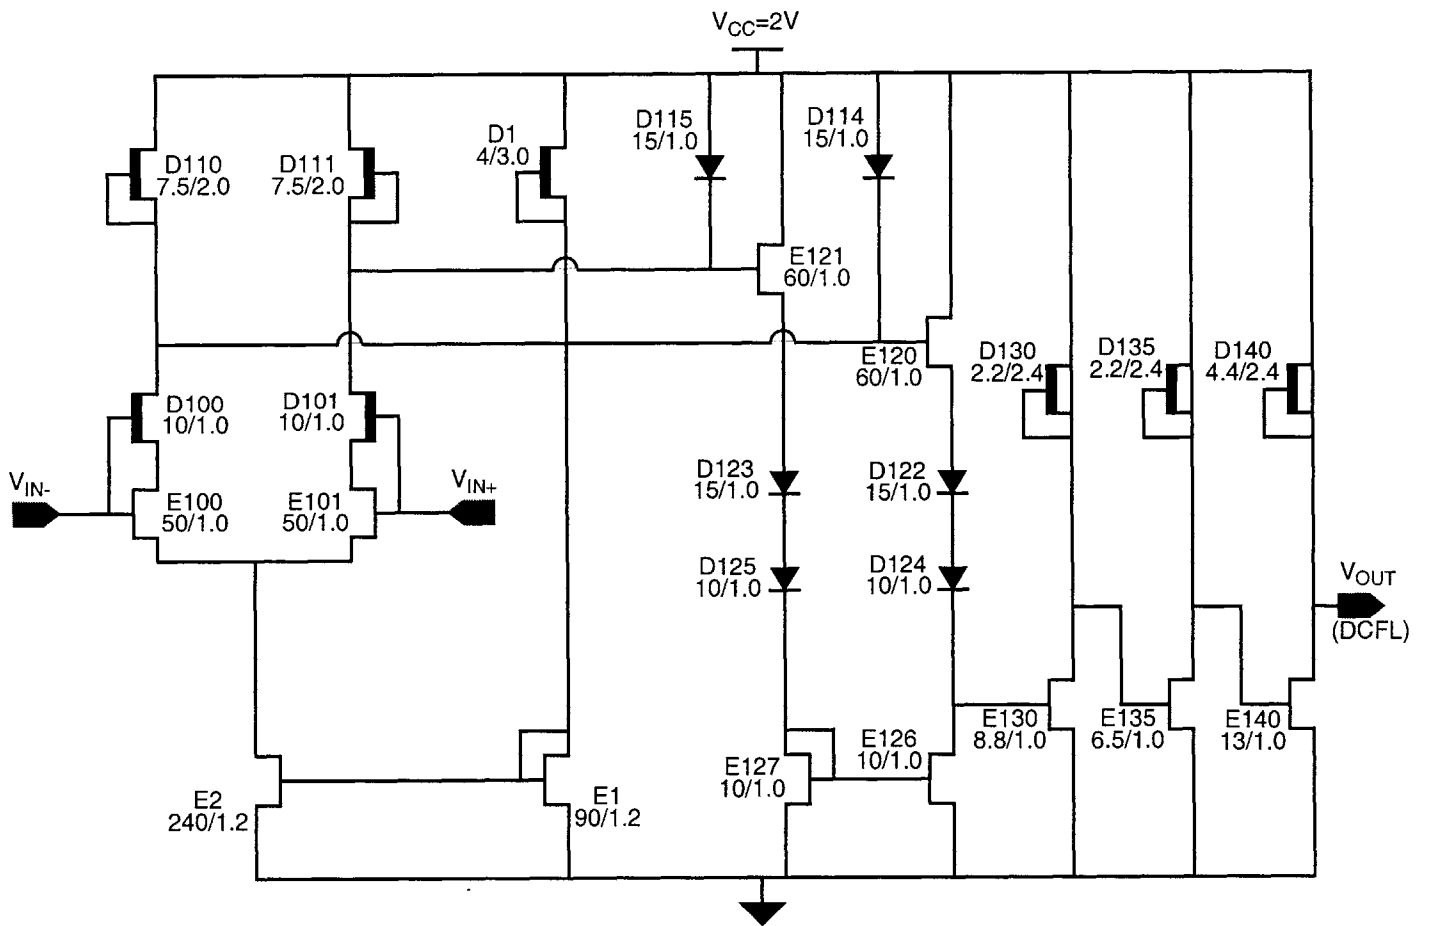
$V_{OD} = +(2A_D)(V_{ID})$ where A_D is defined as the single-sided differential-mode gain. With these definitions, the differential-mode TIA transfer function is $(V_{OD}/i) = -2\{A_D/(1+A_D)\}$ while the common-mode TIA transfer function is $(V_{OC}/i) = -(1/2)\{1/(1+A_D)\}\{A_C/(1+A_C)\}$. From this it is seen that the stability considerations for the differential mode signal are identical to the case in which a constant reference is applied to positive terminal. Namely, provided that A_D itself is stable, the differential mode signal will be stable if the denominator, $1+A_D$, has no poles in the right-half complex plane. This same requirement applies to the common-mode signal since the same denominator term is present. An additional denominator factor, $1+A_C$, also appears in the common-mode transfer function. However, in a well designed differential amplifier the magnitude of the common-mode gain will be less than one so that stability of this term is assured. Thus, the overall feedback arrangement used for the self-referenced first-generation receiver has the same stability considerations as an externally referenced design. This conclusion is borne out by both simulation and measurement. A detailed schematic of the first-stage TIA is given in Figure F.3.

The outputs of the self-referenced TIA may not, by themselves, be used to make a logic-value decision. An offset stage which follows the TIA shifts each of the two signal relative to each other in order to produce a logically complementary pair. The magnitude of the shift determines the decision level of the receiver. The outputs of this stage are further amplified before being combined into a single-ended direct-coupled FET logic (DCFL) logic signal.

A schematic of the offset stage is provided in Figure F.4. The circuit consists of a pair of DFET source followers and a current mirror. The sizing of the source-follower DFETs relative to the DFET which sets the reference current for the mirror result in one source follower having a slightly positive gate-to-source voltage and the other an equal but negative gate-to-source voltage (note that the reference current DFET has zero gate-to-source voltage). The two signal paths are kept nearly identical so that the bias variations and noise appear as common-mode signals.

The second and third gain stages are straightforward differential voltage amplifiers and are shown schematically in Figure F.5 and Figure F.6 respectively. The third stage is followed by a differential-to-single-ended conversion circuit. The resulting single-ended signal is passed through three DCFL inverters to sharpen the edge transitions before outputting a DCFL logic signal.

Note that the use of diode-connected DFET load devices in this circuit was eliminated in later design to achieve greater bias-point stability and better control of drain-lag. The latter issue arises because drain-



F.2 Optical Link Demonstration


An initial study of a fiber-based optical link was completed using the H-GaAs III transmit and receive electronics on MIT-OEIC-6 and is summarized in Figure F.7. The laser driver was a basic switched current mirror design and was discussed in Chapter 6. It occupies an area of $65\ \mu\text{m} \times 70\ \mu\text{m}$. In this demonstration, an external n-side-up oxide-confined 850 nm VCSEL was wire-bonded directly to the laser driver; simulations have shown that the short ($\sim 1\ \text{mm}$) bond wire does not significantly effect the transmitter performance. The VCSEL was of the same type detailed in Chapter 6. The laser driver and VCSEL together consumed less than 3 mW of power in the link tests conducted.

The VCSEL output was coupled into a $62.5\ \mu\text{m}$ graded-index-core multimode-fiber. To prevent external-cavity effects on the VCSEL due to reflections from the fiber a separation of roughly $500\ \mu\text{m}$ was maintained between the VCSEL and the fiber. The other end of the fiber was directly abutted to the receiver. The receiver circuit described above was used in conjunction with an integrated $40\ \mu\text{m} \times 40\ \mu\text{m}$ MSM photodetector. This detector was implemented using the standard H-GaAs III process, as explained in Section 4.1. The entire receiver occupied $170\ \mu\text{m} \times 185\ \mu\text{m}$ and consumed 7.2 mW of power.


This link was operated at 200 Mb/s with a bit error rate of 10^{-12} in response to a $2^{23}-1$ pseudorandom bit sequence.¹ Figure F.7(b) shows the corresponding output eye pattern. Given a valid logic swing at its input, the receiver generates a well-defined output logic level. Under these circumstances, the performance of the link is reflected primarily in the uncertainty in the timing of its output.

Figure F.7(b) understates the jitter due to the limited run-length of the pseudorandom bit sequence. More detailed experiments have shown the timing jitter of this link to be roughly 5 ns. This significant jitter is primarily the result of fluctuation in the backgate voltage of the receiver circuit in response to the applied optical signal. This backgate optical crosstalk is discussed in Section 3.5. As seen in the later receiver designs, the backgate problem is eliminated by the p-contacts and isolation structures provided by the H-GaAs IV process.

1. Measured with a Tektronix GigaBert700 on loan from the manufacturer.



(a)



(b)

Figure F.7: First generation optical link demonstration

The basic switched current mirror laser driver discussed in Chapter 6 and the receiver described in this appendix were used in a preliminary optical link demonstration. The experimental setup is summarized in (a) and an eye diagram of the receiver output operating at 200 Mb/s is shown in (b). This eye pattern was generated in response to a 2^{23} -1-bit pseudorandom bit sequence and achieved a bit error rate of less than 10^{-12} .

G

Receiver-Array OEICs

The OPTOCHIP project, discussed in Chapter 3, sought to aid optical interconnect development by making fully-customized, monolithic, emitter-based OEICs available to system researchers. A complementary effort is being led by the Optoelectronic Industry Development Association (OIDA) and the Japan Optics Program (JOP).¹ Under this initiative, modular components based on a variety of advanced research-grade technologies are being made available to research groups in the United States and Japan. OEICs featuring arrays of the MSM photodetectors and receivers studied in this thesis have been realized as part of the OIDA effort and will be described in this chapter.

G.1 General Considerations

Much of the OIDA effort has been oriented around 8x8 arrays of VCSELs (vertical cavity surface emitting lasers) being provided by one or more industrial participants. Due to the experimental nature of the MSM photodetectors and receivers, and more significantly by cost constraints, the initial receiver array designs were fixed at 4x4. With regard to the cost issue, the primary difficulty in scaling from a 4x4 to an 8x8 array lies in the escalation of die area to roughly 6.5 mm square. This large die area, $> 40 \text{ mm}^2$, is needed to support the many bondpads consumed by the 64 ECL outputs--preferably differential--and the accompanying power and ground connections (at least one power and ground for every four ECL outputs). No other barrier to scaling the 4x4 receiver arrays to 8x8 versions is foreseen.

1. Optoelectronic Industry Development Association, 2010 Massachusetts Avenue, NW, Suite 200, Washington DC, 20036-1023, (202)785-4426.

Two different 4x4 receiver array OEICs have been designed and fabricated on the Vitesse H-GaAs IV process through the MOSIS service. The first of these, MIT-OEIC-8, uses the low power receiver detailed in Chapter 7. Subsequent to its design, a change in the MOSIS fabrication schedule allowed the design of a second receiver array, MIT-OEIC-9. A new receiver circuit, found in Appendix H, was designed specifically for this application and addressed a number of the difficulties encountered in MIT-OEIC-8.

Both receiver arrays use the improved MSM photodetectors analyzed in Chapter 4. Specifically, they feature 75 μm diameter detectors with 2.1 μm finger spacings. The detectors are placed on a 250 μm x 250 μm pitch. The die used are 2.3 mm x 2.3 mm, though some die are slightly larger due to due to handling at MOSIS. The die are 27 mil (686 μm) thick.

Each of the sixteen MSM photodetectors on an array chip are coupled to an adjacent receiver circuit. Dedicated ECL outputs are provided for each of the sixteen receivers. The ECL output levels are referenced to the +2V supply and thus assume nominal low and high values of 0V and 1.2V. They must be terminated to ground (0 V) through a 50 Ω load.


Packaging, supply requirements, and board design guidelines

In-house testing of the array OIEC was carried out using a chip-on-board approach. The board used included four copper layers and was designed at McGill University to support a variety of receiver array applications [257]. Parts have been distributed by OIDA in both bare die and packaged form. Packaged parts made use of the gold-plated glass-wall flat-pack diagrammed in Figure G.1. This package, part number 5M34M-3 from Mini-Systems, Inc. (see figure for contact information), is designed for high speed digital and microwave applications. While the pad layout of the OEICs was designed around this package, other high speed packages may be used. Great attention must be paid to both package selection and board layout in order to successfully operate the components. For example, attempts at a preliminary tests using a 40-pin DIP package failed because the circuits would oscillate uncontrollably upon power-up due to the large parasitic inductances of the package leads. The die-attachment and wire bonding of the OIDA distributed chips was carried out by Micro Hybrids.¹

1. Micro Hybrids, Inc., 2864 Route 112, Medford, New York 11763, <http://www.microhybrids.com>, (516)732-3448.

M S I

MINI-SYSTEMS, INC.
 Electronic Package Division
 P.O. Box 1597 - Plainville, MA 02762
 Tel. (508) 695-2000 - FAX (508) 695-8758



RECOMMENDED COMBINATION LID: LSM-5 SUGGESTED LID DIMENSIONS IF PURCHASED ELSEWHERE: .490 X .350 X .010 THK.

Figure G.1: Diagram of high speed package for use with MIT-OEIC-8 and MIT-OEIC-9
 This is a gold-plated glass-wall flat pack and is intended for high speed digital and microwave applications.

Parts making use of the package of Figure G.1 employ down-bonds to the metallic base of the package for all ground connections. It is important that the base of the package be attached, with both a low electrical and thermal impedance, to the ground plane of the circuit board.

As described below, the array OEICs require +2V and +3.3V supplies in addition to a photodetector bias and a reference voltage. The MSM photodetector bias has been set to +5 V in testing the OEIC. Use of the +3.3 V supply for the MSM bias may be adequate, but this option has not been carefully studied.

It is critical that the ground and the +2V and +3.3V supplies be delivered to the IC with very low inductance. On the board that was used for testing, ground was made available on a dedicated plane and the two supplies were provided on an additional, split plane. Ideally, each of these would have a separate plane. Also, each of the ECL outputs will attempt to switch 22 mA in around 200 ps between the +2V supply and ground. As shown below, each set of eight ECL outputs is powered through three bondpads. In a simultaneous switching event, each of these three +2V connections will need to supply around 60 mA in this 200 ps time frame. This is a large current transient, and in addition to the dedicated supply plane, good bypassing is needed in order to minimize supply bounce. Careful attention must be paid to via inductances and the parasitic inductance of the bypass capacitors. The test boards used a parallel combination of a 0.1 μ F ceramic (1206 surface mount) and a 10 μ F tantalum (also 1206) capacitors placed on each of two connections to each of the supplies. This design was based on convenience; better bypassing arrangements are possible.¹

G.2 Receiver Array “MIT-OEIC-8”

Basic design

MIT-OEIC-8 uses the low power receiver of Chapter 7. A layout plot of MIT-OEIC-8 is shown in Figure G.2, and the layout of one unit cell of the array is shown in Figure G.3. The pad identifiers noted in the figure are defined in Table G.1 and will be described shortly. A bonding diagram for MIT-OEIC-8 using the package of Figure G.1 is shown in Figure G.4. This bonding arrangement, which was used in the OIDA distributed parts, results in the pinout given in Table G.2.

1. For an excellent discussion of high-speed board-level design, including the design of bypass networks, see: H. W. Johnson and M. Graham, *High-Speed Digital Design: A Handbook of Black Magic*, Prentice Hall, New Jersey, 1993.

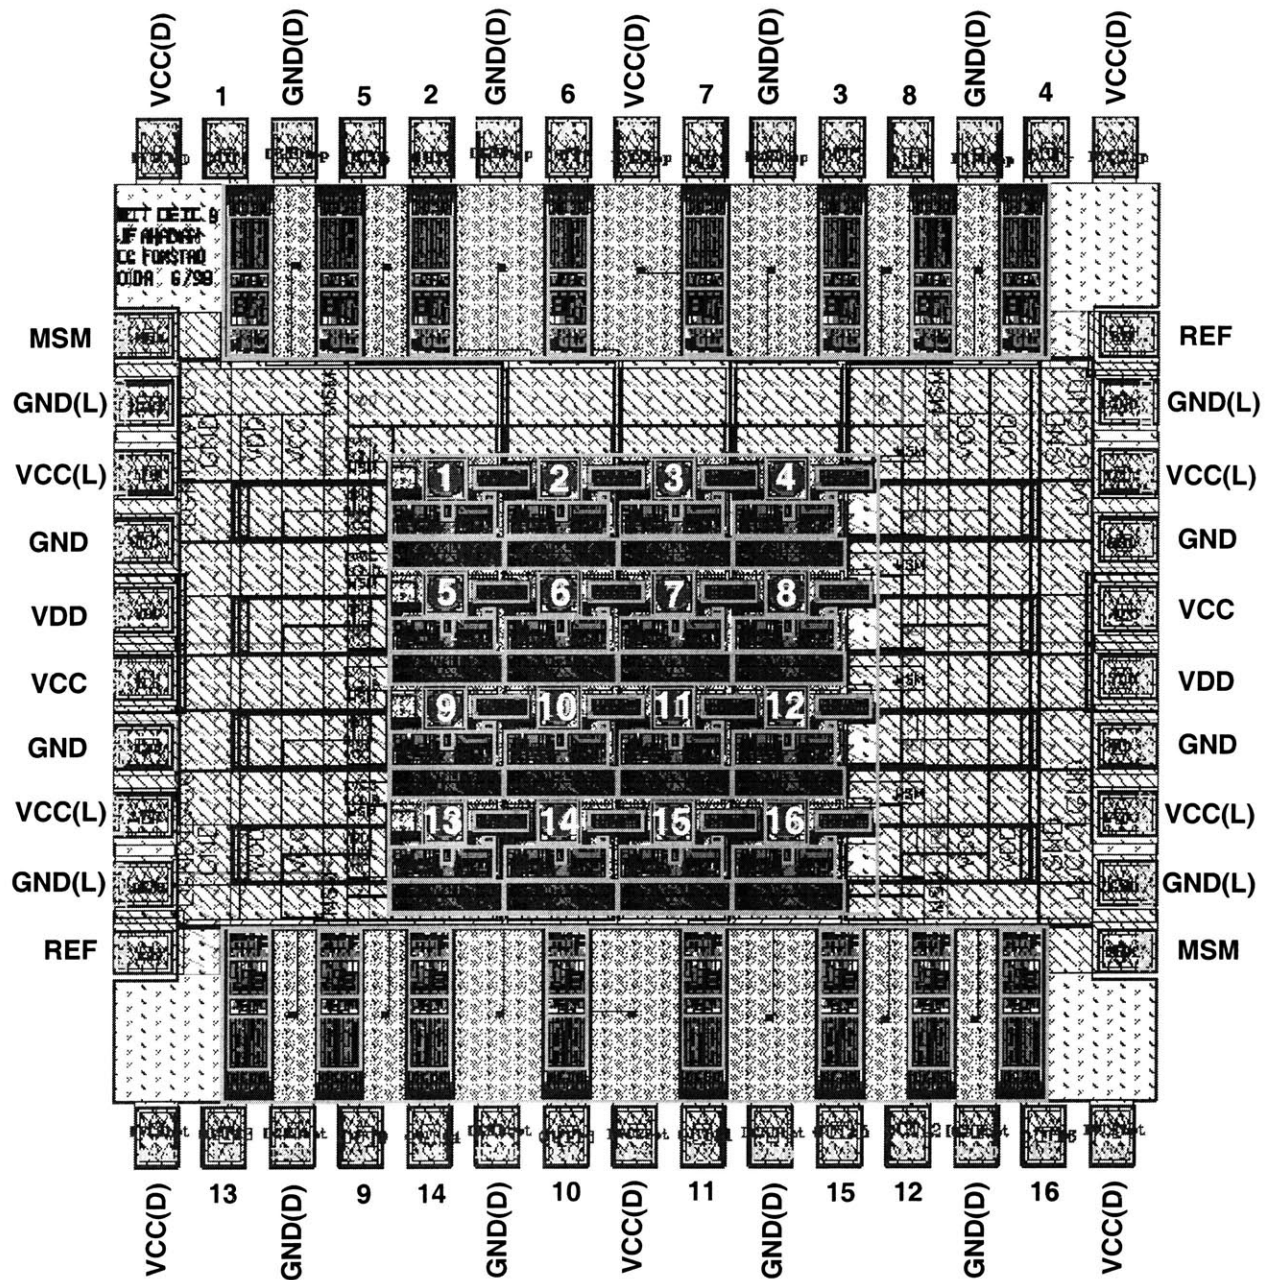


Figure G.2: Layout plot of MIT-OEIC-8

The first of two 4x4 receiver arrays fabricate using the Vitesse H-GaAs IV VLSI MESFET process. The die is 2.3 mm x 2.3 mm and provides an ECL output for each of the sixteen optical inputs. The pad names are defined in Table G.1.

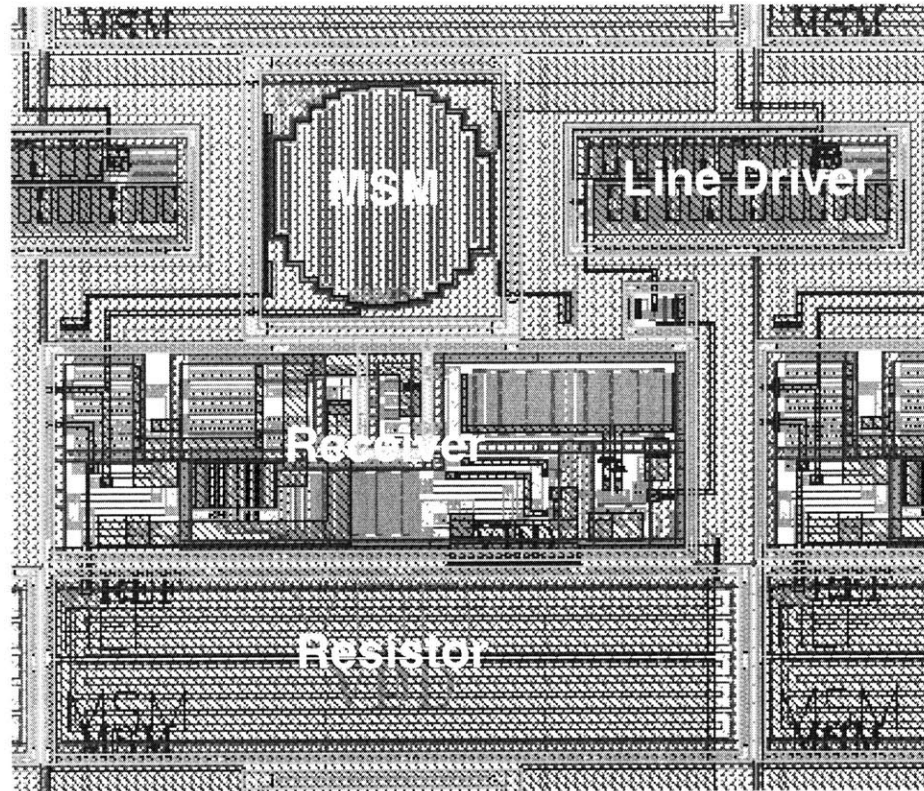


Figure G.3: Unit cell of the MIT-OEIC-8 receiver array

Each of the elements in the 4x4 receiver array includes a 75 μm diameter MSM photodetector, a receiver circuit with a DCFL (direct coupled FET logic) output, and a line driver to deliver the signal to the ECL output buffers along the top and bottom edges of the OEIC. A resistor, nominally 200 k Ω , is also used in each unit cell to set the reference current against which the photojournalist from the MSM is compared to determine the logic state of the output. The unit cells are arrayed on a pitch of 250 μm x 250 μm to form the array.

In Figure G.2, the ECL output drivers are located along the top and bottom of the chip and each set of eight has its own set of three supply and four ground pads. The ECL outputs operate from a +2 V supply. The remaining two sides of the chip provide power and bias connections for the 1 mm x 1 mm core receiver array located near the center of the chip. The core requires both +2V and +3.3V supplies.

The 200 k Ω resistor sets the reference current into the receiver circuit based on the voltage applied to the REF pad of the chip. Note that this resistor was laid out so as to fill the available space. A minimum-area 200 k Ω would occupy the same area as the 75 μm diameter MSM photodetector.

The unit cell also contains a line driver needed to deliver the DCFL (direct-coupled FET logic) output of the receiver to the ECL output drivers at the perimeter of the chip. The line driver is a “squirt-buffer”

Pad	Description
1-16	ECL Outputs
VCC	+2 V Supply to receiver circuits
VDD	+3.3 V Supply to receiver circuits
GND	Ground for receiver circuits
MSM	+5 V Bias for MSM photodetectors
REF	Bias voltage to set logic threshold
VCC(L)	+2 V Supply for line drivers
GND(L)	Ground for line drivers
VCC(D)	+2 V Supply for ECL outputs
GND(D)	Ground for ECL outputs

Table G.1: Pad definitions for MIT-OEIC-8

design and is shown in Appendix C (Figure C.7). The interconnect between the line driver and the ECL output driver has roughly 160 fF of capacitance and 30 Ω of resistance. In simulations, this line was modeled as a ten-segment RC ladder, although simply using a lumped capacitor results in only a small error in delay. The need for a line driver negatively impacts the design of the receiver array because of the large (~5 mA) current transients it generated in order to drive the capacitive load. In order to avoid coupling noise into the sensitive receivers, the line drivers must be maintained on separate power and ground connections.

This further complicates the already difficult task of routing power into the array core--a total of seven supply lines must be routed into the core including +2V and GND to the line drivers, +2V, +3.3V, and GND to the receivers, and the MSM and REF bias voltages.

Pin descriptions

GND, GND(L), and GND(D): These are the ground connections to the receivers, line drivers, and ECL outputs drivers, respectively.




Figure G.4: Bonding diagram for MIT-OEIC-8

MIT-OEIC-8 is shown wire-bonded in the package of Figure G.1. All ground connections are made as down-bonds (D.B.) to the metallic base of the package.

Pin	Function	Pin	Function
Base	GND, GND(L), GND(D)	18	VCC
1	VCC	19	VCC(L)
2	VCC(L)	20	REF
3	REF	21	VCC(D)
4	VCC(D)	22	Output 4
5	Output 13	23	Output 8
6	Output 9	24	Output 3
7	Output 14	25	Output 7
8	Output 10	26	VCC(D)
9	VCC(D)	27	Output 6
10	Output 11	28	Output 2
11	Output 15	29	Output 5
12	Output 12	30	Output 1
13	Output 16	31	VCC(D)
14	VCC(D)	32	MSM
15	MSM	33	VCC(L)
16	VCC(L)	34	VDD
17	VDD		

Table G.2: Pinout for packaged MIT-OEIC-8

VCC, VCC(L), and VCC(D): These are the +2.0 V supply to the receiver electronics, line drivers, and ECL output drivers, respectively. Note that the ECL outputs are operating from 0 V to 2 V rather than from -2 V to 0 V. The output low and high values are thus nominally 0 V and 1.2 V. With the outputs open (not terminated), the current into VCC(D) should be 300 mA. With the ECL outputs terminated to GND through 50 Ω , the current into VCC will increase by ~22 mA for each high output.

VDD: This is the +3.3 V supply to the receiver electronics.

MSM: This is the bias to the MSM photodetectors. This supply should be operated at 5.0 V. Operating at 3.3 V may be possible, but has not been carefully studied. This supply should not require a separate plane of copper, but good bypassing practices should still be followed.

REF: This is a reference voltage which sets the threshold level against which the optical input is compared to determine its logical state. The layout of this supply should follow similar guidelines as for the MSM supply. Selection of the REF bias voltage is described below.

Outputs (1-16): The sixteen pixels are numbered on the layout plot Figure G.6 (left-to-right and top-to-bottom). These are ECL outputs. They must be terminated to ground through 50 ohms at the load (e.g. they may be plugged directly into the 50 Ω input of a high speed scope or into an in-line 50 Ω termination when using a high-impedance scope input). The outputs should swing from ~0V to ~1.2V. The receivers are non-inverting, that is, the outputs should be low (~0 V) with no optical input is applied, and should be high (~1.2 V) when an optical input exceeding the threshold value is applied.

Operating procedure

This procedure assumes the IC is in a test environment where each output may be connected individually.

1) Power up VDD. Raise VDD up to 3.3 V and make sure that the current into this supply goes up to ~10 mA.

2) Power up VCC. It is assumed here that all three VCCs are ultimately connected to a common supply. The current drawn from this supply will be dominated the ECL output drivers through VCC(D). With the outputs left disconnected, raise the VCC's to 2.0 V and make sure that the current into this supply goes up to ~500 mA. The outputs may now be re-attached, i.e. terminated. Once correct operation of the OEIC has been confirmed, the circuit may be powered up with the outputs terminated. In this case, the current into VCC will increase by 22 mA for each high output.

3) Raise MSM to 5.0 V.

4) With no supply attached to REF, look at the voltage at this node with a volt meter. The voltage at this node should be 1.2-1.5 V. With no optical input applied to a photodetector, the corresponding output should be a low. On a properly terminated output, upon powering up the chip, a typical output will generate ~1 V instead of a valid low (~0 V). There appears to be a “bistability”, still unexplained, in the way the reference functions. This behavior is common to both MIT-OEIC-8 and MIT-OEIC-9, but is not observed on isolated receivers of the type used in MIT-OEIC-8. To obtain the correct output functionality, raise REF to a few volts (up to ~6.5 V). The output should snap to a logic high voltage. REF may now be lowered to the desired reference value. But, if REF reduced below ~1 V, the outputs will return to the undefined state and the procedure to set REF must be repeated.

The reliability of these OEICs, when biased above 5 V, has not been studied. Commercial parts in the H-GaAs IV technology are operated at up to 5 V.

Operating characteristics

A typical static input-output curve and eye pattern are shown in Figure G.5. With the exception of the REF bistability issue, the characteristics of receivers in the array are similar to the isolated receiver tested in Chapter 7. In that case, the detector efficiency was found to be 0.23 A/W at 850 nm. The decision level could be varied from 20 μ W to 120 μ W as REF was varied from 1 V to 6 V (reference current of the receiver was varied from 4 μ A to 27 μ A). The receiver was usable up to around 500 MB/s, and was limited by the photodetector.





Figure G.5: Typical characteristics from MIT-OEIC-8

(a) Measured time-average output voltage as a function of 850 nm input optical power for REF=4.5 V. (b) Output eye pattern. A 400 Mb/s $2^{23}-1$ bit pseudorandom bit sequence was used to modulate an 850 nm VCSEL. The reference was set to REF=5.11 V. The laser driver of Chapter 6 was used to modulate the VCSEL ($V_{off}=0.7$ V, $V_{swg}=1.5$ V). The poor ECL output waveforms are due to the packaging used in testing the array.

G.3 Receiver Array “MIT-OEIC-9”

Basic design

MIT-OEIC-9 uses the array-oriented receiver of Appendix H. While the receivers used in MIT-OEIC-8 were intended for general use in dense optical interconnect applications, the current design is geared specifically for the receiver array. Accordingly, the power consumption and area constraints were lifted. The receiver was designed to operate from a single 3.3 V supply in order to reduce supply distribution in the receiver core, and the need for line drivers was eliminated in order to reduce noise generation and further simplify supply distribution.

A layout plot of MIT-OEIC-9 is shown in Figure G.6. The pad identifiers noted in the figure are defined in Table G.3 and will be described shortly. As with MIT-OEIC-8, the ECL output drivers are located along the top and bottom of the chip and each set of eight has its own set of three supply and four ground pads. The remaining two sides of the chip provide the power and bias connections for the 1 mm x 1 mm core receiver array located near the center of the chip and for the remaining receiver circuitry located

Pad	Description
1-16	ECL Outputs
VDD	+3.3 V Supply to receiver circuits
GND	Ground for receiver circuits
MSM	+5 V Bias for MSM photodetectors
REF	Bias voltage to set logic threshold
VCC(D)	+2 V Supply for ECL outputs
GND(D)	Ground for ECL outputs

Table G.3: Pad definitions for MIT-OEIC-9

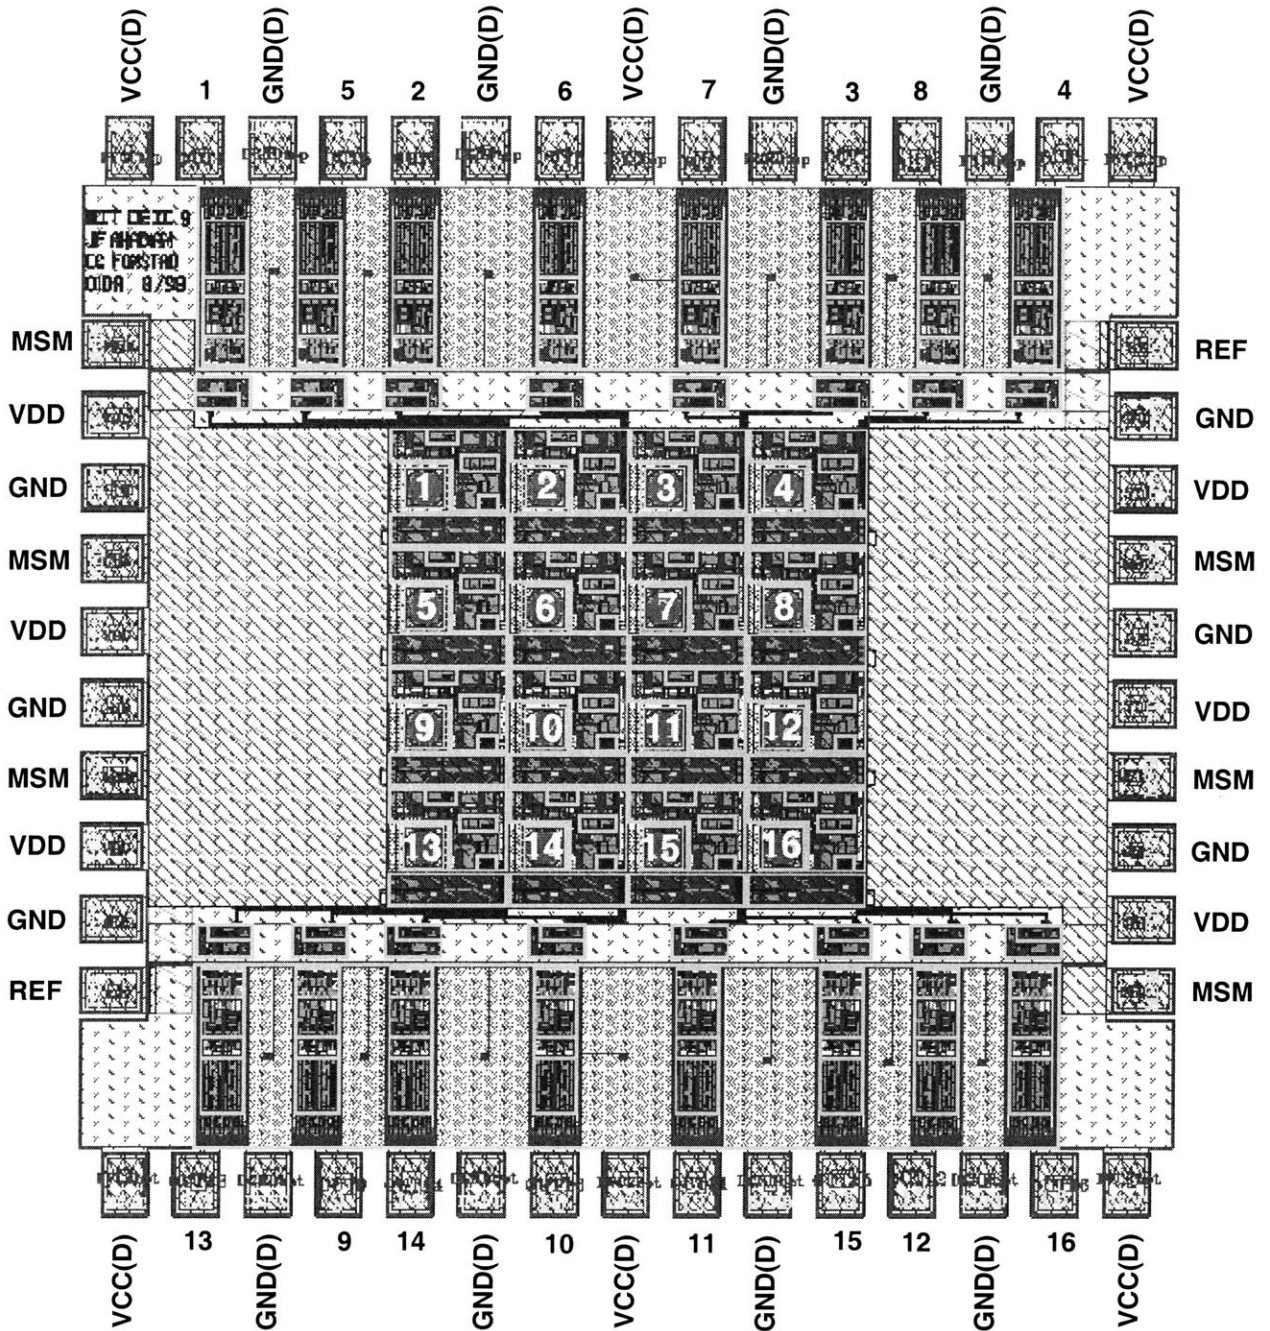


Figure G.6: Layout plot of MIT-OEIC-9

The second of two 4x4 receiver arrays fabricated using the Vitesse H-GaAs IV VLSI MESFET process. The die is 2.3 mm x 2.3 mm and provides an ECL output for each of the sixteen optical inputs. The pad names are defined in Table G.3.

adjacent to the ECL output drivers. The core uses a single +3.3 V supply in addition the MSM bias and a reference voltage. The ECL outputs require a +2 V supply.

The layout of one unit cell of the array is shown in Figure G.3. The 200 k Ω resistor sets the reference current into the receiver circuit based on the voltage applied to the REF pad of the chip. Note that this resistor was laid out to fill the available space. A minimum-sized 200 k Ω would occupy the same area as the 75 μ m diameter MSM photodetector.

A bonding diagram for MIT-OEIC-9 using the selected package, Figure G.1, is shown in Figure G.8. This bonding arrangement results in the pinout given in Table G.4.

Pin descriptions

GND and GND(D): These are the ground connections to the core and to the ECL outputs drivers, respectively.

VDD: This is the +3.3 V supply to the receiver electronics. VDD should draw around 100 mA at +3.3V.

VCC(D): This is a +2V supply to the ECL outputs along the top and bottom of the chip. Note that the ECL outputs are operating from 0V to 2 V rather than -2V to 0V. The output low and high values are thus nominally 0 V and 1.2 V. With the outputs open (not terminated), the current into VCC should be 300 mA. With the ECL outputs terminated to GND through 50 Ω , the current into VCC will increase by ~22 mA for each high output.

MSM: This is the bias to the MSM photodetectors. This supply should be operated at 5.0 V. Operating at 3.3 V may be possible, but has not been carefully studied. This supply should not require a separate plane of copper, but good bypassing practices should still be followed.

REF: This is a reference voltage which sets the threshold level against which the optical input is compared to determine its logical state. The layout of this supply should follow similar guidelines as for the MSM supply. Selection of the REF bias voltage is described below.

Outputs (1-16): The sixteen pixels are numbered on the layout plot Figure G.6 (left-to-right and top-to-bottom). These are ECL outputs. They must be terminated to ground through 50 ohms at the load (e.g. they may be plugged directly into the 50 Ω input of a high speed scope or into an in-line 50 Ω termination when using a high-impedance scope input). The outputs should swing from ~0V to ~1.2V. The receivers

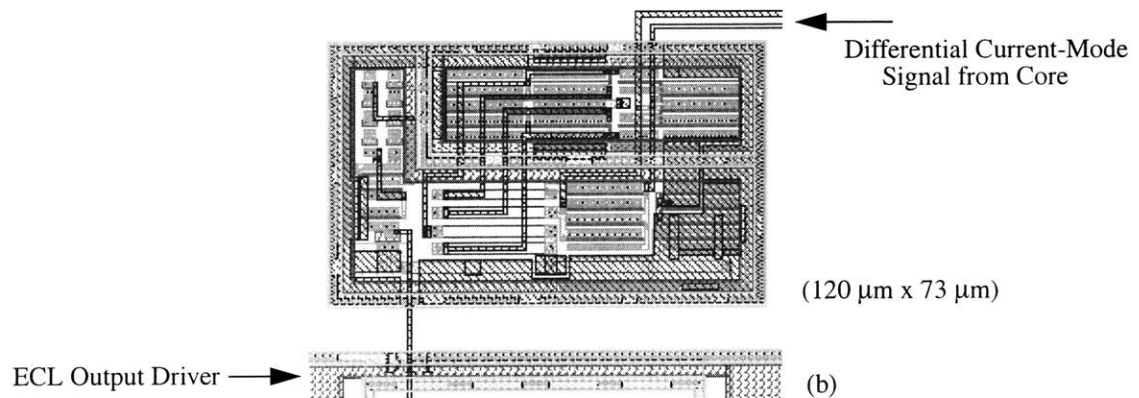



Figure G.7: Unit cell of MIT-OEIC-9 receiver array

(a) is the $250\ \mu\text{m} \times 250\ \mu\text{m}$ unit cell used in the MIT-OEIC-9 receiver array core. It includes a $75\ \mu\text{m}$ diameter MSM photodetector and a $200\ \text{k}\Omega$ resistor used to set the reference current into the receiver. The first two stages and half of the third stage of the receiver are also in the core unit cell. The latter, a transconductance amplifier, drives a differential current-mode signal across the die to the remainder of the receiver circuit, (b), located along the top and bottom of the OEIC, adjacent to the corresponding ECL output buffers.




Figure G.8: Bonding diagram for MIT-OEIC-9

MIT-OEIC-9 is shown wire-bonded in the package of Figure G.1. All ground connections are made as down-bonds (D.B.) to the metallic base of the package. N/C indicates that a pad is not connected when using this package.

Pin	Function	Pin	Function
Base	GND, GND(D)	18	MSM
1	MSM	19	VDD
2	VDD	20	REF
3	REEF	21	VCC(D)
4	VCC(D)	22	Output 4
5	Output 13	23	Output 8
6	Output 9	24	Output 3
7	Output 14	25	Output 7
8	Output 10	26	VCC(D)
9	VCC(D)	27	Output 6
10	Output 11	28	Output 2
11	Output 15	29	Output 5
12	Output 12	30	Output 1
13	Output 16	31	VCC(D)
14	VCC(D)	32	VDD
15	VDD	33	MSM
16	MSM	34	VDD
17	VDD		

Table G.4: Pinout for packaged MIT-OEIC-9

are non-inverting, that is, the outputs should be low (~ 0 V) with no optical input is applied, and should be high (~ 1.2 V) when an optical input exceeding the threshold value is applied.

Operating procedure

This procedure assumes the IC is in a test environment where each output may be connected individually.

1) Power up VDD. Raise VDD up to 3.3 V and make sure that the current into this supply goes up to 100 mA.

2) Power up VCC. With the outputs left disconnected, raise VCC to 2.0 V and make sure that the current into this supply goes up to 300 mA. The outputs may now be re-attached, i.e. terminated. Once correct operation of the OEIC has been confirmed, the circuit may be powered up with the outputs terminated. In this case, the current into VCC will increase by 22 mA for each high output.

3) Raise MSM to 5.0 V.

4) With no supply attached to REF, measure the voltage at this node with a volt meter. The voltage at this node should be 1.2-1.5 V. With no optical input applied to a photodetector, the corresponding output should be a low. On a properly terminated output, upon powering up the chip, a typical output will generate ~ 1 V instead of a valid low (~ 0 V). There appears to be a “bistability”, still unexplained, in the way the reference functions. This behavior is common to both MIT-OEIC-8 and MIT-OEIC-9, but is not observed on isolated receivers of the type used in MIT-OEIC-8. To obtain the correct output functionality, raise REF to ~ 6.5 V. Just before reaching 6.5 V, the output should drop to a low voltage. REF may now be lowered to a desired reference value. But, if REF reduced below ~ 3 V, the outputs will return to the undefined state and the procedure to set REF must be repeated.

The reliability of these OEICs, when biased above 5 V, has not been studied. Commercial parts in the H-GaAs IV technology are operated at up to 5 V.

Threshold and sensitivity

The value of REF sets the threshold optical power level, P_{th} , against which the input signal is compared to determine the output. Input powers below P_{th} produce a low output while inputs above P_{th} yield a high. A family of curves showing the ECL output as a function of the 850 nm optical input power for various REF values is shown in Figure G.9. Notice that in each curve, the output rises sharply, then rises more




Figure G.9: Static characteristics from MIT-OEIC-9

Measured time-average output voltage as a function of 850 nm input optical power for various applied REF voltages.

gradually over a range of input powers, then again abruptly completes its transition. As may be observed on an oscilloscope, within the central region of the transition the ECL output is switching rapidly between its high and low values in response to noise. The curves in Figure G.9, on the other hand, indicate the time averaged value of the output as a function of input power. The width of this transition region, ΔP , may thus be taken as an indication of the sensitivity of the receiver, while the midpoint of this region is defined as P_{th} . Note, however, that the sensitivity derived in this manner is directly applicable in the steady state, whereas larger peak-to-peak swings may be needed to produce valid logic output levels at higher speeds

and in the presence of possible inter-symbol interference effects. Also, measurements made subsequent to those of Figure G.9 have shown that these transition curves are dominated by the noise level of the optical source in use at the time. The intrinsic steady-state sensitivity of the receiver is believed to substantially better than indicated by the data here.

Table G.5 summarizes P_{th} and ΔP parameters based on the curves of Figure G.9 and gives the simultaneously measured value of I_{REF} , the current into REF. I_{REF} and P_{th} are plotted in Figure G.10. A number of observations may be made based on this set of data: With the mean value of ΔP being $22 \mu W$, the power incident on the MSM detector must exceed the specified P_{th} value by $11 \mu W$ in order to produce an output high. That is, the optical signal must have a $22 \mu W$ peak-to-peak swing centered around P_{th} . Furthermore, the efficiency of the MSM photodetectors under the applied bias conditions (MSM=5 V) may be estimated by noting the ratio of the change in I_{REF} per receiver with respect to the change in P_{th} . This yields a value of $0.15 A/W$ or a quantum efficiency of 21%. These values are in line with direct measurements made in Chapter 4. Based on this efficiency, the $22 \mu W$ sensitivity translates into a sensitivity of $3 \mu A$ at the receiver input. As in Chapter 7, this sensitivity measurement is believed to be limited by noise in the opti-

REF (V)	I_{REF} (μA)	P_{TH} (μW)	ΔP (μW)
3.0	168	90	22
3.5	204	105	20
4.0	239	119	20
4.5	275	135	17
5.0	310	148	20
5.5	344	163	20
6.0	378	177	26
6.5	412	195	27

Table G.5: Summary of parameters extracted from Figure G.9

P_{th} is the optical input power at the midpoint of the output transition. ΔP is the range of input power over which the transition occurs. I_{REF} is the current into the REF terminal.




Figure G.10: Plots of P_{th} and I_{REF} vs. the REF voltage

Data of Table G.5 are plotted. The per-receiver change in I_{REF} divided by the corresponding change in P_{th} gives the MSM photodetector efficiency, measure at 850 nm, as 0.15 A/W, or 21%. The slope of I_{REF} vs. REF gives the resistance of the parallel combination of the reference resistors in each receiver. The resulting individual resistance value is 230 k Ω .

cal signal, and is not the intrinsic static sensitivity of the receiver. Finally, the number of receivers times the slope of I_{REF} vs. REF gives the average value of the resistors used within each unit cell (Figure G.7) to set the reference current into the receiver. This value, 230 k Ω , is 15% larger than the 200 k Ω design target and is thus within the 20% resistor value tolerance of the Vitesse process.


High speed characteristics

MIT-OEIC-9 was deployed within a parallel optical link demonstration at McGill university by Maj, *et. al.* [24,25]. The link setup is shown in Figure G.11. Up to 10 outputs from an 8X8 VCSEL array provided by MODE were coupled into a fiber image guide made by Schott Optical Fibers. These signals were imaged onto 10 inputs of the MIT-OEIC-9 receiver array. The eye patterns from four out of the ten signals, operating at 250 Mb/s per channel, is shown in Figure G.12. Parallel operation of the link was limited to 250 Mb/s/ch due to crosstalk within the transmitter.

Single channel measurements were also made to further test the receiver. Transmission of a clock at 600 MHz and 1 GHz are shown in Figure G.13. Data transmission at data rates of 500, 650, 700, and 750 Mb/s is shown in Figure G.14. For reference Figure G.15, shows the input optical signal at 1 Gb/s. The fidelity of the receiver is seen to degrade above 500 Mb/s. Based on the results of Chapter 4 and Chapter 7, this limitation is believed to be due to an excessive MSM photodetector capacitance. Alteration of the detector structure may allow operation of the receiver at up to its designed speed of 1 Gb/s.

Power dissipation


The sixteen receiver circuits draw a total of 100 mA at 3.3 V leading to a power dissipation of 330 mW. With the outputs disconnected, the ECL drivers draw 300 mA at 2 V. This adds 600 mW to the on chip power dissipation. Each terminated high-state output draws an additional 22 mA. In the worst case, this equals a total of 352 mA across 0.8 V of on-chip voltage drop (the remaining 1.2 V is dropped external to the chip). This adds 282 mW of on-chip power dissipation in the worst case. The total worst-case on-chip power dissipation is thus 1.2 W.



MODE VCSEL Array

Fiber Image Guide


Length: 1.35 m
Pixel count: 17000



MIT-OEIC-9 Receiver Array

Figure G.11: Optical link demonstration using MIT-OEIC-9

In this demonstration carried out at McGill university [24,25], ten elements of an 8x8 VCSEL array from MODE are coupled into a fiber image guide from Schott Optical Fibers. The transmitted optical spots are imaged onto an MIT-OEIC-9 receiver array and the ECL outputs were observed.



1 ns/Div

Figure G.12: Parallel data transmission using MIT-OEIC-9

Eye patterns from four out of the ten optical channels being exercised at 250 Mb/s/channel.

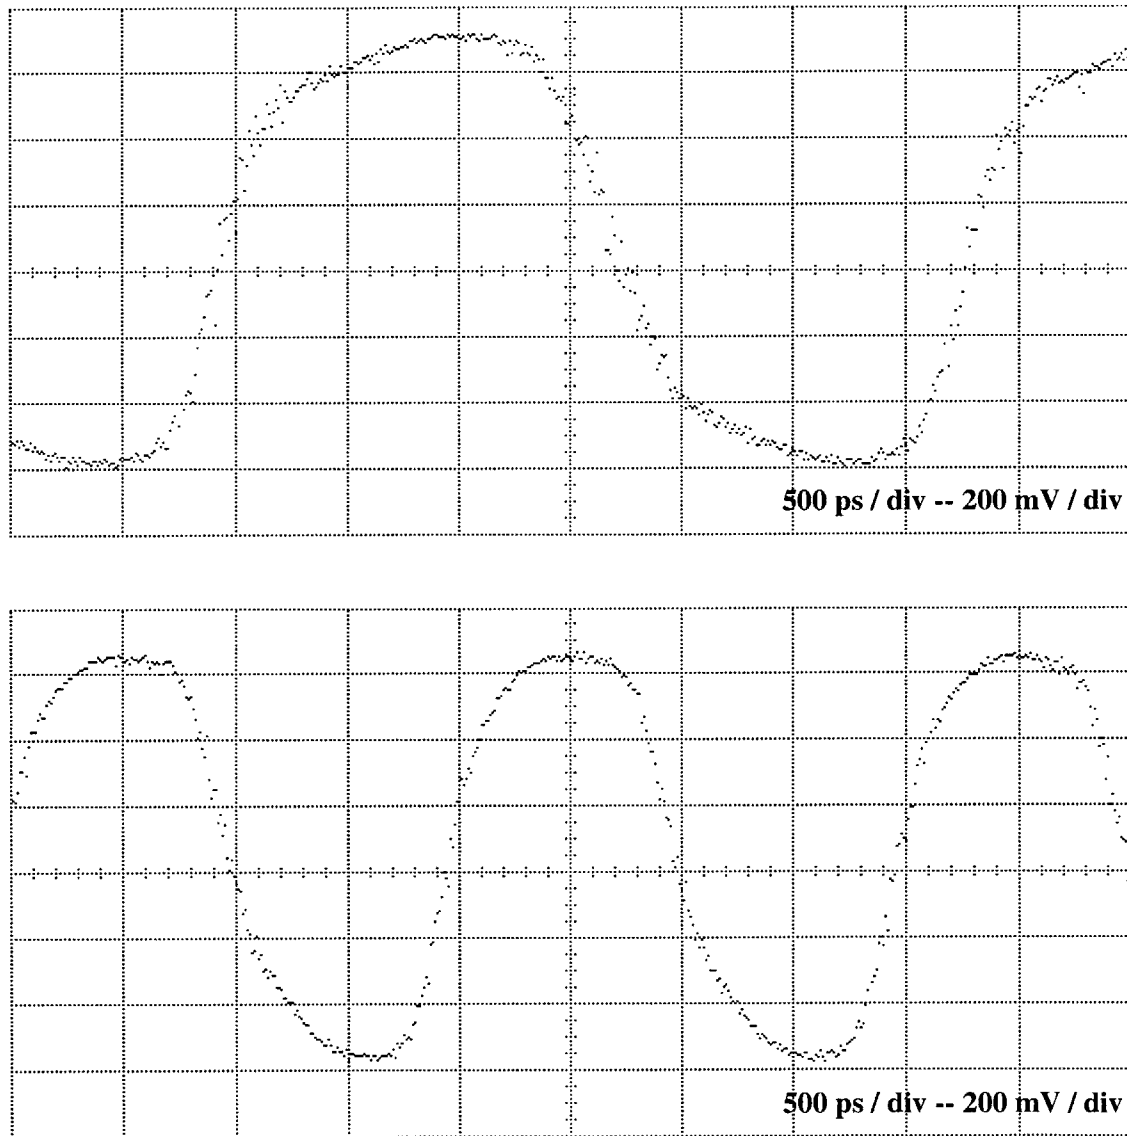


Figure G.13: Transmission of clock signals using MIT-OEIC-9

Output of a single receiver on MIT-OEIC-9 in response input optical clock signals at 500 MHz and 1 GHz.

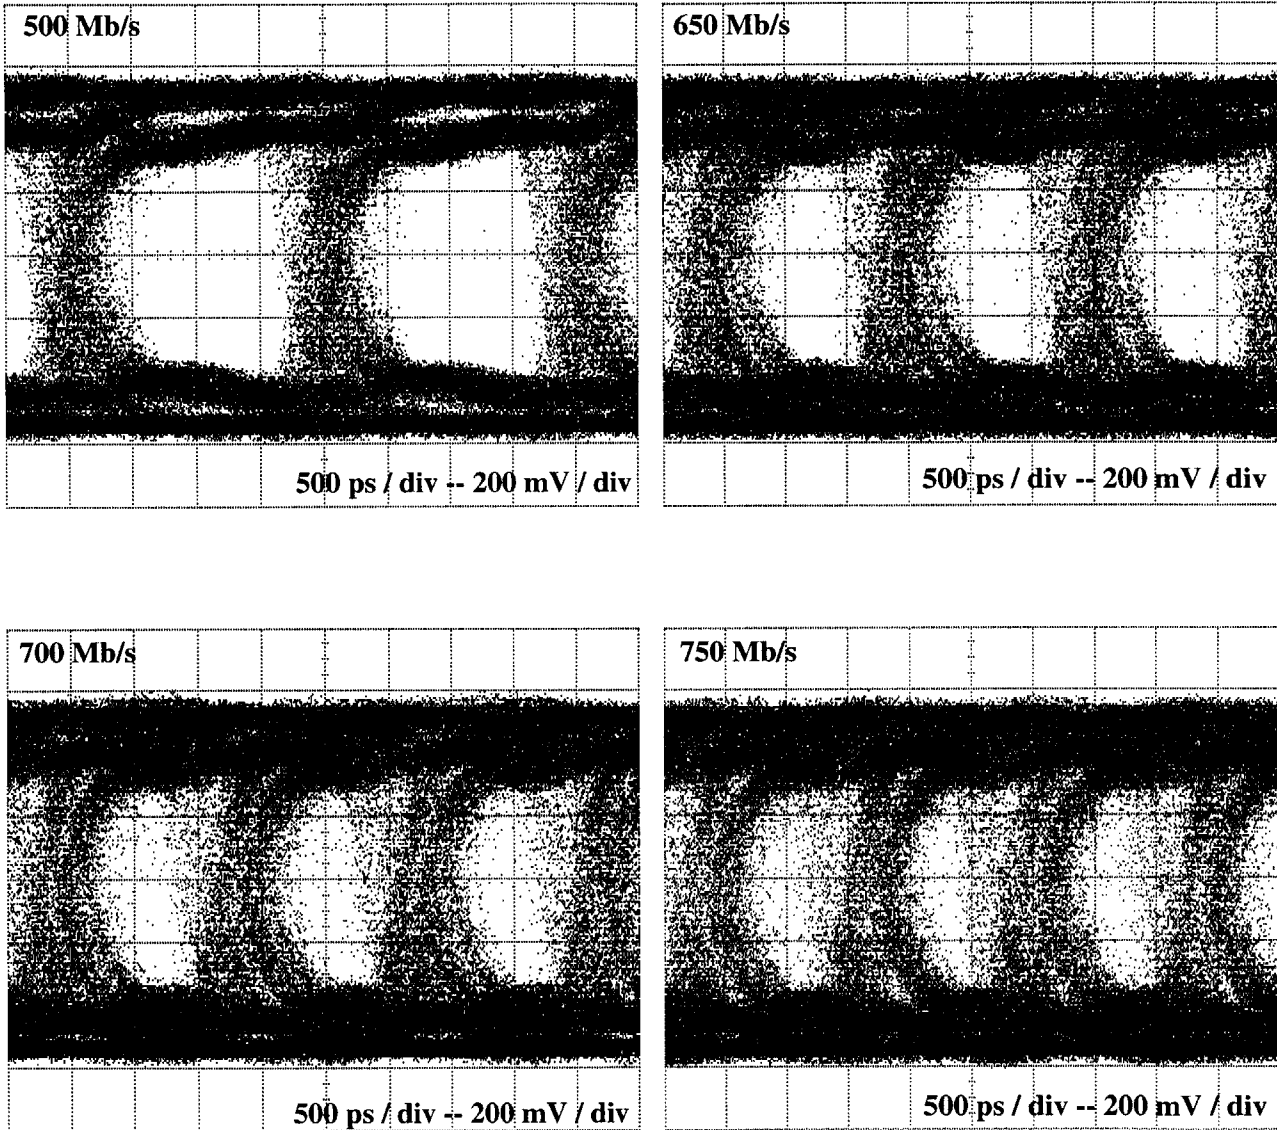


Figure G.14: Single-channel data transmission using MIT-OEIC-9
Eye patterns from a single receiver on MIT-OEIC-9 operated at 500, 650, 700, and 750 Mb/s.




Figure G.15: Eye pattern at 1 Gb/s from optical source used to test MIT-OEIC-9 at McGill University. The receiver output eye patterns shown in Figure G.14 were produced using this optical source.

H

A Fully-Differential Array Receiver

The general principles applied in the design of the low power digital optical interconnect receiver used in the MIT-OEIC-8 array chip were reviewed in Chapter 7. The same principles were applied in the design of the receiver used in the MIT-OEIC-9 array chip. Appendix G discusses the array chip, the physical design of the receiver, and its measured performance. This appendix examines the design of the receiver circuit.

While the receiver design of Chapter 7 was intended for general use in dense optical interconnect applications, the current design is geared specifically for the receiver array. Accordingly, the power consumption and area constraints were relaxed. The typical simulate power dissipation of this circuit is 23 mW (the Chapter 7 design dissipated 4.5 mW). The receiver was designed to operate from a single 3.3 V supply in order to simplify supply distribution in the receiver core. The need for separate line drivers was eliminated in order to reduce noise generation and further simplify supply distribution.

The receiver circuit is shown in block form in Figure H.1 and schematically in Figure H.2, Figure H.3, and Figure H.4. To maximize supply noise immunity, this receiver uses a fully-differential design. The first-stage transimpedance amplifier and second-stage voltage amplifier are followed by cascaded transconductance and a transimpedance amplifiers. This cascade combination is known as a Cherry-Hooper amplifier [258], but in the implementation used here, it is distributed between the receiver array core and the OEIC periphery in order to eliminate the need for a line driver.

The interconnects joining the receiver core the periphery are implemented in minimum-width metal-2, are around 1 mm in length, and have roughly 160 fF of capacitance and 30 Ω of series resistance. Compar-

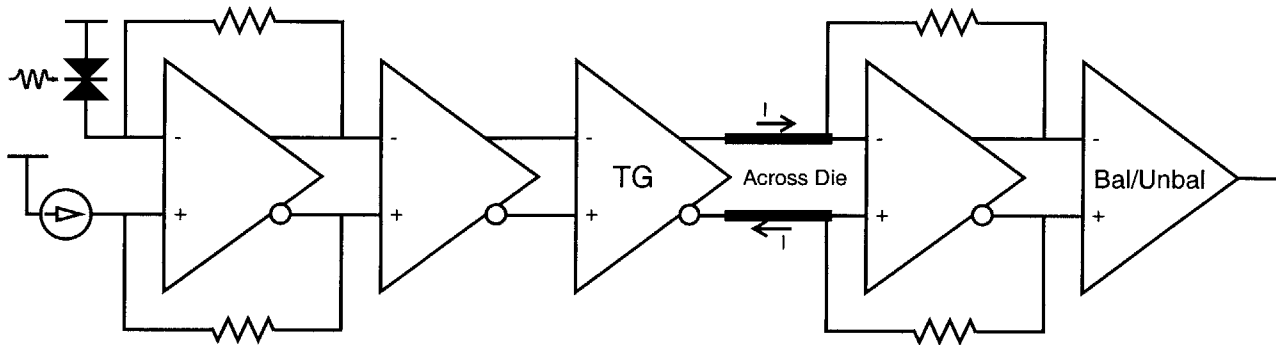


Figure H.1: Block representation of fully-differential receiver

Block representation of the fully differential optical receiver detailed in Figure H.2, Figure H.3, and Figure H.4. The first-stage transimpedance amplifier and second-stage voltage amplifier are followed by the cascade of a transconductance and a transimpedance amplifier. This cascade combination is known as a Cherry-Hooper amplifier [258], but in the implementation used here, it is distributed between the receiver array core and the OEIC periphery in order to eliminate the need for a line driver.

Comparison of simulations in which the interconnect was modeled by a ten stage R-C ladder with those in which only a lumped capacitance was used indicate that the effects of the capacitance are dominant. A current mode signal is used in the across-die interconnect. By reducing the impedance of the node on which the interconnect capacitance appears, the transconductance-transimpedance cascade reduces the time constant on this node and thus the effect of the capacitance.

H.1 Simulation Results

Static simulations

A simulation of the static characteristics of the receiver is shown in Figure H.5. An important feature to observe in this plot is the clamping on nodes 400 and 401. These are the input nodes of the transimpedance amplifier used in the distributed Cherry-Hooper amplifier. V_{400} and V_{401} are clamped relative to each other by the diode-connected EFETs E432 and E433. It is important that these signals be hard limited before the outputs of the TIA, V_{500} and V_{501} , reach their hard limit (which is set by the upper supply). If the TIA were permitted to hard limit, its input impedance would become large since the voltage amplifier would no longer provide gain. This would cause the amplitude of V_{400} and V_{401} and the time constant on

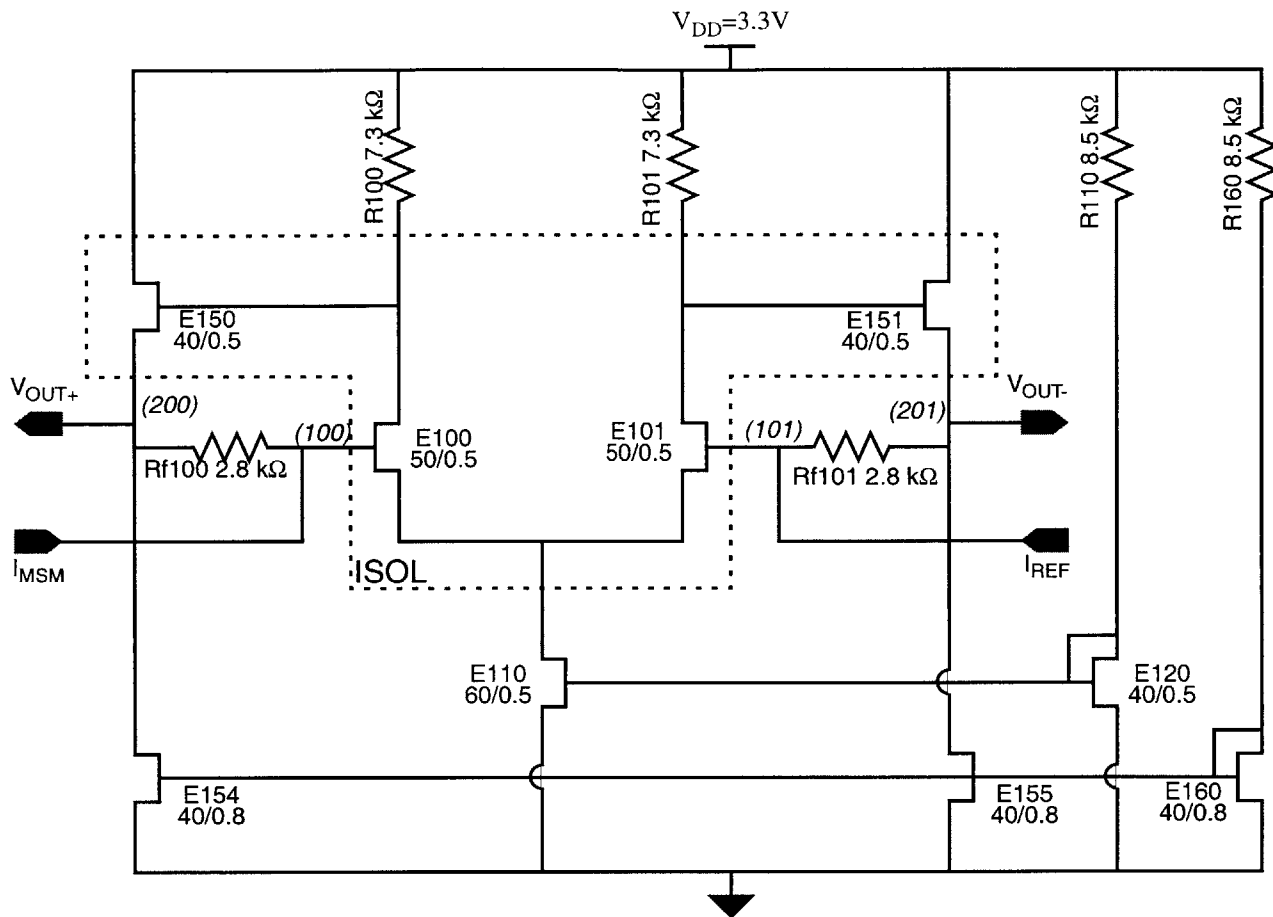


Figure H.2: Schematic of fully-differential receiver; first stage TIA

First stage of the fully differential array receiver used in MIT-OEIC-9. This stage is configured as a transimpedance amplifier.

nodes 400 and 401 to become large, and recovery to the desired mode of operation (small amplitudes and time constants on nodes 400 and 401) would be slow.

Small-signal AC simulations

A small signal frequency response simulation is shown in Figure H.6. There are a number of notable features, including a small amount of peaking on the TIA response (V_{200} and V_{201}). Although a single gain stage is used to construct the TIA, the phase lag through the source follower reduces the phase margin to the point where this peaking is observed. This level of peaking is not problematic, and in fact, it is beneficial to the receiver performance since it reduces the TIA output rise/fall time and delay.

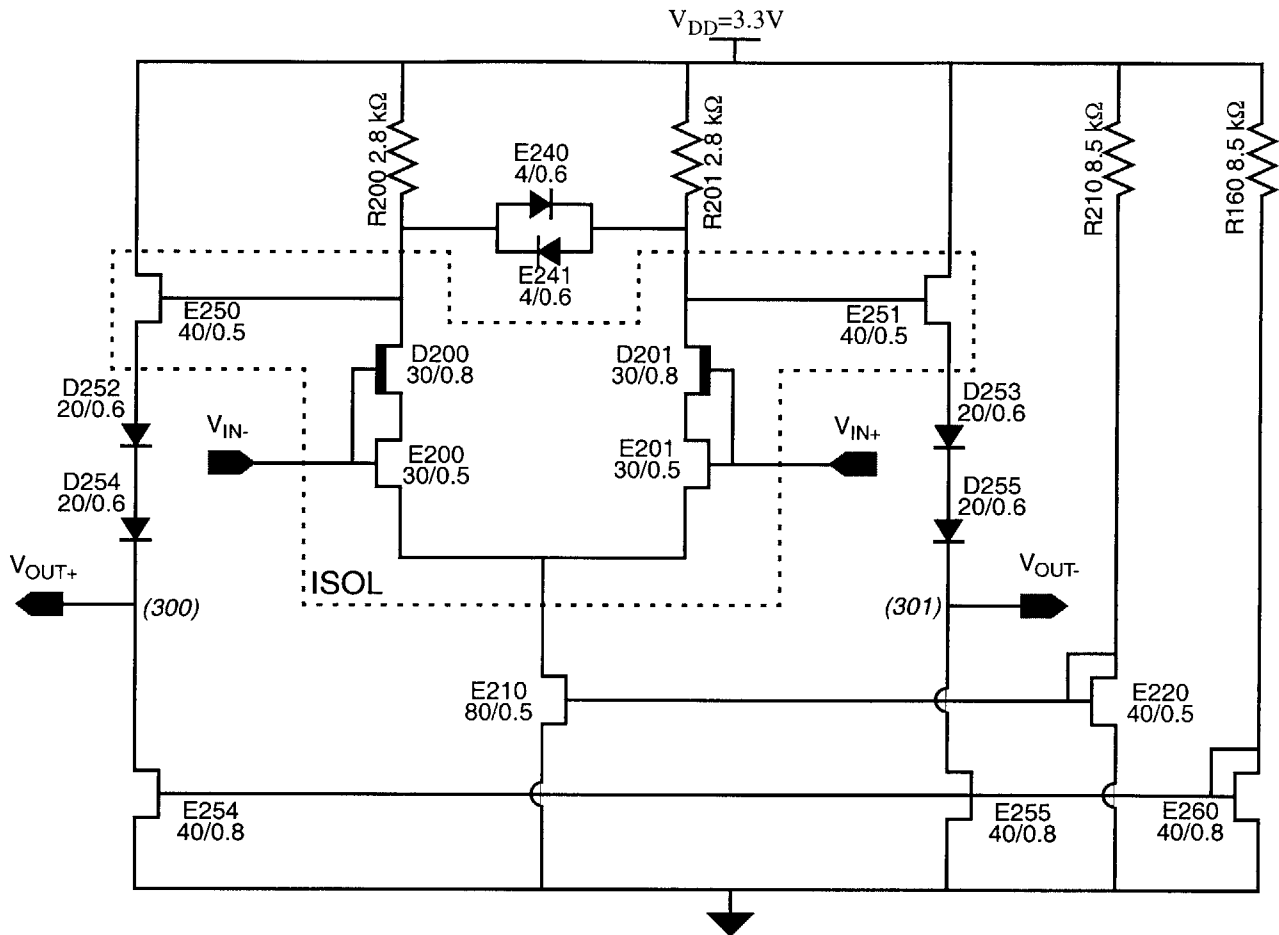


Figure H.3: Schematic of fully-differential receiver; second stage voltage amplifier

Second stage of the fully differential array receiver used in MIT-OEIC-9. The output of this voltage gain stage is clamped by the EFET diodes E240 and E241 to prevent a large input from taking D200/D201 out of saturation.

As explained in Appendix E, the use of feedback in the TIA controls the effects of drain lag on V_{200} and V_{201} . To reduce drain lag effects on V_{300} and V_{301} , the second stage amplifier uses an active cascode topology. The high frequency gain on these nodes is around 1 dB lower than the low frequency value.

If the TIA in the Cherry-Hooper cascade had an ideal, zero input impedance, there would be no voltage response on nodes 400 and 401. Because of the low voltage gain of the last stage, however, the TIA input impedance is not low enough to achieve this ideal. The voltage gain from nodes 300/301 to nodes 400/401 is, however, fairly small (~ 5 dB). The low voltage gain leads to relatively little drain lag effect in the voltage response on nodes 400 and 401. However, the last stage voltage amplifier, although it includes active

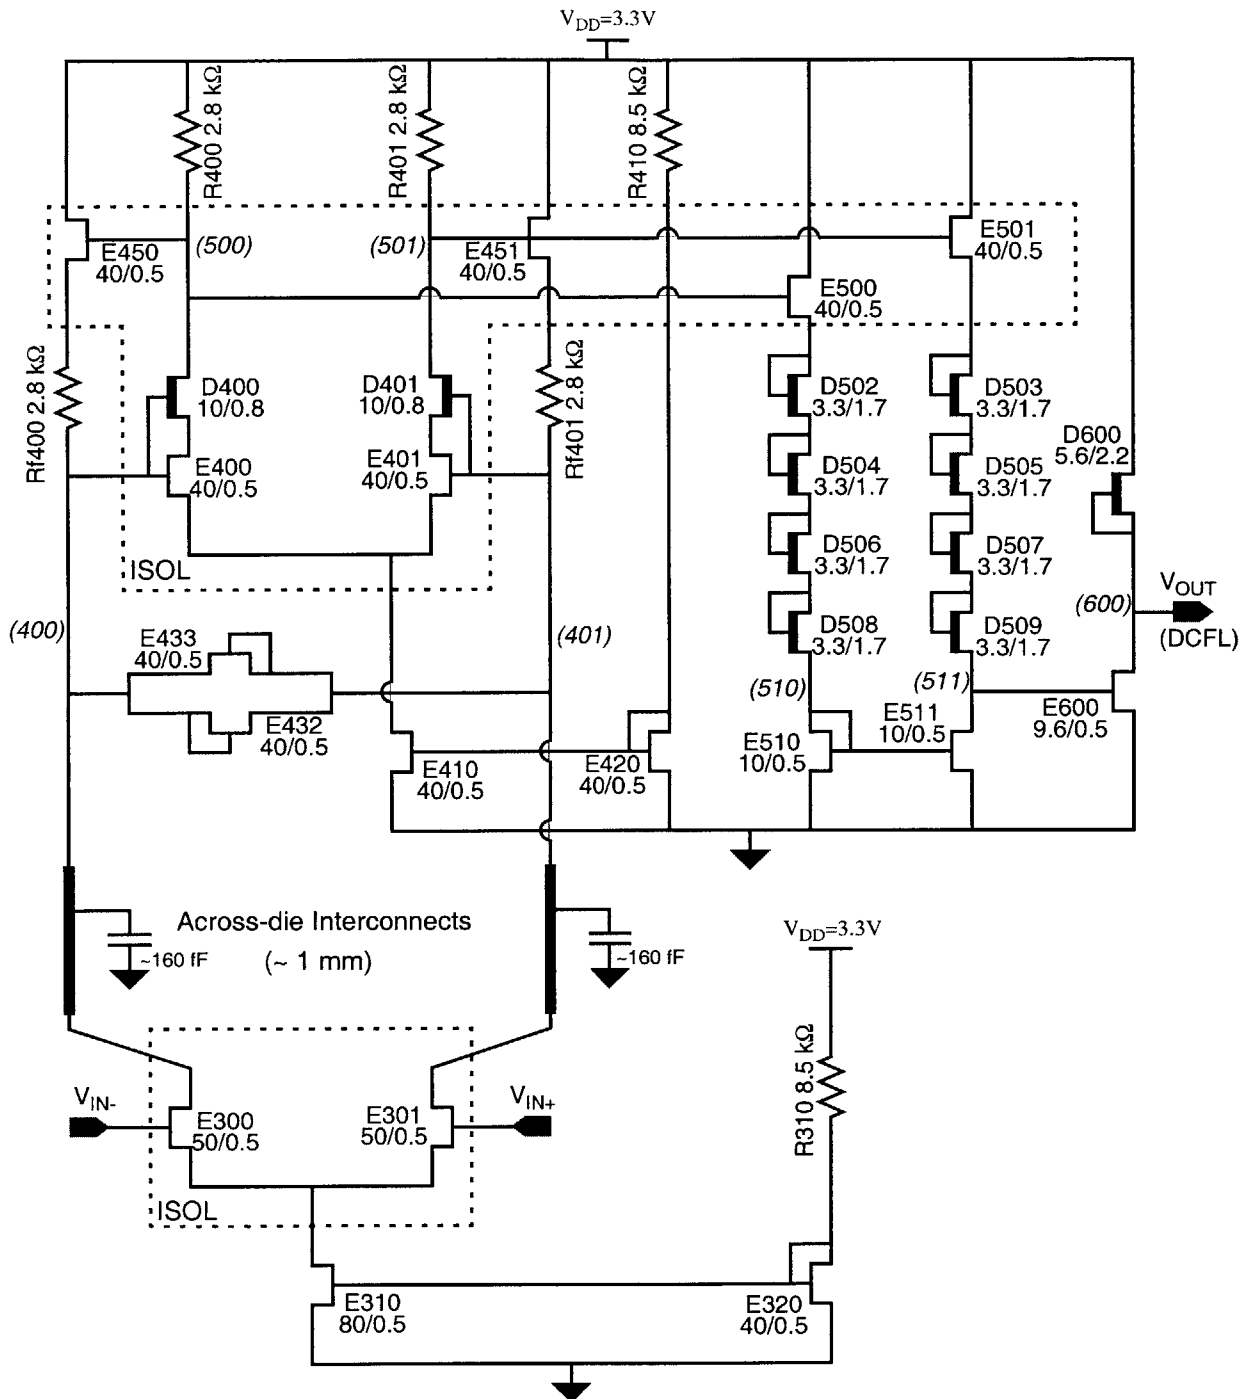


Figure H.4: Schematic of fully-differential receiver; final stages

The transconductance stage in the lower portion of the schematic drives a differential current mode signal across the die into the low impedance inputs of the transimpedance amplified in the upper portion of the schematic. The balanced outputs of this stage are then converted to a DCFL logic signal.




Figure H.5: Simulation of fully-differential receiver: static characteristics

The MSM and REF input currents were swept differentially. The output was loaded by a DCFL inverter.




Figure H.6: Simulation of fully-differential receiver: frequency response

The MSM and REF inputs were loaded by 100 fF capacitors and small-signal AC current sources of opposite polarity were applied to each. The magnitude of the voltage response on the indicated nodes are plotted above in dBΩ.

cascode structure, does not adequately control drain lag. The high frequency gain on nodes 500 and 501 is around 3 dB lower than the low frequency gain. As a TIA, the last stage drain lag effect would be small if its input was an ideal current source. But, the gain variations in the last stage appear as input impedance variations. The current-mode signal arriving at the last stage thus suffers from drain lag rather than the voltage signal. More careful design of the cascode stages on the last stage amplifier would be effective in reducing the drain lag seen in Figure H.6. However, the transient simulation shown below indicates that the drain lag effect is already adequately controlled.

Transient simulations

The results of a transient simulation of the receiver are plotted in Figure H.7. In this simulation, the MSM and REF inputs were loaded by 100 fF capacitors and current sources of opposite polarity switching between plus and minus 6 μA in 200 ps were applied to each. The output was loaded by a DCFL inverter. High and low pulse durations in this simulation were set to 50 ns in order to allow drain lag effects to be observed. An extended response tail associated with drain lag appears on V_{300} and V_{301} and may result in some amount of ISI-related jitter.

The input amplitude in the simulation of Figure H.7, which is equivalent to a 12 μA peak-to-peak single-ended input, has been selected to strongly clamp V_{400} and V_{401} . An extended response tail is observed in these signals, however, it is the common-mode component of these signals which exhibits this response. The differential mode signal, which are of primary importance to the following differential stage, does not show a significant extended response tail at this level of clamping and beyond. For smaller amplitudes, the differential-mode signal also exhibits a long response tail. The long response may be carried over partly from V_{300} and V_{301} , but also depends on the symmetric nature of the clamping on nodes 400 and 401 (see Figure H.5). The important practical observation is that for single-ended input amplitudes of at least 12 μA , no significant drain-lag effect is expected on nodes 400 and 401 or on nodes 500 and 501.

Hard limiting of the signal on node 511 is also used to avoid drain lag problems. The input amplitude needed to adequately hard-limit nodes 400 and 401 exceeds that needed to hard-limit node 511.

A transient simulation carried out on a shorter time scale is reported in Figure H.8. Here the input waveform features 2 ns high and low durations. All rise/fall times are below 500 ps, and the delay from input to output is around 600 ps.

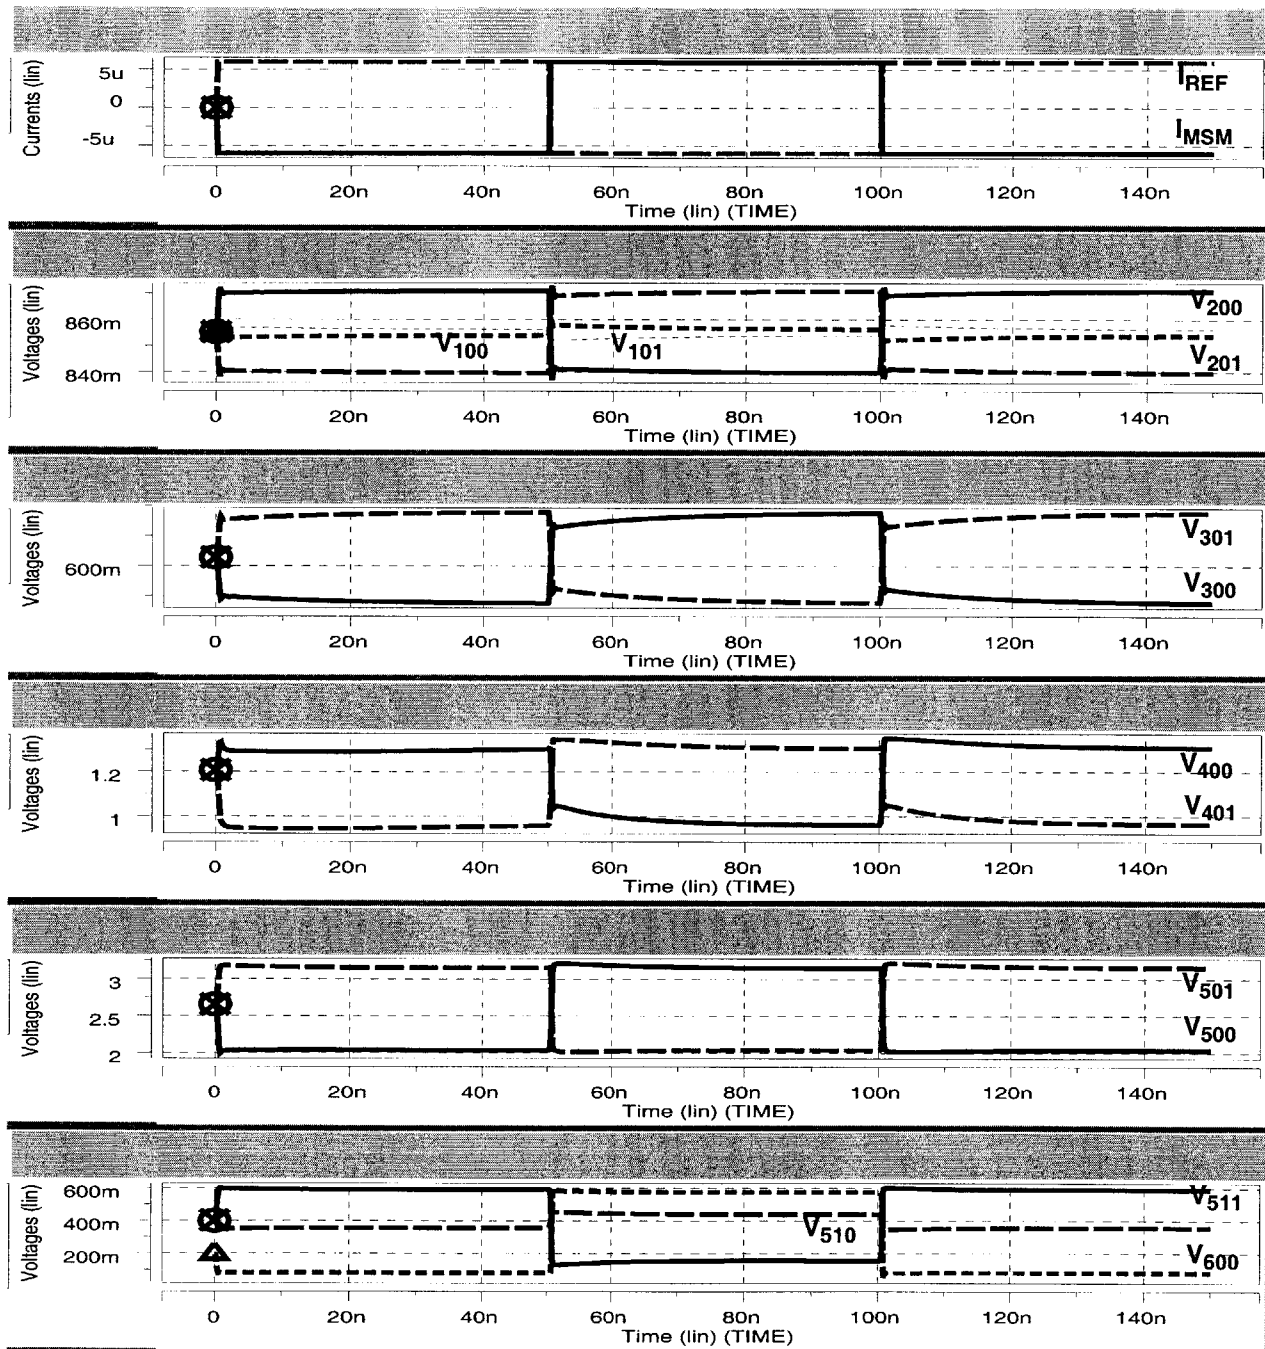


Figure H.7: Simulation of fully-differential receiver: long duration transient response

The MSM and REF inputs were loaded by 100 fF capacitors and current sources of opposite polarity switching between plus and minus 6 μ A in 200 ps were applied to each. The output was loaded by a DCFL inverter. The initial 20 ns of the simulation does not represent a valid input condition and was constructed for simulation convenience.




Figure H.8: Simulation of fully-differential receiver: short duration transient response

The MSM and REF inputs were loaded by 100 fF capacitors and current sources of opposite polarity switching between plus and minus $6\mu\text{A}$ in 200 ps were applied to each. The output was loaded by a DCFL inverter. The initial 1 ns of the simulation does not represent a valid input condition and was constructed for simulation convenience.

Supply noise simulations

The simulated frequency response of various nodes of the receiver to small-signal variation in the 3.3 V supply are shown in Figure H.9. The use of a differential topology is effective in preventing the noise that is coupled into the early stages from being amplified and fed forward as would be expected in a single-ended design. The noise reaching node 511, which is the input to the final comparator (DCFL inverter) of this receiver, originates mainly at nodes 500 and 501. These noise signals appear symmetrically on nodes 500 and 501 and is largely rejected by the differential-to-single-ended conversion circuit.

It is worthwhile comparing the supply response of nodes 500, 501, and 511 of this receiver with the corresponding nodes 400, 401, and 411 of the receiver presented in Chapter 7. In the former case, the last gain stage biasing circuit was designed to maintain constant bias voltages in spite of supply voltage variations. This was accomplished by matching the biasing resistor on the current mirror with the differential-pair load resistors and using a diode chain in series with the biasing resistor to set the bias voltage. This was successful in reducing the response on nodes 400 and 401 to -9 dB (see Figure 7.17) below 100 kHz, but above this frequency the biasing circuit did not track the supply variations and the response on nodes 400 and 401 rose to -1.5 dB. By contrast, the present circuit does not use a matched biasing resistor or a diode chain. As a result, response on nodes 500 and 501 is largely flat at -2 dB.

On the other hand, the differential-to-single-ended conversion circuit of the present receiver design appears to be superior in terms of its common-mode signal rejection. It suppresses the response on node 511 to around 20 dB below that on nodes 500 and 501 while the corresponding circuit of the Chapter 7 design provides around 10 dB less rejection. Thus, the receiver shown here has better supply rejection on node 511 even though more noise appears on nodes 500 and 501. The primary difference in the two conversion circuits is the use of EFET source followers in the present circuit rather than DFETs in the Chapter 7 design. The DFETs give better matching of the V_{411} bias voltage with the DCFL trip point and thus help to reduce static delay variation, but the present circuit makes up for reduced bias-point tracking with higher gain and sharper signal transitions.

The simulated transient response of the receiver to a step variation in the supply is shown in Figure H.10. A 100 mV step in the supply, switched in 10 ps, is seen to result in around 14 mV transients in V_{511} . This is around one-half the corresponding result of the Chapter 7 receiver which responded with 30 mV transients under the same test conditions.




Figure H.9: Simulation of fully-differential receiver: supply frequency response

A small signal AC voltage source was superimposed on the 3.3 V supply. The MSM and REF inputs were loaded by 100 fF capacitors. The magnitude of the voltage response on various nodes is plotted above.




Figure H.10: Simulation of fully-differential receiver: supply transient response

The supply voltage was stepped between 3.3 V and 3.4 V in 10 ps. The MSM and REF inputs were loaded by 100 fF capacitors. The output was loaded by a DCFL inverter.

H.2 Conclusion

This appendix has presented the design and simulation results of a fully-differential digital optical receiver used in the MIT-OEIC-9 receiver array OEIC. Experimental results are found in Appendix G. This receiver dissipates 23 mW, as compared to 4.5 mW in the case of the receiver presented in Chapter 7, while the performance of the two designs is largely similar. Both are designed for operation at 1 Gb/s, and the present design has a dynamic sensitivity, limited by hard-limiting requirements, of 12 μA peak-to-peak as compared to 16 μA for that in Chapter 7. A key feature of this design, however, is the use of a distributed Cherry-Hooper amplifier to eliminate the need for line drivers within the receiver array core. The larger high-frequency gain of this receiver, 112 dB Ω rather than 92 dB Ω , translates into sharper transition on hard-limited signal nodes and thus leads to reduced static delay variations. This fact, in turn, allows the use of a differential-to-single-ended conversion circuit with improved common-mode signal and supply rejection. However, the bulk of the power increase in the present case is due to the fully-differential topology which, by itself, does not provide a significant improvement over the use of symmetric single-ended stages in the Chapter 7 design. Future receiver designs may benefit by incorporating elements from each of these circuits.


I

A Clocked Optical Receiver

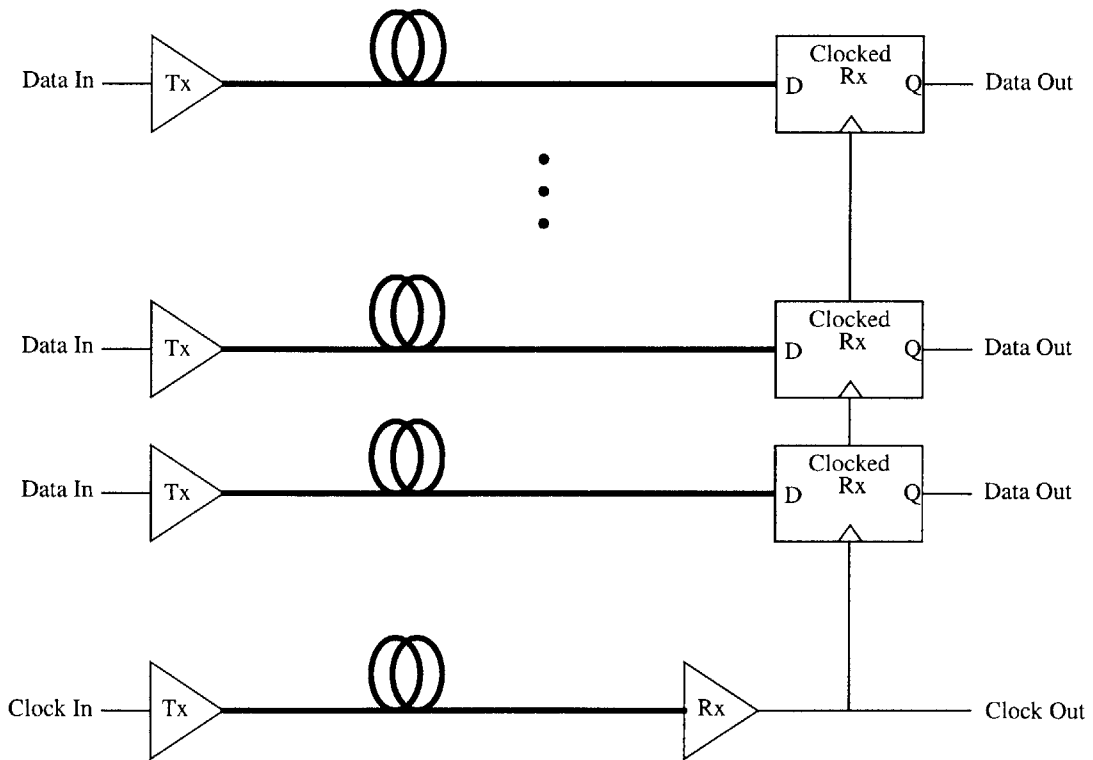
The digital optical interconnect receiver analyzed in Chapter 5, Chapter 7, and Appendix H is a continuous-time receiver. These receivers produce a binary-valued output which reflects the instantaneous value of their input. This appendix looks at a discrete-time receiver with a binary-valued output that reflects the state of its input at the time of the last active clock edge applied to it. That is, this “clocked optical receiver” has the functionality of an edge-triggered D-type flip-flop. At the heart of the clocked receiver is a sense-amplifier which uses positive-feedback to generate a logic output based on a small input. The use of positive feedback allows a compact, low-power circuit to replace the multiple broadband amplifier stages in a continuous-time receiver. The small size and power consumption of the clocked optical receiver is its major advantage relative to the continuous-time receiver.

Since it requires a clock input, the clocked optical receiver is not a general-purpose replacement for the continuous-time receiver. However, it may be used in applications where a clock is available. This may be a local clock or it may be transmitted optically with the use of a continuous-time receiver. Figure I.1 compares a digital optical interconnect systems using only continuous-time receivers with one that uses both clocked and continuous-time receivers.

The clocked optical receiver is conceptually equivalent to the receiver in an optical communication system. Thus, the design framework examined in Chapter 5 for the continuous-time digital optical interconnect receiver is not directly relevant here. Rather, the sensitivity and ISI considerations which go into the design of communication receivers may be applied [194-198], with the understanding that an error probability in the 10^{-15} - 10^{-20} range is required. Furthermore, the clocked optical receiver in a digital opti-



(a) Digital optical interconnect system using all continuous-time receivers



(a) Digital optical interconnect system using both continuous-time and clocked receivers

Figure I.1: Optical bus using continuous time and clocked receivers

An digital optical interconnect system may be assembled using only continuous time receivers, as shown in (a), or it may use a combination of continuous-time and clocked receivers, as in (b).

cal interconnect differ significantly from telecommunication receivers in that they must operate in a noisy VLSI environment and must be very compact and power-efficient.

Clocked optical receivers for use in free-space optical interconnects have been implemented using CMOS electronics [222,259,260]. These circuit topologies are not readily translated into the enhancement/depletion MESFET process. Standard NMOS sense amplifiers are also difficult to implement with MESFETs because of the use of pass transistors [261].

I.1 Design and Simulation

A clocked optical receiver designed for compatibility with direct-coupled FET logic (DCFL) MESFET electronics is shown symbolically in Figure I.2 and schematically in Figure I.3. The heart of this circuit is a pair of cross-connected DCFL NOR gates formed by E100-E104 and D100/D101. Unlike standard NOR gates, however, the gates of the load devices, D100/D101, are driven by a matched pair of transimpedance amplifiers. When CLK is high the outputs of both NOR gates are pulled low. When CLK falls, the state




Figure I.2: Block representation of clocked optical receiver

Low-level voltage signals are generated by compact transimpedance amplifiers at the front end. A clocked sense-amplifier formed from drive-modulated cross-connected NOR gates makes a decision on the low-level signals. Its outputs are latched by an additional flip-flop. The overall functionality is that of a negative-edge triggered D-type flip-flop having an optical input.

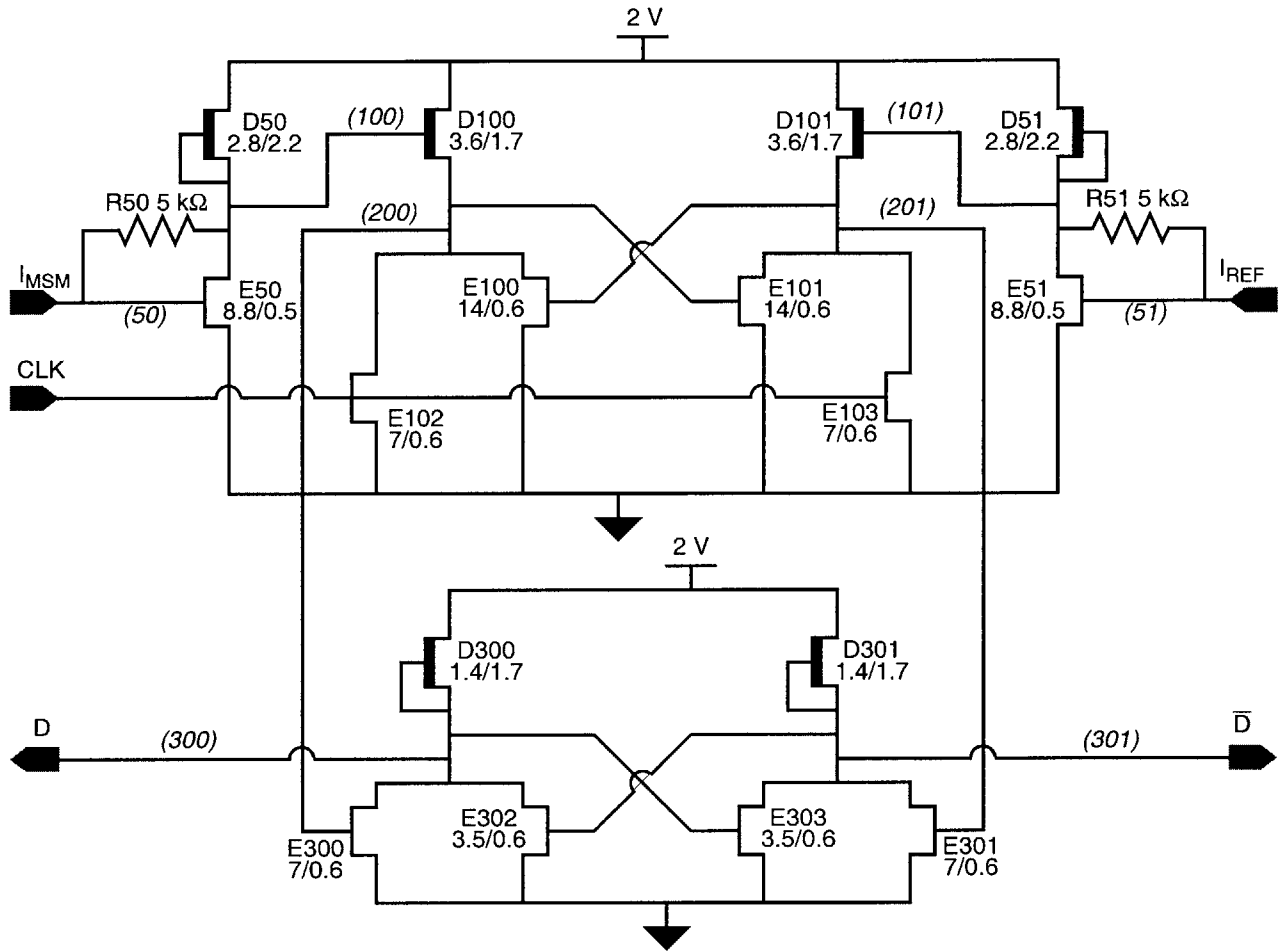


Figure I.3: Schematic of clocked optical receiver

MSM and reference currents are input to a matched pair of simple transimpedance amplifiers which drive the DFET loads in a pair of cross-connected NOR gates. When CLK is switched low this latch settles according to the relative strength of the loads. A second NOR latch holds this decision for the remainder of the clock cycle. The values indicated in parentheses are node numbers referred to in Figure I.4.

into which this latch settles is determined by the relative strength of the load devices, and is thus controlled by the relative level of the current input to the two transimpedance amplifiers. The outputs of the sense amp are applied to a second cross-connected NOR pair, E300-E304 and D300/D301. This latch holds the sense amp output when CLK returns high. The overall circuit thus performed the function of an optical-input, negative-edge-triggered D-flip-flop.

The transimpedance amplifiers (TIAs) in this circuit are implemented using DCFL inverters as single-ended voltage amplifiers. The bias point of the TIAs are thus expected to track those to the DCFL-based sense-amplifier thereby reducing process and temperature dependent variability in the circuit performance. Because of the symmetry of the TIAs and sense amplifier, supply noise, which appears as a common-mode signal on the two halves of the circuit, should be rejected to a large extent.

A simulation of this circuit, including interconnect capacitances, source/drain resistances, and source/drain junction diodes/capacitances, is shown in Figure I.4. A assumed detector capacitance of 100 fF is included at the inputs, and the outputs are loaded by DCFL inverters with a total fanout of four. In this simulation, the reference current, I_{REF} , is set to 10 μ A and the input current, which is switched in 100 ps, I_{MSM} , is modulated by 10 μ A centered around the 10 μ A reference level. The data logic level is reversed every 0.9 ns while a 1 GHz clock is applied. In this way, the effect of various set-up and hold time combinations may be observed. The clock-to-Q delay in this simulation, measured from the midpoint of the falling clock edge to the midpoint of the falling output edge, varies between 330 ps and 400 ps. The rising output edge is typically around 100 ps later. The circuit does not appear to have a significant hold-time requirement. This is seen in the event at 2.2 ns in which the falling edge is nearly coincident with the data edge and the proceeding data value is clocked in. Following this event, the receiver fails to sense the correct value due to set-up time violations. This continues up to the clock event at 6.2 ns. The set-up time is around 300 ps at this point, but the output falling edge is distorted, indicating this set-up time is marginal. In the following clock event, a 400 ps set-up is seen to be quite effective.

Notice in Figure I.4 that the TIA output and input are significantly distorted due to capacitive coupling between nodes 100 and 200. The receiver is functional despite this distortion because it appears as a common-mode signal. This distortion may, however, be contributing to the set-up time. The distortion may be largely eliminated by placing a source follower between the TIA and the sense amplifier.

The clocked sense amplifier simulated here dissipates 1.8 mW. A layout plot is shown in Figure I.5.

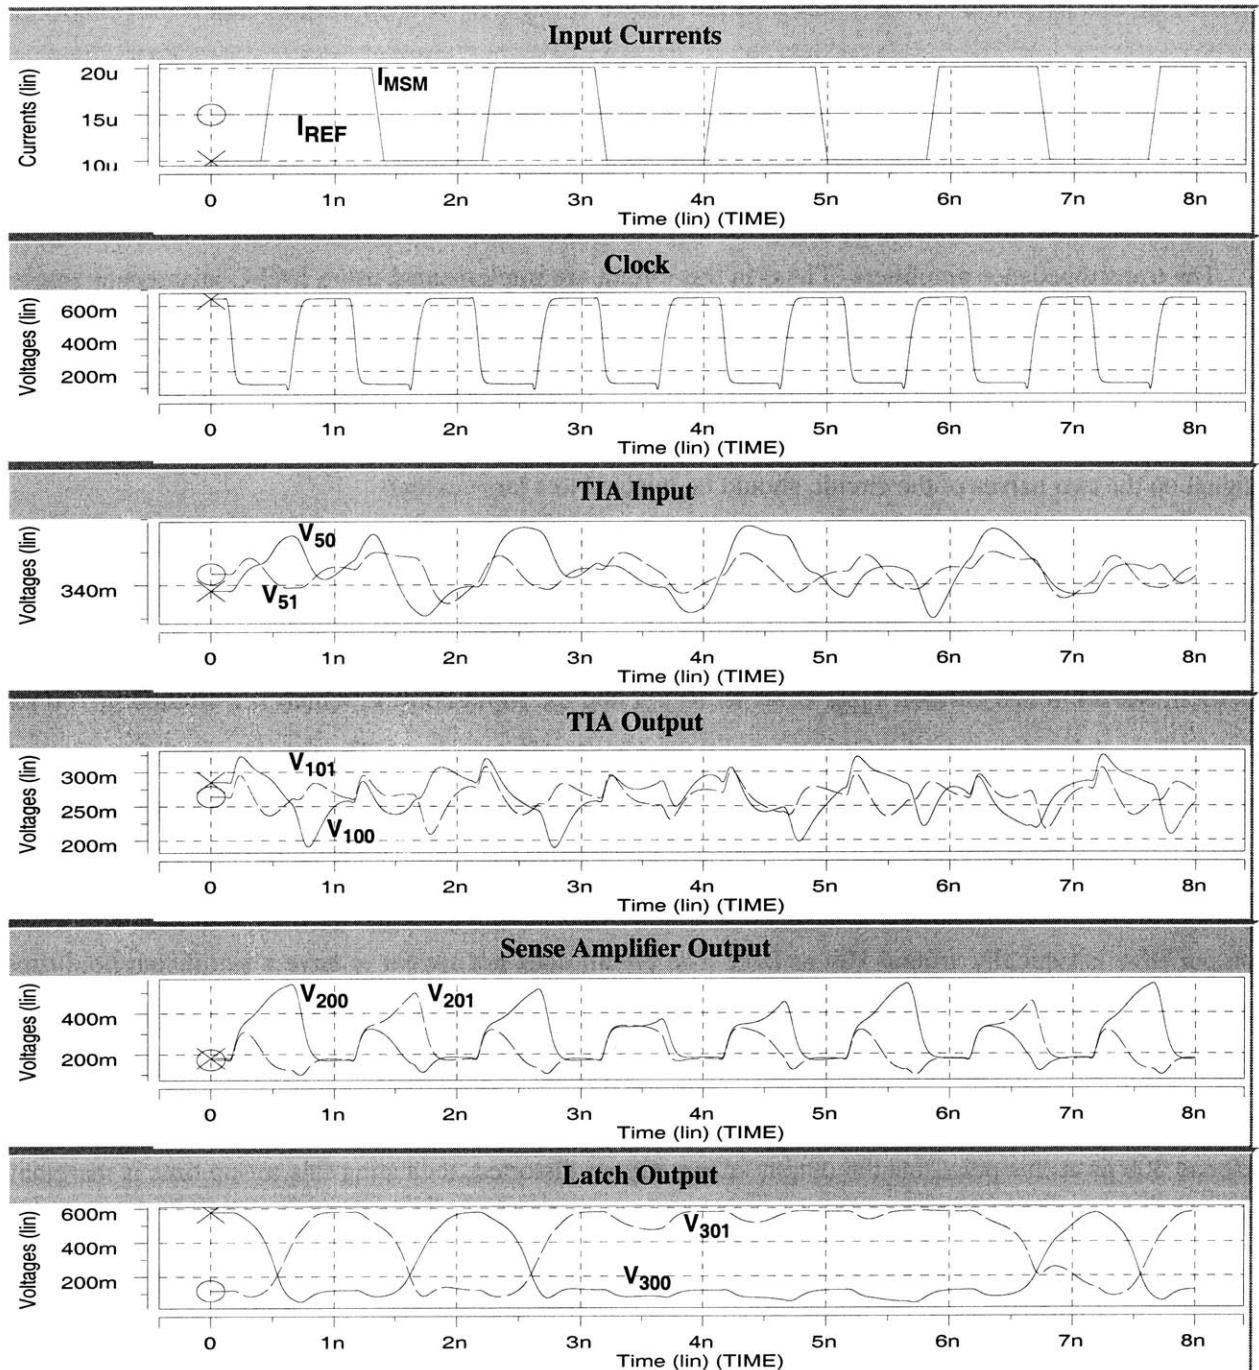


Figure I.4: Simulation of the clocked optical receiver

The simulation includes parasitic interconnect capacitances, source/drain resistances, and source/drain junction diodes. The node numbers refer to the schematic in Figure I.3. An detector capacitance of 100 fF is assumed at the inputs.




Figure I.5: Layout of the clocked optical receiver

The inputs, clock, and output are connected on metal-2 at the points indicated. Metal-3 and metal-4, not shown, are present over the entire area and are used for the 2 V supply and ground, respectively. An n^+ isolation ring surrounds the cell. The isolation structure should be connected (to the most positive supply available) at the perimeter of the cell using metal-1.

I.2 Characterization

The clocked optical receiver appeared in a test circuit on MIT-OEIC-7 (Appendix C) in conjunction with the MSM27 photodetector (Chapter 4). Because of the poor performance of the photodetector, characterization of the set-up and hold times of the circuit has not been pursued. A test cell containing the clocked optical receiver also appeared with 200 k Ω source/drain implant resistors at its inputs to set the MSM and reference currents. Because of the large capacitance of these resistors, high speed operation of the test circuit was not possible. Static characterization of the receiver was completed giving a bound on its sensitivity.

Figure I.6 summarizes an experiment in which the reference current, I_{REF} , was fixed at 5.03 μA while the MSM current, I_{MSM} , was swept while the receiver was clocked at 10 MHz. The receiver output was buffered by a Vitesse ECL output driver which switches the output between ~ 0 V and ~ 1.2 V. When I_{MSM} is well below or above I_{REF} , the receiver consistently sets the output low or high, respectively. When the difference between I_{MSM} and I_{REF} is comparable to the input referred noise level of the receiver, the output logic value generated following each falling edge of the clock becomes random, with the fraction of high- and low-level outputs reflecting the noise distribution

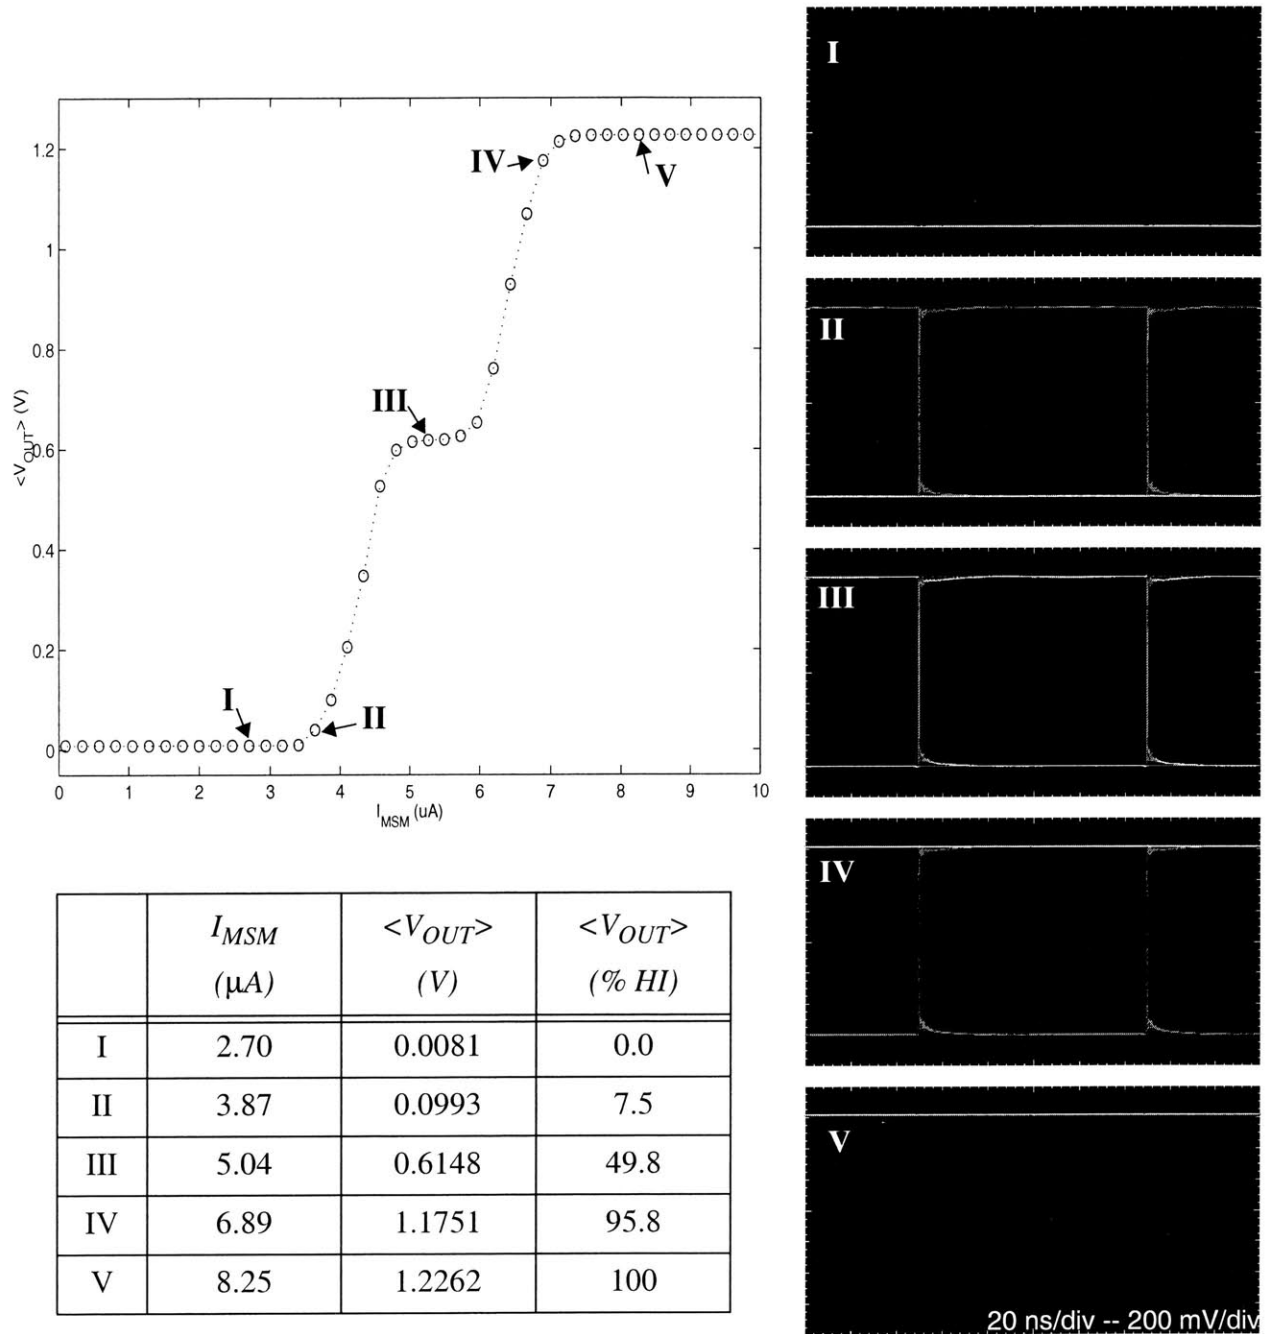


Figure I.6: Static characteristics of clocked optical receiver

The reference current was held at $I_{REF}=5.03 \mu\text{A}$ while I_{MSM} was swept. A 10 MHz clock was applied. The plot above shows the time-averaged output voltage, $\langle V_{OUT} \rangle$. Five operating points are identified and the output eye pattern at each is shown to the right. Driven by noise, the receiver selects at random between high and low values when the input and reference are similar in value. For each operating point, the “ $\langle V_{OUT} \rangle$ % HI” value in the table indicates the fraction of time the output spends in the high state.




Figure I.7: Additional static characteristics of clocked optical receiver

The measurement described in Figure I.6 is repeated here for various values of I_{REF} .

The behavior described above is observed in the output eye patterns shown in Figure I.6. When I_{MSM} is well above or below I_{REF} (operating points I and V), the output is a constant low or high, whereas when $I_{REF} = I_{MSM}$ (point III) the eye pattern is equally weighted in the low and high states. In intermediate cases (points II and IV) the eye pattern is weighted more heavily in one or the other state. The fraction of high- and low-level outputs determines the time-average value of the output voltage. The time-average output voltage is plotted in Figure I.6 indicates that the width of the input noise distribution is around $3.5 \mu A$. Figure I.7 repeats the above measurement for various different values of I_{REF} . The result in each case is a curve which is nearly identical to that of Figure I.6 but shifted to reflect the value of I_{REF} . This is a lower bound on the sensitivity of the receiver, and is below the $\sim 10 \mu A$ signal level used in the simulation in Figure I.4.

I.3 Conclusion

Compared to the continuous time receiver on Chapter 7, the clocked optical receiver presented in this appendix dissipates roughly one-third the power and occupies around one-fourth the area. These savings are particularly beneficial in densely interconnected systems, such as smart pixel arrays. The small power dissipation of this type of circuit also makes it a particularly attractive candidate for scaling to higher bit rates. The static characterization results above have confirmed the basic functionality of the clocked optical receiver presented here. Future high-speed characterization of this design, including bit-error-rate testing, is needed to determine its performance limits.

Bibliography

- [1] K. V. Shenoy, "VLSI GaAs MESFETs and laser diodes for monolithic optoelectronic circuits," *S.M. Thesis*, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1992.
- [2] K. V. Shenoy, "Monolithic optoelectronic VLSI circuit design and fabrication for optical interconnects," *Ph.D. Thesis*, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1995.
- [3] K. V. Shenoy, C. G. Fonstad, and J. M. Mikkelsen. "High-temperature stability of refractory-metal VLSI GaAs MESFETs," *IEEE Electron Device Letters*, 15:106-108, 1994.
- [4] A. C. Grot, D. Psaltis, K. V. Shenoy, and C. G. Fonstad, Jr. "Integration of LED's and GaAs circuits by MBE regrowth," *IEEE Photonics Technology Letters*, 6:819-821, 1994.
- [5] K. V. Shenoy, C. G. Fonstad, Jr., A. C. Grot, and D. Psaltis. "Monolithic optoelectronic circuit design and fabrication by epitaxial growth on commercial VLSI GaAs MESFETs," *Photonics Technology Letters*, 7:1-3, 1995.
- [6] E. K. Braun, K. V. Shenoy, C. G. Fonstad, Jr., and J. M. Mikkelsen. "Elevated temperature stability of GaAs digital integrated circuits," *IEEE Electron Device Letters*, 17:37-39, 1996.
- [7] E. K. Braun, "Elevated temperature stability of GaAs integrated circuits," *S.M. Thesis*, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1995.
- [8] R. J. Aggarwal, K. V. Shenoy, and C. G. Fonstad, Jr., "A technology for monolithic integration of high-indium-fraction resonant tunneling diodes with commercial MESFET VLSI electronics," Proceedings of the InP and Related Compounds Conference, 85-88, Japan, May 1995. And, R. J. Aggarwal, "Design of resonant-tunneling diodes for a GaAs integrated SRAM," *Ph.D. Thesis*, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1996.
- [9] H. Wang, J. Luo, K. V. Shenoy, Y. Royter, and C. G. Fonstad, Jr. "Monolithic integration of SEED's and VLSI GaAs circuits by Epitaxy on Electronics," *IEEE Photonics Technology Letters*, 9:607-609, 1997.
- [10] J. F. Ahadian, "Phosphide-based optical emitters for monolithic integration with GaAs MESFETs," *S.M. Thesis*, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1995.

-
- [11] J. F. Ahadian, P. T. Vaidyanathan, S. G. Patterson, Y. Royter, G. S. Petrich, L. A. Kolodziejski, S. Prasad, and C. G. Fonstad, Jr., "Demonstration of sub-475°C Epitaxy-on-Electronics process for monolithic very large scale OEIC fabrication," *IEEE/LEOS Annual Meeting*, Boston, MA, Nov. 18-21, 1996, PDP 1.7.
- [12] J. F. Ahadian, S. G. Patterson, P. T. Vaidyanathan, Y. Royter, D. E. Mull, G. S. Petrich, W. D. Goodhue, S. Prasad, L. A. Kolodziejski, and C. G. Fonstad, Jr., "Monolithic OEICs using GaAs VLSI technology," *Light-Emitting Diodes: Research, Manufacturing, and Applications*, Proceedings of SPIE, Vol. 3002, E. F. Shubert, Ed., pp. 180-185, Jan Jose, CA, 13-14 February 1997.
- [13] J. F. Ahadian, S. G. Patterson, P. T. Vaidyanathan, Y. Royter, D. E. Mull, G. S. Petrich, W. D. Goodhue, S. Prasad, L. A. Kolodziejski, and C. G. Fonstad, Jr., "Epitaxy-on-Electronics: Building monolithic OEICs on commercial GaAs VLSI," *Proceedings of the 1997 International Conference on GaAs Manufacturing Technology (MANTECH)*, June 1997, San Francisco, CA, pp. 70-73.
- [14] P. T. Vaidyanathan, "Fabrication and characterization of GaAs/InGaP double heterostructure light emitting diodes for monolithic optoelectronic integrated circuits," *M.S. Thesis*, Department of Electrical and Computer Engineering, Northeastern University, 1997.
- [15] J. F. Ahadian, S. G. Patterson, Y. Royter, G. S. Petrich, L. A. Kolodziejski, C. G. Fonstad, Jr., P. T. Vaidyanathan, S. Prasad, D. E. Mull, and W. D. Goodhue, "Monolithic LED-based OEICs and the OPTOCHIP research foundry," *Proceedings of the 1998 Conference on Lasers and Electro-Optics (CLEO)*, San Francisco, CA, 3-8 May 1998, CThG1, pp. 349-350.
- [16] J. F. Ahadian, P. T. Vaidyanathan, S. G. Patterson, Y. Royter, D. E. Mull, G. S. Petrich, W. D. Goodhue, S. Prasad, L. A. Kolodziejski, and C. G. Fonstad, Jr., "Practical OEICs based on the monolithic integration of GaAs/InGaP LEDs with commercial GaAs VLSI electronics," *IEEE Journal of Quantum Electronics*, 34:1117-1123, 1998.
- [17] J. F. Ahadian and C. G. Fonstad, Jr., "Epitaxy-on-Electronics technology for monolithic optoelectronic integration: foundations, development, and status," *Optical Engineering*, 37:3161-3174, 1998.
- [18] J. F. Ahadian, P. T. Vaidyanathan, S. G. Patterson, Y. Royter, D. E. Mull, G. S. Petrich, W. D. Goodhue, Jr., S. Prasad, L. A. Kolodziejski, and C. G. Fonstad, "Feasibility of tera-bps smart pixel arrays using monolithic emitter-based optoelectronic VLSI circuits," *Proceedings of the 1998 IEEE/LEOS Summer Topical Meetings*, Monterey, CA, 20-24 July 1998, WA6, pp. IV/13-14.
- [19] J. F. Ahadian and C. G. Fonstad, Jr., "Development of monolithically integrated transceivers for single- and multi-channel fiber-based optical interconnects," *Workshop Notes for the 10th Annual Workshop on Interconnects Within High Speed Digital Systems (IHSDS)*, Santa Fe, NM, 9-12 May 1999, 2.7.
- [20] D. A. Hall, B. L. Shoop, J. R. Loy, G. B. Tait, E. K. Ressler, J. F. Ahadian, and C. G. Fonstad, Jr., "Experimental demonstration of a 3x3 monolithically integrated smart pixel neural array based on Epi-
-

- taxy-on-Electronics technology," *Postdeadline presentation at the 1997 OSA Annual Meeting*, Long Beach, CA, October 12-17, 1997.
- [21] D. A. Hall, B. L. Shoop, J. R. Loy, G. B. Tait, E. K. Ressler, J. F. Ahadian, and C. G. Fonstad, Jr., "Toward real-time digital image halftoning on the army after next battlefield: experimental demonstration of a novel EoE smart pixel array using GaAs circuitry and GaAs/InGaP LEDs," *Proceedings of the 21st Army Science Conference*, Norfolk, VA, June 1998, pp. 173-178.
- [22] D. A. Hall, B. L. Shoop, J. R. Loy, G. B. Tait, E. K. Ressler, J. F. Ahadian, and C. G. Fonstad, "Experimental demonstration of OPTOCHIP: a GaAs E-o-E smart pixel neural array for digital halftoning," *Postdeadline presentation at the 1998 IEEE/LEOS Summer Topical Meetings*, Monterey, CA, 20-24 July 1998, PD001.
- [23] D. A. Hall, B. L. Shoop, J. R. Loy, E. K. Ressler, J. F. Ahadian, and C. G. Fonstad, Jr., "Performance of GaAs smart pixel components before and after monolithic integration of InGaP LEDs using Epitaxy-on-Electronics technology," *Optics Express*, 4:151-160, 1999, <http://epubs.osa.org/opticsexpress/>
- [24] T. Maj, A. Kirk, D. Plant, J. Ahadian, C. Fonstad, K. Lear, K. Tatah, M. Robinson, and J. A. Trezza, "Interconnection of a 2D VCSEL array to a receiver array via a fiber image guide," *Technical Digest of the OSA Optics in Computing Conference*, 1999, PD2.
- [25] T. Maj, A. Kirk, D. Plant, J. Ahadian, C. Fonstad, K. Lear, K. Tatah, M. Robinson, and J. A. Trezza, "Interconnection of a two-dimensional VCSEL array to a receiver array via a fiber image guide," *Applied Optics*, Vol. 39, April 2000, To be published.
- [26] Vitesse Semiconductor Corporation, 741 Calle Plano, Camarillo, CA, 93012.
- [27] Fiber optic network map for 1 Terabit, Inc. <http://www.terabit.net/>.
- [28] J. Greason, "An analytic interconnect density model for process targetting", *MIT Microsystems Technology Laboratory VLSI Seminar*, September 16, 1997, Cambridge MA.
- [29] 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, December 1971.
- [30] P. Christie, "A fractal analysis of interconnection complexity", *Proceedings of the IEEE*, Vol. 81, No. 10, pp. 1492-1499, October 1993.
- [31] W. E. Donath, "Placement and average interconnections lengths of computer logic", *IEEE Transactions on Circuits and Systems*, Vol. CAS-26, No. 4, pp. 272-277, April 1979.
- [32] J. A. Davis, V. K. De, and J. D. Meindl, "A priori wiring estimation and optimal multilevel wiring networks for portable ULSI systems", *Proceedings of the 46th Electronics Components and Technology Conference*, Orlando, Florida, 1996, pp. 1002-1008.

-
- [33] Anis Husein, DARPA. See for instance, C. Brown, "DARPA looks to optical interconnects," *Electronic Engineering Times*, May 26, 1997.
- [34] T. Nakahara, S. Matsuo, S. Fukushima, and T. Kurokawa, "Performance comparison between multiple-quantum-well modulator-based and vertical-cavity-surface-emitting laser-based smart pixels," *Applied Optics*, 35:860-871, 1996.
- [35] A. Chenevas-Paule, T. Collette, P. Scheer, J-P. Bousinac, P. Churoux, and S. Paineau, "From inter-system to ICs level, VCSEL as a key component for optoelectronic active interconnects, a survey," *Proceedings of the 47th Electronic Components and Technology Conference (ECTC)*, San Jose, CA, 17-22 May 1997, pp. 363-367.
- [36] A. L. Lentine, "Qualitative comparison of MQW modulator and VCSEL based OE-VLSI: a systems perspective," *IEEE/LEOS Annual Meeting*, 1997, WJ4.
- [37] A. V. Krishnamoorthy and K. W. Goosen, "Status of optoelectronic-VLSI technology based on electroabsorption modulators," *Heterogeneous Integration: Systems on a Chip*, Proceedings of SPIE, Vol. CR70, A. Husain and M. Fallahi, Eds., pp. 229-252, Jan Jose, CA, 24-30 January 1998.
- [38] M. K. Hibbs-Brenner, Y. Liu, R. Morgan, and J. Lehman, "VCSEL based optoelectronic smart pixels," *Heterogeneous Integration: Systems on a Chip*, Proceedings of SPIE, Vol. CR70, A. Husain and M. Fallahi, Eds., pp. 206-228, Jan Jose, CA, 24-30 January 1998.
- [39] G. J. Grimes, G. L. Sonnier, C. J. Sherman, W. K. Honea, and J. S. Helton, "Optoelectronics for large scale optical backplanes," *Proceedings of the 45th Electronic Components and Technology Conference (ECTC)*, Las Vegas, NV, 21-24 May 1995, pp. 548-551.
- [40] G. J. Grimes, J. P. C. Markush, Y. M. Wong, P. J. Anthony, W. R. Holland, E. G. Priest, C. J. Sherman, S. R. Peck, D. J. Muehlner, C. C. Faudskar, J. S. Nyquist, J. S. Helton, G. L. Sonnier, J. V. Gates, W. K. Honea, and J. R. Bortolini, "Photonic packaging using laser/receiver arrays and flexible optical circuits" *Proceedings of the 46th Electronic Components and Technology Conference (ECTC)*, Orlando, FL, 28-30 May 1996, pp. 500-505.
- [41] G. J. Grimes, "Packaging of emerging optical interconnection technologies in telecom platforms," *Proceedings of the 47th Electronic Components and Technology Conference (ECTC)*, San Jose, CA, 17-22 May 1997, pp. 399-403.
- [42] C. DeCusatis, "Optical data communication: fundamentals and future directions," *Optical Engineering*, 37:3082-3099, 1998.
- [43] L. J. Norton, F. Carney, N. Choi, K. Y. Chun, R. K. Denton, Jr., D. Dias, J. Knapp, M. Meyering, C. Ngo, S. Planer, G. Raskin, E. Reyes, J. Sauvageau, D. B. Schwartz, S. G. Shook, J. Yoder, and Y. Wen, "OPTOBUS I: a production parallel fiber optical interconnect," *Proceedings of the 47th Electronic Components and Technology Conference (ECTC)*, San Jose, CA, 17-22 May 1997, pp. 204-209.
-

-
- [44] W. Jiang, L. J. Norton, P. Kiely, D. B. Schwartz, B. Gable, M. Lebby, and G. Raskin, "Vertical cavity surface emitting laser-based parallel optical data link," *Optical Engineering*, 37:3113-3118, 1998.
- [45] Y.-M. Wong, D. J. Muehler, C. C. Faudskar, D. B. Bucholz, M. Fishteyn, J. L. Brander, W. J. Parzygat, R. A Morgan, T. Mullally, R. E. Leibenguth, G. D. Guth, M. W. Focht, K. G. Glogovsky, J. L. Zilko, J. V Gates, P. J. Anthony, B. H. Tyrone, Jr., T J. Irelant, D. H. Lewis, D. F. Smith, S. F. Nati, D. K. Lewis, D. L. Rogers, H. A. Aispain, S. M. Gowda, S. G. Walker, Y. H. Kwark, R. J. S Bates, D. M. Kuchta, and J. D. Crow, "Technology development for a high-density 32-channel 16-Gb/s optical data link for optical interconnection applications for the optoelectronic technology consortium (OETC)," *Journal of Lightwave Technology*, 13:995-1016, 1995.
- [46] J. D. Crow, J. Choi, M. S. Cohen, G. Johnson, D. Kuchta, D. Lacey, S. Ponnappalli, P. Papeljugoski, K. Stawiasz, J. Trehella, P. Xiao, S. Tremblay, S. Ouimet, A. Lacerte, M. Gauvin, D. Booth, W. Nation, T. L. Smith, B. A. DeBaun, G. D. Henson, S. M. Igl, N. A. Lee, A. J. Peikarczyk, A. S. Kuczma, and S. L. Spanoudis, "The Jitney parallel optical interconnect" *Proceedings of the 46th Electronic Components and Technology Conference (ECTC)*, Orlando, FL, 28-30 May 1996, pp. 292-300.
- [47] K. H. Hahn, K. S. Giboney, R. E. Wilson, J. Straznicky, E. G. Wong, M. R. Tan, R. T. Kaneshiro, D. W. Dolfi, E. H. Mueller, A. E. Plotts, D. D. Murray, J. E. Marchegiano, B. L. Booth, B. J. Sano, B. Madhavan, B. Raghavan, and A. F. J. Levi, "Gigabyte/s data communications with the POLO parallel optical link" *Proceedings of the 46th Electronic Components and Technology Conference (ECTC)*, Orlando, FL, 28-30 May 1996, pp. 301-307.
- [48] J. Nishikido, S. Fujita, Y. Arai, Y. Akahori, S. Hino, and K. Yamasaki, "Multigigabit multichannel optical interconnection module for broadband switching system," *Journal of Lightwave Technology*, 13:1104-1110, 1995.
- [49] H. Karstensen, C. Hanke, M. Honsberg, J.-R. Kropp, J. Wieland, M. Blaser, P. Weger, and J. Popp, "Parallel optical interconnect for uncoded data transmission with 1 Gb/s-per-channel capacity, high dynamic range, and low power consumption," *Journal of Lightwave Technology*, 13:1017-1030, 1995.
- [50] Gore parallel fiber optic data link, W. L. Gore & Associates, Inc., 750 Ott's Chapel Road, Newark DE 19714, <http://www.gore.com/>.
- [51] H. Karstensen, J. Wieland, R. Dal'Ara, and M. Blaser, "Parallel optical link for multichannel interconnections at gigabit rate," *Optical Engineering*, 37:3119-3123, 1998.
- [52] The PAROLI parallel optical link, Infineon Technologies AG, PO Box 800949, 81609 Munich, Germany, <http://www.infineon.com/>.
- [53] A. Neyer, B. Wittmann, and M. Johnck, "Plastic-optical-fiber-based parallel optical interconnects," *IEEE Journal of Selected Topics in Quantum Electronics*, 5:193-200, 1999.
-

-
- [54] R. A. Nordin, W. R. Holland, and M. A. Shahid, "Advanced optical interconnection technology in switching equipment," *Journal of Lightwave Technology*, 13:987-99, 1995.
- [55] M. A. Shahid, N. R. Lampert, A. W. Carlisle, D. A. Hendrickson, D. M. Emmerich, T. E. McNeil, and J. E. George, "Small and efficient connector systems" *Proceedings of the 49th Electronic Components and Technology Conference (ECTC)*, San Diego, CA, 1-4 June 1999, pp. 375-380.
- [56] M. Niburg, P. Eriksen, K. Gustafsson, H. Hesselborn, G. Palmkog, M. Robertsson, N. Tran, K. Akermark, A. Djupsjobacka, T. Juhola, B. Kerzar, M. Mokhtari, and G. Schuppener, "A complete sub-system of parallel optical interconnects for telecom applications," *Proceedings of the 46th Electronic Components and Technology Conference (ECTC)*, Orlando, FL, 28-30 May 1996, pp. 259-236.
- [57] T. Hashimoto, Y. Nakasuga, Y. Yamada, H. Terui, M. Yanagisawa, Y. Akahori, Y. Tohmori, K. Kato, and Y. Suzuki, "Multichip optical hybrid integration technique with planar lightwave circuit platform," *Journal of Lightwave Technology*, 16:1249-1258, 1998.
- [58] B. L. Booth, J. E. Marchegiano, C. T. Chang, R. J. Furmanak, D. M. Graham, and R. G. Wagner, "Polyguide polymeric technology for optical interconnect circuits and components," *Optoelectronic Interconnects and Packaging IV*, Proceedings of SPIE, Vol. 3005, R. T. Chen and P. S. Guilfoyle, Eds., pp. 238-251, Jan Jose, CA, 12-14 February 1997.
- [59] T. Li, S. Tang, R. Wickman, L. Wu, F. Li, M. Dubinovsky, and R. T. Chen, "Polymer waveguide based high speed clock signal distribution system," *Optoelectronic Interconnects and Packaging IV*, Proceedings of SPIE, Vol. 3005, R. T. Chen and P. S. Guilfoyle, Eds., pp. 128-135, Jan Jose, CA, 12-14 February 1997.
- [60] J. E. Thomson, H. Levesque, E. Savov, F. Horwitz, B. L. Booth, and J. E. Marchegiano, "Optical waveguide circuit board with a surface-mounted optical receiver array," *Optical Engineering*, 33:939-945, 1994.
- [61] M. E. Robertsson, O.-J. Hagel, G. Gustafsson, A. Dabek, M. Popall, L. Cergel, P. Wennekers, P. Kiely, M. Leby, and T. Lindahl, "O/E-MCM packaging with new patternable dielectric and optical materials" *Proceedings of the 48th Electronic Components and Technology Conference (ECTC)*, Seattle, WA, 25-28 May 1998, pp. 1413-1421.
- [62] K. W. Jelly, G. T. Valliath, and J. W. Stafford, "High-speed chip-to-chip optical interconnect," *IEEE Photonics Technology Letters*, 4:1157-1159, 1992.
- [63] F. A. P. Tooley, "Challenges in optically interconnecting electronics," *Journal of Selected Topics in Quantum Electronics*, 2:3-13, 1996.
- [64] A. C. Walker, M. P. Y. Desmulliez, M. G. Forbes, S. J. Fancey, G. S. Buller, M. R. Taghizadeh, J. A. B. Dines, C. R. Stanley, G. Pennelli, A. R. Boyd, P. Horan, D. Byrne, J. Hegarty, S. Eitel, H.-P. Gauggel, K.-H. Gulden, A. Gauthier, P. Benabes, J.-L. Gutzwiller, and M. Goetz, "Design and construction of an
-

- optoelectronic crossbar switch containing a terabit per second free-space optical interconnect," *IEEE Journal of Selected Topics in Quantum Electronics*, 5:236-249, 1998.
- [65] A. V. Krishnamoorthy, J. E. Ford, F. E. Kiamilev, R. G. Rozier, S. Hunsche, K. W. Goosen, B. Tseng, J. A. Walker, J. E. Cunningham, W. Y. Jan, and M. C. Nuss, "Design and construction of an optoelectronic crossbar switch containing a terabit per second free-space optical interconnect," *Journal of IEEE Selected Topics in Quantum Electronics*, 5:261-275, 1999.
- [66] C. B. Kuznia, J.-M. Wu, C.-H. Chen, B. Hoanca, L. Cheng, A. G. Weber, and A. A. Sawchuk, "Two-dimensional parallel pipeline smart pixel array cellular logic (SPARCL) processors--chip design and system implementation," *IEEE Journal of Selected Topics in Quantum Electronics*, 5:376-386, 1999.
- [67] R. A. Nordin, A. F. J. Levi, R. N. Nottenburg, J. O'Gorman, T. Tanbun-Ek, and R. A. Logan, "A systems perspective on digital interconnect technology," *Journal of Lightwave Technology*, 10:811-827 1992.
- [68] C. R. Theorin, S. P. Kilcoyne, F. H. Peters, R. D. Martin, and M. N. Donhowe, "A 'seamless migration' to VCSEL-based optical data links," in *Vertical-Cavity Surface-Emitting Lasers*, K. D. Choquett and D. G. Deppe, Eds., Proceedings of SPIE, Vol. 3003, pp. 120-130, Jan Jose, CA, 13-14 February 1997.
- [69] D. A. B. Miller, "Physical reasons for optical interconnection," *International Journal of Optoelectronics*, 11:155-168, 1997.
- [70] D. A. B. Miller, "Dense two-dimensional integration of optoelectronics and electronics for interconnections," *Heterogeneous Integration: Systems on a Chip*, Proceedings of SPIE, Vol. CR70, A. Husain and M. Fallahi, Eds., pp. 80-109, Jan Jose, CA, 24-30 January 1998.
- [71] M. R. Feldman, S. C. Esener, C. C. Guest, and S. H. Lee, "Comparison between optical and electrical interconnects based on power and speed considerations," *Applied Optics*, 27:1742-1751, 1988.
- [72] J. W. Goodman, "Optics as an interconnect technology", *Optical processing and computing*, H. H. Arsenault, T. Szoplik, and B. Bacukow, Eds., Academic Press, San Diego, 1989.
- [73] J. W. Goodman, F. J. Leonberger, S. Kung, and R. A. Athale, "Optical interconnections for VLSI systems," *Proceedings of the IEEE*, 72:850, 1984.
- [74] M. W. Haney and M. P. Christensen, "Fundamental geometric advantages of free-space optical interconnects," *Proceedings of the 3rd International Conference on Massively Parallel Processing using Optical Interconnections (MPPOI)*, Maui, Hawaii, 27-29 October 1996, pp. 16-23.
- [75] S. D. Mukherjee, "OEICs for optical interconnects," *Optoelectronic Integration--Physics, Technology and Applications*, O. Wada, Ed., Kluwer Academic Publishers, Norwell, MA, 1994.

-
- [76] T. J. Drabik, "Optoelectronic integrated systems based on free-space interconnects with an arbitrary degree of space variance," *Proceedings of the IEEE*, 82:1595-1622, 1994.
- [77] H. S. Hinton, T. J. Cloonan, F. B. McCormick, Jr., A. L. Lentine, and F. A. P. Tooley, "Free-space digital optical systems," *Proceedings of the IEEE*, 82:1632-1649, 1994.
- [78] I. Hayashi, "ULSI optical interconnection--technologies for future photo-electronic integrated systems," in *Integrated Optics and Optoelectronics*, K. Wong and M. Razeghi, Eds., Proceedings of SPIE, Vol. CR45, pp. 3-14, Los Angeles, CA, 21-23 January 1993.
- [79] M. S. Goodman and E. Arthurs, "Telecommunications systems applications for optoelectronic integrated circuits", *Integrated Optoelectronics*, M. Dagenais, R. F. Leheny, and J. Crow, pp. 3-38, Academic Press, San Diego, 1995.
- [80] J. Crow, J. M. Jaffe, and M. W. Sachs, "Computing system applications", *Integrated Optoelectronics*, M. Dagenais, R. F. Leheny, and J. Crow, pp. 39-80, Academic Press, San Diego, 1995.
- [81] D. A. B. Miller and H. M. Ozaktas, "Limit to the bit-rate capacity of electrical interconnects from the aspect ratio of the system architecture," *Journal of Parallel and Distributed Computing*, 41:4252, 1997.
- [82] R. Walker, "Tb/s chip I/O -- how close are we to practical reality?," *Workshop Notes for the 10th Annual Workshop on Interconnects Within High Speed Digital Systems (IHSDS)*, Santa Fe, NM, 9-12 May 1999, 1.3, email: walker@opus.hpl.hp.com.
- [83] J. R. Broomall and H. Deusen, "Extending the usefull range of copper interconnects for high data rate signal transmission," *Proceedings of the 47th Electronic Components and Technology Conference (ECTC)*, San Jose, CA, 17-22 May 1997, pp. 196-203.
- [84] M. Chang, M. Nesnidal, G. Emmel, R. Marsland, R. Williamson, M. Rodwell, M. Degerstrom, G. Fokken, and D. Schwab, "Low-cost, high-speed fiber-optic tranceiver," *Proceedings of the DARPA/ETO 1999 Optoelectronics Review*, San Diego, CA, 2-5 August 1999, Section: Optical Micro Networks.
- [85] K. L. Lear, H. Q. Hou, V. M. Hietala, K. D. Choquette, and J. R. P. Schneider, "Engineering high performance vertical cavity lasers," *Proceedings of the Conference on Optoelectronic and Microelectronic Materials and Devices (COMMAD)*, Canberra, Australia, 1996, pp. 33-41.
- [86] S. Sasaki, N Tanaka, S. Yamaguchi, M. Nakamura, and T. Hayashi, "An optical active connector: and optical interconnect module with an delectrical connector interface," *Proceedings of the 46th Electronic Components and Technology Conference (ECTC)*, Orlando, FL, 28-30 May 1996, pp. 512-519.
- [87] Q. Tan, G. De Pestel, G. Willems, J. Cannaerts, W. Rehm, N. Kaiser, K. Vandeputte, A. Van Hove, and J. Van Koetsem, "A parallel optical link for intra- and inter-rack interconnections," *Proceedings of the*
-

-
- 47th *Electronic Components and Technology Conference (ECTC)*, San Jose, CA, 17-22 May 1997, pp. 234-239.
- [88] M. P. Christensen and M. W. Haney, "Two-bounce free-space arbitrary interconnection architecture," *Proceedings of the 4th International Conference on Massively Parallel Processing using Optical Interconnections (MPPOI)*, Montreal, Canada, 22-24 June 1997, pp. 61-67.
- [89] M. W. Haney and M. P. Christensen, "Sliding banyan network," *Journal of Lightwave Technology*, 14:703-710, 1996.
- [90] R. R. Michael, M. P. Christensen, and M. W. Haney, "Experimental evaluation of the 3-D optical shuffle interconnection module of the sliding banyan architecture," *Journal of Lightwave Technology*, 14:1970-1978, 1996.
- [91] M. W. Haney and M. P. Christensen, "The application of smart pixels to high performance Viterbi decoders," *Proceedings of the 1994 IEEE/LEOS Summer Topical Meetings*, 6-13 July 1994, P7, pp. 2/74-75.
- [92] K. V. Guinn and R. C. Frye, "Flip-chip and chip-scale I/O density requirements and printed wiring board capabilities," *Proceedings of the 47th Electronic Components and Technology Conference (ECTC)*, San Jose, CA, 17-22 May 1997, pp. 649-655.
- [93] S. Siala, A. P. Kanjamala, R. N. Nottenburg, and A. F. J. Levi, "Low skew multimode ribbon fibers for parallel optical communication," *Proceedings of the 1995 Optical Fiber Conference (OFC)*, 1995.
- [94] S. Tang, R. T. Chen, L. Garrett, D. Gerold, and M. M. Li, "Design limitations of highly parallel free-space optical interconnects based on arrays of vertical cavity surface-emitting laser diodes, microlenses, and photodetectors," *Journal of Lightwave Technology*, 12:1971-1975, 1994.
- [95] H. W. Johnson and M. Graham, *High-Speed Digital Design: A Handbook of Black Magic*, Prentice Hall, New Jersey, 1993.
- [96] Y. Chen, Z. Chen, and J. Fang, "Optimum placement of decoupling capacitors on packages and printed circuit boards under the guidance of electromagnetic field simulation," *Proceedings of the 46th Electronic Components and Technology Conference (ECTC)*, Orlando, FL, 28-30 May 1996, pp.756-760.
- [97] A. Deutsh, G. V. Kopcsay, P. Restle, G. Katopis, W. D. Becker, H. Smith, P. W. Coteus, C. W. Surovic, B. J. Rubin, R. P. Dunne, T. Gallo, K. A. Jenkins, L. M. Terman, R. H. Dennard, G. A. Sai-Halasz, and D. R. Knebel, "When are transmission-line effects important for on-chip interconnections," *Proceedings of the 47th Electronic Components and Technology Conference (ECTC)*, San Jose, CA, 17-22 May 1997, pp. 704-712.
- [98] Electromagnetic Cmpatibility Laboratory, University of Missouri-Rolla, <http://www.emclab.umr.edu/>
-

-
- [99] J. Bartley, "A user's view of MCM-D/C packaging: is it worth the trouble?" *Proceedings of the 46th Electronic Components and Technology Conference (ECTC)*, Orlando, FL, 28-30 May 1996, pp.144-148.
- [100] B. Chappel, "The fine art of IC design," *IEEE Spectrum*, 36:30-34, 1999.
- [101] D. Herrell, "Power to the package," *IEEE Spectrum*, 36:46-53, 1999.
- [102] M. Chastain, "The role of optics in balanced computer system design," *Workshop Notes*, 10th Annual Workshop on Interconnects Within High Speed Digital Systems, Santa Fe, NM, 9-12 May 1999, 2.2, email: chastain@rsn.hp.com.
- [103] P. Singh and D. L. Landis, "Optimal chip sizing for multi-chip modules," *IEEE Transactions on Components, Packaging, and Manufacturing Technology--Part B*, 17:369-375, 1994.
- [104] G. George, J. P. Krusius, and R. Granitz, "Tiled silicon: a packaging alternative to large silicon chips," *Proceedings of the 45th Electronic Components and Technology Conference (ECTC)*, Las Vegas, NV, 21-24 May 1995, pp. 608-613.
- [105] D. L. Rogers, "The role of OEICs in commercial applications," *IEEE/LEOS 1995 Digest of the LEOS Summer Topical Meetings*, Keystone, CO, August 7-11, 1995, ThD2, pp. 29-30.
- [106] E. A. Fitzgerald and L. C. Kimerling, "Si-based microphotronics and integrated optoelectronics," *MRS Bulletin*, 23(4):39-47, 1998.
- [107] M. Sugo, H. Mori, Y. Sakai, and Y. Itoh, "Stable cw operation at room temperature of a 1.5- μ m wavelength multiple quantum well laser on a Si substrate," *Applied Physics Letters*, 60:472-473, 1992.
- [108] J. M. London, A. H. Loomis, J. F. Ahadian, and C. G. Fonstad, Jr., "Preparation of Si-on-gallium arsenide wafers for monolithic optoelectronic integration," *IEEE Photonics Technology Letters*, 11:958-960, 1999.
- [109] J. M. London, "Wafer bonding for monolithic integration of Si CMOS VLSI electronics with III-V optoelectronic devices," *S.M. Thesis*, Massachusetts Institute of Technology, 1999.
- [110] J. M. London, A. Postige, and C. G. Fonstad, Jr., "Quantum well heterostructures grown by molecular beam epitaxy on silicon-on-gallium arsenide substrates," *Applied Physics Letters*, Vol. 75, November 1999, To be published.
- [111] H. S. Hinton. "Progress in the smart pixel technologies," *IEEE Journal of Selected Topics in Quantum Electronics*, 2:14-23, 1996.
- [112] S. T. Wilkinson, J. Tabler, N. M. Jokerst, M. Brooke, and R. P. Leavitt, "An 8x8 array of resonant cavity enhanced light emitting diodes integrated onto silicon grayscale (32 level) driver circuitry," *IEEE/LEOS Annual Meeting*, Boston, MA, Nov. 18-21, 1996, ThA5.
-

-
- [113] R. Pu, E. M. Hayes, R. Jurrat, C. W. Wilmsen, K. D. Choquette, H. Q. Hou, and K. M. Geib. "VCSEL's bonded directly to foundry fabricated GaAs smart pixel arrays," *IEEE Photonics Technology Letters*, 9:1622-1624, 1997.
- [114] R. Pu, C. Duan, and C. W. Wilmsen, "Hybrid integration of VCSEL's to CMOS integrated circuits," *IEEE Journal of Selected Topics in Quantum Electronics*, 5:201-208, 1999.
- [115] A. V. Krishnamoorthy, L. M. F. Chirovsky, W. S. Hobson, R. E. Leibenguth, S. P. Hui, G. J. Zydzik, K. W. Goosen, J. D. Wynn, B. J. Tseng, J. Lopata, J. A. Walker, J. E. Cunningham, and L. A. D'Asaro, "Operation of 970 nm VCSELs flip-chip bonded to CMOS circuits," IEEE/LEOS Summer Topical Meeting, Monterey, CA, July 20-24, 1998, PD004.
- [116] A. V. Krishnamoorthy, L. M. F. Chirovsky, W. S. Hobson, R. E. Leibenguth, S. P. Hui, G. J. Zydzik, K. W. Goosen, J. D. Wynn, B. J. Tseng, J. Lopata, J. A. Walker, J. E. Cunningham, and L. A. D'Asaro, "Vertical-cavity surface-emitting lasers flip-chip bonded to gigabits-per-second CMOS circuits," *IEEE Photonics Technology Letters*, 11:128-130, 1999.
- [117] L. M. F. Chirovsky, W. S. Hobson, R. E. Leibenguth, S. P. Hui, J. Lopata, G. J. Zydzik, G. Giaretta, K. W. Goosen, J. D. Wynn, A. V. Krishnamoorthy, B. J. Tseng, J. M. Vandenberg, , and L. A. D'Asaro, "Implant-apertured and index-guided vertical-cavity surface-emitting lasers (I^2 -VCSEL's)," *IEEE Photonics Technology Letters*, 11:500-502, 1999.
- [118] S. Matsuo, T. Nakahara, Y. Kohama, Y. Ohiso, and T. Kurokawa. "A monolithic integrated smart pixel using an MSM-PD, MESFET's, and a VCSEL," *IEEE Journal of Selected Topics in Quantum Electronics*, 2:121-127, 1996.
- [119] T. Jakobus, W. Bronner, A. Gaymann, F. Grotjahn, J. Hornung, V. Hurm, K. Kohler, M. Ludwig, and Z. G. Wang, "Monolithically integrated optoelectronic receivers and transmitters on GaAs-wafers," GaAs Manufacturing Technology Conference, June 1997, San Francisco, CA.
- [120] J. Naber, "Digital GaAs Integrated Circuits," *Gallium Arsenide IC Applications Handbook*, D. Fisher and I. Bahl, Eds., Vol. 1, pp. 63-70, Academic Press, San Diego, 1995.
- [121] J. M. Mikkelson, CTO Vitesse Semiconductor, Personal communication of major Vitesse clients, 1997.
- [122] The MOSIS foundry service, USC/ISI, 4676 Admiralty Way, Marina del Rey, CA 90292, email mosis@mosis.edu, <http://www.mosis.org/>.
- [123] K. V. Shenoy, MIT, Personal communicaiton of outgassing results on Vitesse ICs, 1995.
- [124] SIMS analysis of residues performed at AT&T Bell Laboratories complements of J. E. Cunningham, 1995.
-

-
- [125] J. M. Mikkelsen, Vitesse Semiconductor Corporation, Personal communication of observation regarding RIE residues, 1995.
- [126] Shipley SC-1827 photoresist, Shipley Microelectronic Products, 2300 Washington St., Newton, MA 02162-1469, (617)969-5500.
- [127] Futurrex NR8-3000 photoresist, Futurrex, Inc., 12 Cork Hill Rd., Franklin, NJ 07416, (973)209-1563, <http://www.eclipse.net/~futurrex/>.
- [128] Hydrogen cracker manufactured by EPI, 1290 Hammond Rd., Saint Paul, MN 55110. See "On the use of atomic hydrogen in MBE", EPI Application Note, 1994.
- [129] T. Achnich, G. Burri, and M. Ilegems. "Study of oxygen incorporation in AlGaAs layers grown by molecular beam epitaxy," *Journal of Vacuum Science Technology*, 7A:2536-2541, 1989.
- [130] S. V. Iyer, H. P. Meier, S. Ovidia, and C. Parks. "Role of substrate temperature in molecular beam epitaxy growth of high power GaAs/AlGaAs lasers," *Applied Physics Letters*, 60:416-418, 1992.
- [131] W. T. Tsang. "The influence of bulk nonradiative recombination in the wide bandgap regions of molecular beam epitaxy grown GaAs/Al_xGa_{1-x}As DH lasers," *Applied Physics Letters*, 33:245-248, 1978.
- [132] K. Akimoto, M. Kamada, K. Taira, M. Arai, and N. Watanabe, "Photoluminescence killer center in AlGaAs grown by molecular beam epitaxy," *Journal of Applied Physics*, 59:2833, 1986.
- [133] J. M. Olson, R. K. Ahrenkiel, D. J. Dunlavy, B. Keyes, and A. E. Kibbler. "Ultralow recombination velocity at Ga_{0.5}In_{0.5}P/GaAs heterostructures," *Applied Physics Letters*, 55:1208-1210, 1989.
- [134] G. Zhang, J. Nappi, K. Vantinen, H. Asonen, and M. Pessa. "Low threshold current InGaAs/GaAs/GaInP lasers grown by gas source molecular beam epitaxy," *Applied Physics Letters*, 61:96-98, 1992.
- [135] SDL, Inc., 80 Rose Orchard Way, San Jose, CA 95134-1365, (408)943-9411, <http://www.sdli.com/>.
- [136] J. D. Woodhouse, C. A. Wang, J. P. Donnelly, D. Z. Tsang, R. J. Bailey, D. E. Mull, K. Rauschenbach, and O. A. Popov. "Uniform linear arrays of strained-layer InGaAs-AlGaAs quantum-well ridge-waveguide diode lasers fabricated by ECR-IBAE", *IEEE Journal of Quantum Electronics*, 31:1357-1363, 1995.
- [137] FIBE conducted by K. Edinger and J. Melngailis, Laboratory for Ion Beam Research and Application, University of Maryland, College Park, MD.
- [138] P. K. Footner, B. P. Richards, and R. B. Yates, "Purple plague: eliminated or just forgotten?" *GEC Technical Journal*, 4:174-180, 1986.
- [139] *AT&T/DARPA Workshop on Hybrid-SEED Smart Pixels*, George Mason University, 1995.
-

-
- [140] IBOOST This cell was credited to P. R. Nuytkens in [2], but Mr. Nuytkens has disavowed any such involvement.
- [141] Standard H-GaAs III MSM photodetectors characterized by M. Tachikawa, NTT Opto-electronics Laboratories, 1996, email: tachi@aecl.ntt.jp.
- [142] M. Azadeh, R. B. Darling, W. R. Babbitt, "Smart pixels with built-in illumination and optical feedback," *Proceedings of the 1998 Conference on Lasers and Electro-Optics (CLEO)*, San Francisco, CA, 3-8 May 1998, CThG2, p. 350.
- [143] M. Azadeh, R. B. Darling, W. R. Babbitt, "Characteristics of optoelectronic feedback for smart pixels with smart illumination," *IEEE Journal of Selected Topics in Quantum Electronics*, 5:172-177, 1999.
- [144] C.-H. Chen, C. B. Kuznia, B. Hoanca, J.-M. Wu, and A. A. Sawchuk, "Design and demonstration of smart pixel array network interface for 2D parallel data packet networks," OSA Annual Meeting, Long Beach, CA, October 12-17, 1997, WS1.
- [145] Smart pixel array circuit design by J. Drolet, California Institute of Technology.
- [146] P. Bhattacharya, *Semiconductor Optoelectronic Devices*, Prentice Hall, Englewood Cliffs, New Jersey, 1994, p. 214.
- [147] Number display circuit designed by Y. Royter, Massachusetts Institute of Technology.
- [148] J. Campbell, "Photodetectors for optoelectronic integrated circuits", *Integrated Optoelectronics*, M. Dagenais, R. F. Leheny, and J. Crow, pp. 430-441, Academic Press, San Diego, 1995.
- [149] T. Sugeta and T. Urisu, "High-gain metal-semiconductor-metal photodetectors for high-speed optoelectronic circuits," *IEEE Transactions on Electron Devices*, 26:1855-1856, 1979.
- [150] T. Sugeta, T. Urisu, S. Sakata, and Y. Mizushima, "Metal-semiconductor-metal photodetector for high speed optoelectronic circuits," *Japanese Journal of Applied Physics*, Volume 19 Supplement 19-1:459-464, 1980.
- [151] C. W. Slayman and L. Figueroa, "Frequency and pulse response of a novel high speed interdigital surface photoconductor (IDPC)," *IEEE Electron Device Letters*, 2:112-114, 1981.
- [152] W. Roth, H. Schumacker, J. Kluge, H. J. Geelen, and H. Beneking, "The DSI diode -- a fast large-area optoelectronic detector," *IEEE Transactions on Electron Devices*, ED-32:1035-1036, 1985.
- [153] M. Ito and O. Wada, "Low dark current GaAs metal-semiconductor-metal (MSM) photodiodes using WSi_x contacts," *IEEE Journal of Quantum Electronics*, 22:1073-1077, 1986.
-

- [154] M. Makiuchi, H. Hamaguchi, T. Kumai, M. Ito, O. Wada, and T. Sakurai, "A monolithic four-channel photoreceiver integrated on a GaAs substrate using metal-semiconductor-metal photodiodes and FET's," *IEEE Electron Device Letters*, 6:634-635, 1985.
- [155] W. S. Lee, G. R. Adams, J. Mun, and J. Smith, "Monolithic GaAs photoreceiver for high-speed signal processing applications," *Electronics Letters*, 22:147-148, 1986.
- [156] D. L. Rogers, "Monolithic integration of a 3-GHz detector/preamplifier using a refractory-gate ion-implanted MESFET process," *IEEE Electron Device Letters*, 7:600-602, 1986.
- [157] J. D. Crow, C. J. Anderson, S. Bermon, A. Callegari, J. F. Ewen, J. D. Feder, J. H. Greiner, E. P. Harris, P. D. Hoh, H. J. H. Hovel, J. H. Magerlein, T. E. McKoy, A. T. S. Pomerene, D. L. Rogers, G. J. Scott, M. Thomas, G. W. Mulvey, B. K. Ko, T. Ohashi, M. Scontras, and D. Widiger, "A GaAs MESFET IC for optical multiprocessor networks," *IEEE Transactions on Electron Devices*, 36:263-268, 1989.
- [158] R. Bates, J. Crow, J. Ewen, and D. Rogers, "Scaling and system issues of OEIC design", *Integrated Optoelectronics*, M. Dagenais, R. F. Leheny, and J. Crow, pp. 502-525, Academic Press, San Diego, 1995.
- [159] D. L. Rogers, "Integrated optical receivers using MSM detectors," *Journal of Lightwave Technology*, 9:1635-1638, 1991.
- [160] G. E. Stillman, V. M. Robbins, and N. Tabatabaie, "III-V compound semiconductor devices: optical detectors," *IEEE Transactions on Electron Devices*, ED-31:1643-1655, 1984.
- [161] K. D. Li Dessau, "Application Note 1: Insights into High-Speed Detectors and High-Frequency Techniques", New Focus, Inc., 2630 Walsh Avenue, Santa Clara, CA 95051, tel. 408-980-8088.
- [162] V. Hurn, W. Benz, W. Bronner, A. Hulsmann, T. Jakobus, K. Kohler, A. Leven, M. Ludwig, B. Raynor, J. Rosenzweig, M. Schlechtweg, and W. Thiede, "40 Gbit/s 1.55 μm *pin*-HEMT photoreceiver monolithically integrated on 3 in GaAs substrate," *Electronics Letters*, 34:2060-2062, 1998.
- [163] V. Hurn, W. Benz, W. Bronner, M. Dammann, T. Jakobus, G. Kaufel, K. Kohler, Z. Lao, M. Ludwig, B. Raynor, J. Rosenzweig, and M. Schlechtweg, "10 Gbit/s long wavelength *pin*-HEMT photoreceiver monolithically grown on GaAs," *Electronics Letters*, 33:1653-1654, 1997.
- [164] Z. Lao, V. Hurn, W. Benz, W. Bronner, A. Hulsmann, T. Jakobus, K. Kohler, M. Ludwig, B. Raynor, J. Rosenzweig, and M. Schlechtweg, and W. Thiede, "20-Gb/s 14-k Ω transimpedance long-wavelength MSM-HEMT photoreceiver OEIC," *IEEE Photonics Technology Letters*, 10:710-712, 1998.
- [165] V. Hurn, W. Benz, W. Bronner, T. Fink, D. Fritsch, M. Haupt, P. Hofmann, T. Jakobus, K. Kohler, M. Ludwig, K. Mause, B. Raynor, and J. Rosenzweig, "1.3 μm monolithic integrated optoelectronic receiver using an InGaAs MSM photodiode and AlGaAs/GaAs HEMTs grown on GaAs," *Electronics Letters*, 31:67-68, 1995.

-
- [166] H. Wang, "Monolithic integration of 1.55 micron photodetectors with GaAs electronics for high speed optical communications," *Ph.D. Thesis*, Department of Material Science and Engineering, Massachusetts Institute of Technology, 1998.
- [167] T. Knodl, H. K. H. Choy, J. L. Pan, R. King, R. Jager, G. Lullo, J. F. Ahadian, R. J. Ram, C. G. Fonstad, Jr., and K. J. Ebeling, "RCE photodetectors based on VCSEL structures," *IEEE Photonics Technology Letters*, 11:1289-1291, 1999.
- [168] G. G. Ortiz, C. P. Hains, J. Cheng, H. Q. Hou, and J. C. Zolper, "Monolithic integration of $\text{In}_{0.2}\text{Ga}_{0.8}\text{As}$ vertical-cavity surface-emitting lasers with resonance-enhanced quantum well photodetectors," *Electronics Letters*, 32:1205-1207, 1996.
- [169] J. M. Mikkelson, Vitesse Semiconductor Corporation, Personal communication of thermally enhanced plastic packaging of 20 W Vitesse ICs, 1997.
- [170] T. M. Pinkston and C. Kuznia, "Smart-pixel-based network interface chip," *Applied Optics*, 36:4871-4880, 1997.
- [171] T. K. Woodward and A. V. Krishnamoorthy, "1-Gb/s integrated optical detectors and receivers on commercial CMOS technologies," *IEEE Journal of Selected Topics in Quantum Electronics*, 5:146-156, 1999.
- [172] D. Wiedenmann, R. King, C. Jung, R. Jager, R. Michalzik, P. Schnitzer, M. Kicherer, and K. J. Ebeling, "Design and analysis of single-mode oxidized VCSEL's for high-speed optical interconnects," *IEEE Journal of Selected Topics in Quantum Electronics*, 5:503-511, 1999.
- [173] R. Jager, M. Grabherr, C. Jung, R. Michalzik, G. Reiner, B. Weigl, and K. J. Ebeling, "57% wallplug efficiency oxide-confined 850 nm wavelength GaAs VCSEL's," *Electronics Letters*, 33:330-331, 1997.
- [174] E. R. Hegblom, N. M. Margalit, A. Fiore, and L. A. Coldren, "High-performance small vertical-cavity lasers: a comparison of measured improvements in optical and current confinement in devices using tapered apertures," *IEEE Journal of Selected Topics in Quantum Electronics*, 5:553-560, 1999.
- [175] B. J. Thibeault, K. Bertilsson, E. R. Hegblom, M. Strzelecka, P. D. Floyd, and L. A. Coldren, "High-speed characteristics of low-optical loss oxide-apertured vertical-cavity lasers," *IEEE Photonics Technology Letters*, 9:11-13, 1997.
- [176] A. E. Bond, P. D. Dapkus, and J. D. O'Brien, "Design of low-loss single-mode vertical-cavity surface-emitting lasers," *IEEE Journal of Selected Topics in Quantum Electronics*, 5:574-581, 1999.
- [177] D. L. Huffaker, L. A. Graham, H. Deng, D. G. Deppe, "Sub-40 μA continuous-wave lasing in an oxidized vertical-cavity surface emitting laser with dielectric mirrors," *IEEE Photonics Technology Letters*, 8:974-976, 1996.
-

-
- [178] K. D. Choquette and H. Q. Hou, "Vertical-cavity surface emitting lasers: moving from research to manufacturing," *Proceedings of the IEEE*, 85:1730-1739, 1997.
- [179] H. Q. Hou, R. P. Schneider, S. W. Corzine, B. E. Hammons, and S. Steward, "n-up 850-nm VCSEL arrays with very low and uniform operation voltage for multimode fiber data communications," *Proceedings of the 1998 Conference on Lasers and Electro Optics (CLEO)*, San Francisco, CA, May 1998, pp. 369-370, CThK5.
- [180] K. L. Lear, H. Q. Hou, J. J. Banas, B. E. Hammons, J. Furioli, and M. Osinski, "Vertical cavity lasers on p-doped substrates," *Electronics Letters*, 33:783-784, 1997.
- [181] V. M. Hietala, K. L. Lear, M. G. Armendariz, C. P. Tigges, H. Q. Hou, and J. C. Zolper, "Electrical characterization and application of very high speed vertical cavity surface emitting lasers (VCSELs)," *Proceedings of the 1997 IEEE MTT-S International Microwave Symposium*, Denver, CO, 1997, pp. 355-358.
- [182] K. L. Lear, M. Ochiai, V. M. Hietala, H. Q. Hou, B. E. Hammons, J. J. Banas, and J. A. Nevers, "High-speed vertical cavity surface emitting lasers," *Proceedings of the 1997 IEEE LEOS Summer Topical Meetings*, 1996, pp. 53-54.
- [183] M. Ochiai, K. L. Lear, V. M. Hietala, H. Q. Hou, and H. Temkin, "Modulation properties of high-speed vertical cavity surface emitting lasers," *Proceedings of the 1997 Device Research Conference*, 1997, pp. 102-103.
- [184] K. L. Lear, A. Mar, K. D. Choquette, S. P. Kilcoyne, R. P. Schneider, Jr., and K. M. Geib, "High-frequency modulation of oxide-confined vertical cavity surface emitting lasers," *Electronics Letters*, 32:457-458, 1996.
- [185] K. D. Choquette, K. L. Lear, R. P. Schneider, Jr., K. M. Geib, J. J. Figiel, and R. Hull, "Fabrication and performance of selectively oxidized vertical cavity lasers," *IEEE Photonics Technology Letters*, 7:1237-1239, 1995.
- [186] K. L. Lear, K. D. Choquette, R. P. Schneider, Jr., S. P. Kilcoyne, and K. M. Geib, "Selectively oxidized vertical cavity surface emitting lasers with 50% power conversion efficiency," *Electronics Letters*, 31:208-209, 1995.
- [187] K. D. Choquette, R. P. Schneider, Jr., K. L. Lear, and K. M. Geib, "Low threshold voltage vertical cavity lasers fabricated by selective oxidation," *Electronics Letters*, 30:2043-2044, 1995.
- [188] J. Heinrich, E. Zeeb, and K. J. Ebeling, "Butt-coupling efficiency of VCSEL's into multimode fibers," *IEEE Photonics Technology Letters*, 9:1555-1557, 1997.
- [189] J. Yuan, "A fiber optic probe," *6.161 Final Project*, MIT, 1997.
-

-
- [190] C. Kuznia, University of Southern California, Personal communication of loss in free-space optical systems, 1998.
- [191] M. Derstine, Optivision, Personal communication of loss in free-space optical systems, 1997.
- [192] E. A. Lee and D. G. Messerschmitt, *Digital Communication, Second Edition*, Kluwer Academic Publishers, Boston, 1994.
- [193] L. W. Couch, Jr., *Digital and Analog Communication Systems, Fourth Edition*, Macmillian Publishing Company, New York, 1993.
- [194] S. D. Personick, "Receiver design for digital fiber optic communication systems, I," and "Receiver design for digital fiber optic communication systems, II," *Bell System Technical Journal*, 52:843-885, 1973.
- [195] S. D. Personick, "Receiver design for optical fiber systems" *Proceedings of the IEEE*, 65:1670-1678, 1977.
- [196] R. G. Smith and S. D. Personick, "Receiver design for optical fiber communication systems", *Semiconductor Devices for Optical Communication*, H. Kressel, Ed., Topics in Applied Physics, Volume 39, pp. 89-160, Springer-Verlag, Berlin, 1980.
- [197] T. V. Muoi, "Receiver design for high-speed optical-fiber systems" *Journal of Lightwave Technology*, 2:243-267, 1984.
- [198] J. N. Hollenhorst, "Fundamental limits on optical pulse detection and digital communication" *Journal of Lightwave Technology*, 13:1135-1145, 1995.
- [199] T. H. Lee, *The Design of CMOS Radio Frequency Integrated Circuits*, pp. 167-171, Cambridge University Press, Cambridge, 1998.
- [200] W. C. Elmore, "The transient response of damped linear networks with particular regard to wideband amplifiers," *Journal of Applied Physics*, 19:55-63, 1948.
- [201] D. A. Van Blerkom, C. Fan, M. Blume, and S. C. Esener, "Transimpedance receiver design optimization for smart pixel arrays," *Journal of Lightwave Technology*, 16:119-126, 1998.
- [202] O. Kibar, D. A. Van Blerkom, C. Fan, and S. C. Esener, "Power minimization and technology comparisons for digital free-space optoelectronic interconnection," *Journal of Lightwave Technology*, 17:546-555, 1999.
- [203] H.-M. Rein and M. Moller, "Design considerations for very-high-speed Si-bipolar IC's operating up to 50 Gb/s," *IEEE Journal of Solid-state Circuits*, 31:1076-1090, 1996.
- [204] H. Rein, "Multi-gigabit-per-second silicon bipolar IC's for future optical fiber transmission systems," *IEEE Journal of Solid-state Circuits*, 23:664-675, 1988.
-

-
- [205] K. Runge, D. Daniel, R. D. Standley, J. L. Gimlett, R. B. Nubling, R. L. Pierson, S. M. Beccue, K. Wang, N. Sheng, M. F. Chang, D. M. Chen, and P. M. Asbeck, "AlGaAs/GaAs HBT IC's for high-speed lightwave transmission systems," *IEEE Journal of Solid-state Circuits*, 27:1332-1341, 1992.
- [206] Z. Wang, M. Berroth, U. Nowotny, M. Ludwig, P. Hofmann, A. Hulsmann, K. Kohler, B. Raynor, and J. Schneider, "Integrated laser-diode voltage driver for 20-Gb/s optical system using 0.3- μ m gate length quantum-well HEMT's," *IEEE Journal of Solid-state Circuits*, 28:829-834, 1993.
- [207] J. Riishoj, "5 Gbit/s laser-driver GaAs IC," *Electronics Letters*, 30:774-776, 1994.
- [208] H-GaAs IV Foundry Design Manual, Release 1.0, Vitesse Semiconductor Corporation, 1998.
- [209] R. N. Deming, Vitesse Semiconductor Corporation, Personal communication regarding optimal source/drain length and usable range of amplifier gain.
- [210] S_{11} measurements of VCSEL performed by A. Ramamurthy, Northeastern University, 1998.
- [211] J. M. Mikkelson, Vitesse Semiconductor Corporation, Personal communication regarding frequency dependence of MESFET drain conductance, 1998.
- [212] L. A. Coldren and S. W. Corzine, *Diode Lasers and Photonic Integrated Circuits*, pp. 217-219, John Wiley & Sons, Inc., New York, 1995.
- [213] A. Tanabe, M. Soda, Y. Nakahara, T. Tamura, K. Yoshida, and A. Furukawa, "A single-chip 2.4-Gb/s CMOS optical receiver IC with low substrate cross-talk preamplifier," *IEEE Journal of Solid-State Circuits*, 33:2148-2153, 1998.
- [214] M. Nakamura, N. Ishihara, and Y. Akazawa, "A 156-Mb/s CMOS optical receiver for burst-mode transmission," *IEEE Journal of Solid-State Circuits*, 33:1179-1187, 1998.
- [215] S. M. Park and C. Toumazou, "Low noise current-mode CMOS transimpedance amplifier for gigabit optical communication," *Proceedings of the 1998 IEEE International Symposium on Circuits and Systems (ISCAS '98)*, 31 May - 3 June 1998, Vol. 1, pp. 293-296.
- [216] S. M. Park and C. Toumazou, "Gigahertz low noise CMOS transimpedance amplifier," *Proceedings of the 1997 IEEE International Symposium on Circuits and Systems (ISCAS '97)*, 9-12 June 1997, Vol. 1, pp. 209-212.
- [217] C. Toumazou and S. M. Park, "Wideband low noise CMOS transimpedance amplifier for gigaHertz operation," *Electronics Letters*, 32:1194-1196, 1996.
- [218] C. C. Cadenas and F. P. Martinez, "New analysis technique of transimpedance amplifiers with cascode structure: applications to a DC-2.5 GHz MMIC amplifier," *Proceedings of the 7th IEEE Mediterranean Electrotechnical Conference*, 12-14 April 1994, Vol. 2, pp. 613-616.
-

- [219] J. A. Casao-Perez, P. Dorta, F. Perez, and J. Perez, "MMIC transimpedance amplifier for an optical front-end with high gain and minimum ripple," *Proceedings of the 1994 IEEE MTT-S International Microwave Symposium*, 1994, Vol. 2, pp. 241-244.
- [220] J. A. Casao, P. Dorta, J. L. Caceres, M. Salazar-Palma, and J. Perez, "An enhanced GaAs monolithic transimpedance amplifier for low noise and high speed optical communications," *Proceedings of the 1992 IEEE MTT-S International Microwave Symposium*, 1992, Vol. 1, pp. 85-88.
- [221] N. Schienberg, R. J. Bayruns, and T. M. Laverick, "Monolithic GaAs transimpedance amplifiers for fiber-optic receivers," *IEEE Journal of Solid-State Circuits*, 26:1834-1839, 1991.
- [222] T. K. Woodward, A. V. Krishnamoorthy, A. L. Lentine, and L. M. Chirovsky, "Optical receivers for optoelectronic VLSI," *IEEE Journal of Selected Topics in Quantum Electronics*, 2:106-116, 1996.
- [223] A. Van Der Ziel, *Noise: Sources, Characterization, Measurement*, pp. 69-81, Prentice-Hall, Englewood Cliffs, New Jersey, 1970.
- [224] W. Beachtold, "Noise behavior of GaAs field-effect transistors with short gate lengths," *IEEE Transactions on Electron Devices*, 19:674-680, 1972.
- [225] T. Witkowicz, "Design of low-noise fiber-optic receiver amplifiers using J-FETs," *IEEE Journal of Solid State Circuits*, 13:195-197, 1978.
- [226] T. Witkowicz, "Transistor noise model for photodetector amplifiers," *IEEE Journal of Solid State Circuits*, 13:722-724, 1978.
- [227] K. Drogemuller, "PAROLI: a synchronous interconnection link with a throughput of 13 Gbit/s," *Workshop Notes for the 10th Annual Workshop on Interconnects Within High Speed Digital Systems (IHSDS)*, Santa Fe, NM, 9-12 May 1999, 1.2, email: karsten.drogemueller@infion.com.
- [228] R. Landauer. "Need for critical assessment," *IEEE Transactions on Electron Devices*, 43(10):1637-1639, 1996.
- [229] S. I. Long and S. E. Butner, *Gallium Arsenide Digital Integrated Circuit Design*, McGraw-Hill Publishing Company, New York, 1990.
- [230] O. Wing, *Gallium Arsenide Digital Circuit*, McGraw-Hill Publishing Company, New York, 1990.
- [231] C. Toumazou and D. Haigh, "Gallium arsenide analogue integrated circuit design techniques", *Analogue IC Design: the Current-Mode Approach*, C. Toumazou, F. J. Lidgey, and D. G. Haigh, Eds., pp. 327-369, Peter Peregrinus Ltd., London, 1990.
- [232] D. P. Hornbuckle and R. L. Van Tuyl, "Monolithic GaAs direct-coupled amplifiers," *IEEE Transactions on Electron Devices*, ED-28:175-182, 1981.

-
- [233] M. Shur, T. A. Fjeldly, T. Ytterdal, and K. Lee, "Unified GaAs MESFET model for circuit simulation," *International Journal of High Speed Electronics*, 3:201-233, 1992.
- [234] T. Ytterdal, B. Moon, T. A. Fjeldly, and M. S. Shur, "Enhanced GaAs MESFET CAD model for a wide range of temperatures," *IEEE Transactions on Electron Devices*, 42:1724-1734, 1994.
- [235] Foundry Design Manual, Version 6.0, Vitesse Semiconductor Corporation, 1993. Note: This is the H-GaAs III Foundry Design Manual.
- [236] P. Alfke, "Efficient shift registers, LFSR counters, and long pseudorandom sequence generators," Application Note XAPP 052, July 7, 1996, Ver. 1.1, Xilinx, Inc., 2100 Logic Drive, San Jose, CA, 95124, 1-800-255-7778, <http://xilinx.com/>.
- [237] Alberta Printed Circuits, Unit 3, 1112-40th Ave. NE, Calgary, AB, Canada T2E 5T8, (403)250-3406, <http://www.apcircuits.com/>.
- [238] P-S. Lin and C. A. Zukowski, "Jitter due to signal history in digital logic circuits and its control strategies," *Proceedings of the 1993 IEEE International Symposium on Circuits and Systems (ISCAS)*, New York, May 1993, Vol. 3, pp. 2114-2117.
- [239] D. Middleton, *An Introduction to Statistical Communication Theory*, McGraw-Hill Publishing Company, New York, 1960.
- [240] D. K. Campbell, "First zero passage time statistics of a ramp plus noise," *M.S. Thesis*, Massachusetts Institute of Technology, 1968.
- [241] B. W. Stuck, "Switching-time jitter statistics of bipolar transistor threshold-crossing detectors," *M.S. Thesis*, Massachusetts Institute of Technology, 1968.
- [242] A. A. Abidi and R. G. Meyer, "Noise in relaxation oscillators," *IEEE Journal of Solid-State Circuits*, SC-18:794-802, 1983.
- [243] T. C. Weigandt, B. Kim, and P. R. Gray, "Analysis of timing jitter in CMOS ring oscillators," *Proceedings of the 1994 IEEE International Symposium on Circuits and Systems (ISCAS)*, June 1994, pp. 27-30.
- [244] J. A. McNeil, "Jitter in ring oscillators," *IEEE Journal of Solid-State Circuits*, 32:870-879, 1997.
- [245] F. Herzel and B. Razavi, "A study of oscillator jitter due to supply and substrate noise," *IEEE Transactions on Circuits and Systems--II: Analog and Digital Signal Processing*, 46:56-62, 1999.
- [246] A. Hajimiri, S. Limotyrakis, and T. H. Lee, "Jitter and phase noise in ring oscillators," *IEEE Journal of Solid-State Circuits*, 34:790-804, 1997.
- [247] J. A. Phillips and S. J. Harrold, "GaAs sampled analogue integrated circuits", *GaAs Integrated Circuits*, J. Mun, pp. 353-356, BSP Professional Books, Oxford, 1988.
-

-
- [248] C. Camacho-Penalosa and C. S. Aitchison, "Modelling frequency dependence of output impedance of a microwave MESFET at low frequencies," *Electronics Letters*, 21:528-529, 1985.
- [249] J. M. Golio, M. G. Miller, G. N. Maracas, and D. J. Johnson, "Frequency-dependent electrical characteristics of GaAs MESFET's," *IEEE Transactions on Electron Devices*, 37:1217-1227, 1990.
- [250] N. Scheinberg, R. Bayruns, and R. Goyal, "A low-frequency GaAs MESFET circuit model," *IEEE Journal of Solid State Circuits*, 23:605-608, 1988.
- [251] N. Scheinberg, R. J. Bayruns, P. W. Wallace, and R. Goyal, "An accurate MESFET model for linear and microwave circuit design," *IEEE Journal of Solid State Circuits*, 24:532-539, 1989.
- [252] C. Canfield, J. Medinger, and L. Forbes, "Buried-channel GaAs MESFET's with frequency-independent output conductance," *IEEE Electron Device Letters*, EDL-8:88-89, 1987.
- [253] S. Choi and M. B. Das, "Origin and modeling of the frequency dependent output conductance in microwave GaAs MESFET's with buried p layer," *IEEE Transactions on Electron Devices*, 41:1725-1733, 1994.
- [254] L. E. Larson, G. C. Temes, and S. Law, "Comparison of amplifier gain enhancement techniques for GaAs MESFET analogue integrated circuits," *Electronics Letters*, 22:1138-1139, 1986.
- [255] W. S. Lee and J. Mun, "Improved negative feedback technique to reduce drain conductance of GaAs MESFETs for precision analogue ICs," *Electronics Letters*, 23:705-707, 1987.
- [256] P. R. Gray and R. G. Meyer, *Analysis and Design of Analog Integrated Circuits, Third Edition*, John Wiley & Sons, New York, 1993, pp. 535-596.
- [257] M. Mony and D. Kabal, McGill University, Generic Test Board Documentation, 1997.
- [258] E. M. Cherry and D. E. Hooper, "The design of wide-band transistor feedback amplifier," *Proceedings of the IEE*, 110:375-389, 1963.
- [259] T. K. Woodward, A. L. Lentine, K. W. Goosen, J. A. Walker, B. T. Tseng, S. P. Hui, and R. E. Leibenguth, "Demultiplexing 2.48-Gb/s optical signals with a CMOS receiver array based on clocked-sense-amplifiers," *IEEE Photonics Technology Letters*, 9:1146-1148, 1997.
- [260] J. A. B. Dines, "Smart pixel optoelectronic receiver based on a charge sensitive amplifier design," *IEEE Journal of Selected Topics in Quantum Electronics*, 2:117-120, 1996.
- [261] L. A. Glasser and D. W. Dobberpuhl, *The Design and Analysis of VLSI Circuits*, Addison-Wesley Publishing Company, Reading, Massachusetts, 1985.
-