CIRCUITS FOR THE REALIZATION OF HIGH SPEED SERIAL TRANSMITTERS USING SIGE HBTS

By

Peter F. Curran

A Thesis Submitted to the Graduate

Faculty of Rensselaer Polytechnic Institute

in Partial Fulfillment of the

Requirements for the Degree of

DOCTOR OF PHILOSOPHY

Major Subject: Electrical Engineering

Approved by the
Examining Committee:

_____
John F. McDonald, Thesis Adviser

_____
Yannick L. LeCoz, Member

_____
Shivkumar Kalyanaraman, Member

_____
Mukkai S. Krishnamoorthy, Member

Rensselaer Polytechnic Institute
Troy, New York

August 2005
(For Graduation August 2005)

CIRCUITS FOR THE REALIZATION OF HIGH SPEED SERIAL TRANSMITTERS USING SIGE HBTS

By

Peter F. Curran

An Abstract of a Thesis Submitted to the Graduate

Faculty of Rensselaer Polytechnic Institute

in Partial Fulfillment of the

Requirements for the Degree of

DOCTOR OF PHILOSOPHY

Major Subject: Electrical Engineering

The original of the complete thesis is on file
in the Rensselaer Polytechnic Institute Library

Examining Committee:

John F. McDonald, Thesis Adviser
Yannick L. LeCoz, Member
Shivkumar Kalyanaraman, Member
Mukkai S. Krishnamoorthy, Member

Rensselaer Polytechnic Institute
Troy, New York

August 2005
(For Graduation August 2005)

© Copyright 2005
by
Peter F. Curran
All Rights Reserved

CONTENTS

LIST OF TABLES iv

LIST OF FIGURES v

ACKNOWLEDGMENT vi

ABSTRACT vii

1. Introduction & History 1

    1.1 SERDES Overview 2

        1.1.1 Serial Communication Systems 2

            1.1.1.1 Transmitter 3

            1.1.1.2 Receiver 5

            1.1.1.3 Formatting/Encoding 7

            1.1.1.4 Framing 8

    1.2 SERDES Research Project at RPI 9

        1.2.1 Research Timeline/Description 9

    1.3 State of the Art 12

    1.4 Contributions to the Field 14

    1.5 Overview of this document 15

2. Previous Work 16

    2.1 SERDES I 16

        2.1.1 Current Starving VCO 17

        2.1.2 Transmitter 18

            2.1.2.1 Data “Shuffling” Scheme 18

            2.1.2.2 LFSR 22

        2.1.3 Receiver 25

        2.1.4 Layout 29

        2.1.5 Design for Testing 29

    2.2 SERDES II 32

        2.2.1 Symmetric Multiplexer 32

        2.2.2 Leap & FFI VCO 34

        2.2.3 SERDES II Testing Scheme 37

2.2.4 Transmitter . . . 39
2.2.4.1 Transmitter LFSR . . . 40
2.2.5 Receiver . . . 42
2.2.5.1 4-to-16 Demultiplexer . . . 43
2.2.5.2 Receiver LFSR . . . 44
2.2.6 SERDES II Receiver/Transmitter Chip Layout . . . 46
2.2.6.1 Results . . . 47
2.2.7 Second Prototype Corrections . . . 48
3. Voltage Controlled Oscillators . . . 55
3.1 5DM Dual Control VCO . . . 55
3.1.1 Design . . . 56
3.1.2 Layout . . . 59
3.1.3 Testing . . . 61
3.2 SIII 5AM Oscillator . . . 63
3.2.1 Design . . . 63
3.2.2 Layout . . . 66
3.2.3 Performance . . . 67
3.3 Second Revision 5AM Oscillator . . . 69
3.3.1 Layout . . . 70
3.3.2 Performance . . . 71
3.4 Conclusions . . . 74
4. Edge Steering Multiplexer Architecture . . . 75
4.1 SERDES III Symmux . . . 76
4.2 Edge Steering Multiplexer . . . 81
5. Multiphase LFSR . . . 86
5.1 Multiphase Design . . . 86
5.2 Test Clock . . . 90
5.3 Problems . . . 92
6. LFSR . . . 93
6.1 Bit Pattern . . . 93
6.2 State Machine Generation . . . 94
6.3 Advance-by-16 implementation . . . 95
6.4 Layout . . . 98

7. 16-to-4 Round-Robin Multiplexer . . . 100
7.1 Quad 4-bit Shift Register . . . 101
7.2 16-to-4 Multiphase Multiplexer . . . 104
7.2.1 Hidden state . . . 106
7.2.2 Solution . . . 108
8. Transmitter . . . 114
8.1 ESMA Implementation Details . . . 115
8.2 SERDES III Transmitter Chip . . . 119
8.3 Testing . . . 121
8.4 SERDES III Transmitter Revision . . . 125
8.4.1 PLL . . . 125
8.4.2 Layout . . . 126
8.4.3 Simulations . . . 127
9. Performance Limitations . . . 129
9.1 Multiplexer Overview . . . 129
9.2 Switching Noise . . . 131
9.3 Weak Switching Distortion . . . 132
9.4 Multiplexer Bandwidth . . . 134
9.5 Pass Through . . . 136
9.6 Technology Comparison . . . 137
9.7 Bandwidth Improvements . . . 138
9.8 Eye Comparisons . . . 141
10. Conclusion . . . 143
10.1 Future Work . . . 144
LITERATURE CITED . . . 145
APPENDICES
A. Technology . . . 150
A.1 Initial Choices . . . 150
A.2 IBM SiGe 5HP BiCMOS Process . . . 150
A.2.1 The SiGe HBT . . . 150
A.3 Resistors . . . 155

A.4 Interconnect . . . 156
A.4.0.1 Capacitors . . . 157
A.4.0.2 Inductors . . . 157
A.5 Other Components . . . 158
A.6 Logic . . . 158
A.7 CML Gates . . . 159
A.8 Current Sources . . . 167
A.9 Performance . . . 169
A.10 High Frequency Response . . . 172
A.10.1 Linearized Differential Buffers . . . 174
A.11 Pad Drivers . . . 177
A.12 Power Rails . . . 180
B. LFSR . . . 183
B.1 Linear Feedback Shift Registers . . . 183
B.1.1 Characteristic Polynomial . . . 184
B.1.2 Sequence Characteristics . . . 186
C. Clocking Components . . . 188
C.1 Phase Locked Loops . . . 188
C.2 Phase Detectors . . . 189
C.3 Loop Filters . . . 192
C.4 VCOs . . . 193
C.5 Jitter . . . 196
C.6 Phase Noise . . . 199
C.7 Clock Recovery . . . 201
C.8 Data Retiming . . . 201
D. Source Code . . . 203
D.1 bits31.c . . . 203
E. Equipment Specifications . . . 205
E.1 Signal Source SML01 . . . 205
E.2 Probes . . . 205
E.3 Oscilliscope 11801C . . . 206

LIST OF TABLES

3.1 VCO Characteristics 74

A.1 Metal Interconnect Characteristics 156

B.1 Selected LFSR coefficients 186

LIST OF FIGURES

1.1 Serial Communication System 2
1.2 SERDES I Transmitter Block Diagram 3
1.3 SERDES I Receiver Block Diagram 6
1.4 NRZ .vs. RZ Data Formats 7
1.5 Research Timeline 10
2.1 Ring Oscillators 17
2.2 SERDES quarter-clock shuffle multiplexing 19
2.3 Block diagram of the SERDES I Transmitter 20
2.4 SERDES I LFSR Schematic 22
2.5 SERDES I LFSR Layout 23
2.6 SERDES I Transmitter Control 23
2.7 SERDES I Control Layout 24
2.8 Block diagram of the SERDES I receiver 25
2.9 Receiver CDR Operation 26
2.10 Receiver Sample and Detection Schematic 27
2.11 Sampling Block Schematic 28
2.12 Receiver Sample and Detection Layout 29
2.13 Proportional Current Generator 30
2.14 SERDES I Integrator Layout 30
2.15 Microphotograph of SERDES I 31
2.16 Symmux Circuit 33
2.17 Duty cycle plot of Symmux 33
2.18 Symmux Layout 34
2.19 “Leapfrog” interpolation plot 35

2.20 Feed Forward Interpolated VCO . . . 36
2.21 FFI VCO control plot. . . . 37
2.22 Performance of the PLL . . . 38
2.23 Layout of the FFI VCO . . . 38
2.24 SERDES II Transmitter Block Diagram . . . 40
2.25 SERDES II Transmitter LFSR Schematic . . . 40
2.26 SERDES II Transmitter LFSR Layout . . . 41
2.27 SERDES II Receiver Block Diagram . . . 42
2.28 SERDES II Receiver 4-to-16 Demultiplexer . . . 43
2.29 SERDES II Demultiplexer Control . . . 44
2.30 SERDES II Receiver 4-to-16 Demultiplexer Layout . . . 44
2.31 SERDES II Receiver LFSR Schematic . . . 45
2.32 SERDES II LFSR High Fan-in AND . . . 45
2.33 SERDES II Receiver LFSR Layout . . . 46
2.34 SERDES II Prototype Layout . . . 47
2.35 SERDES II 20Gbps Eye Diagram . . . 48
2.36 Dual Eye Diagram. . . . 49
2.37 Simulated edge position plot . . . 50
2.38 Plot showing edge placement errors . . . 52
2.39 Radial phase plot . . . 53
2.40 Simulated 40Gbps SERDES II output . . . 54
3.1 5DM VCO Architecture . . . 56
3.2 Test VCO element fabricated in 5DM . . . 57
3.3 Variable voltage reference . . . 58
3.4 5DM Voltage Control Plot . . . 59
3.5 10 GHz test VCO Layout . . . 60

3.6 5DM 10GHz VCO test chip . . . 61
3.7 Microphotograph of the 5DM 10GHz VCO . . . 61
3.8 5DM VCO voltage.vs.frequency plot . . . 62
3.9 5DM 10GHz VCO Output . . . 63
3.10 5AM VCO Architecture . . . 64
3.11 5AM VCO Delay Element . . . 64
3.12 5AM VCO Characteristic . . . 65
3.13 5AM VCO Phase Noise . . . 66
3.14 5AM VCO Layout . . . 67
3.15 5AM VCO at 50C . . . 68
3.16 5AM 10GHz VCO Tuning . . . 69
3.17 5AM 10GHz VCO Buffer . . . 70
3.18 5AM 10GHz VCO Layout . . . 71
3.19 10GHz 5AM VCO with Parasitics . . . 72
3.20 7.5GHz 5AM VCO with Parasitics . . . 73
4.1 Four to one architecture . . . 76
4.2 Symmetric 2-to-1 multiplexer . . . 77
4.3 Symmux Gain . . . 78
4.4 Symmux 40Gbps Output . . . 79
4.5 Potential edges in serial stream . . . 81
4.6 First Tier Multiplexers . . . 81
4.7 Proposed new multiplexer architecture . . . 82
4.8 Eight phase clocks . . . 83
4.9 Steered edge architecture multiplexer outputs . . . 84
5.1 Test generation architecture . . . 87
5.2 State dependency diagram . . . 87

5.3 Combined XOR and MS latch . . . 88
5.4 Multiphase LFSR schematic . . . 89
5.5 Test generation circuitry . . . 90
5.6 Divide-by-4095 chain . . . 91
5.7 Divide-by-three counter circuit . . . 91
6.1 Schematic of LFSR . . . 94
6.2 Next state functions . . . 95
6.3 LFSR Building Block . . . 96
6.4 LFSR full row schematic . . . 96
6.5 SIII LFSR schematic . . . 97
6.6 31-bit LFSR layout . . . 99
7.1 Slow Clocking Errors . . . 101
7.2 Quad 4-bit Shift Register . . . 102
7.3 Shift Register Details . . . 103
7.4 Shift Register Timing Diagram . . . 103
7.5 Latched 4-to-1 Multiplexer . . . 104
7.6 CML Latched 4-to-1 Multiplexer . . . 105
7.7 CML Divide By Two . . . 106
7.8 Distributed divide-by-two Clock . . . 108
7.9 Distributed divide-by-two Plus Sampling . . . 109
7.10 Clock Division Details . . . 109
7.11 Schematic of 4-to-1 Multiplexer . . . 110
7.12 LFSR Clock Generation . . . 111
7.13 Schematic of 16-to-4 Multiplexer . . . 112
7.14 Complete Layout of 16-to-4 Multiplexer . . . 113
8.1 Transmitter Block Diagram . . . 114

8.2 Edge Steering Multiplexer Schematic . . . 115
8.3 Top Tier 2-to-1 multiplexer . . . 116
8.4 Top tier 2-to-1 emitter followers . . . 116
8.5 ESMA Final multiplexer schematic . . . 117
8.6 ESMA Layout for SERDES III Transmitter chip . . . 118
8.7 SERDES III Transmitter Schematic . . . 119
8.8 SERDES III Transmitter Layout . . . 120
8.9 SERDES III Transmitter Microphotograph . . . 121
8.10 LFSR 1.6Gbps Eye Output . . . 122
8.11 Transmitter 27Gbps Eye Output . . . 122
8.12 Captured LFSR waveform . . . 123
8.13 Multiple Eye Output . . . 124
8.14 Revised Transmitter Block Diagram . . . 125
8.15 Loop filter layout . . . 126
8.16 Revised Transmitter Layout . . . 126
8.17 Revised Transmitter Simulations . . . 128
9.1 Standard CML Multiplexer . . . 130
9.2 Switching noise . . . 131
9.3 Simulated switching noise . . . 132
9.4 Single Ended Multi-Eye Plot . . . 133
9.5 Small Signal Multiplexer Bandwidth . . . 134
9.6 Large Signal Multiplexer Bandwidth . . . 135
9.7 Multiplexer Signal Leakage . . . 136
9.8 Multiplexer as part of a D-Latch . . . 138
9.9 Generic Inductive Peaking Circuit . . . 139
9.10 Transformer Layout . . . 139

9.11 Large Signal Inductive Peaking Sweep . . . 140
9.12 Large Signal Inductive Coupling Sweep . . . 141
A.1 Layout of the 5HP SiGe HBT . . . 151
A.2 Spectre VBIC Model . . . 152
A.3 Gummel plot of SiGe HBT . . . 153
A.4 SiGe HBT forward active curve . . . 154
A.5 SiGe HBT f_T plot . . . 154
A.6 CML buffer . . . 159
A.7 Current Steering . . . 161
A.8 Normalized VTC . . . 162
A.9 Noise Margin Plot . . . 163
A.10 CML gate with two inputs . . . 164
A.11 CML D-Latch . . . 165
A.12 CML gate with two inputs . . . 165
A.13 Voltage levels . . . 166
A.14 Current Mirror Schematic . . . 167
A.15 Interconnect.vs.buffer performance . . . 170
A.16 Interconnect .vs. buffer rise/fall times . . . 170
A.17 Gate delay reduction . . . 171
A.18 Emitter follower impact on rise/fall times . . . 171
A.19 Large signal bandwidth of buffer . . . 173
A.20 CML buffer phase response . . . 173
A.21 Linearized buffer circuit . . . 174
A.22 Inverter transfer characteristic . . . 175
A.23 Linearized buffer transfer characteristic . . . 176
A.24 Single-ended 1GHz pad driver . . . 177

A.25 Single-ended DC pad receiver . . . 178
A.26 Analog pad receiver . . . 179
A.27 High-speed differential pad receiver circuit . . . 180
A.28 Power rail dimensions and characteristics . . . 181
B.1 Fibonacci LFSR . . . 183
B.2 Galois LFSR . . . 183
B.3 Galois 6th order {6,4,3} LFSR . . . 185
C.1 Phase locked loop . . . 189
C.2 XOR phase detector . . . 190
C.3 Three state phase detector . . . 191
C.4 Active loop filter . . . 192
C.5 Ring Oscillators . . . 194
C.6 "Forward leap" architecture . . . 195
C.7 Types of jitter measurement . . . 196
C.8 Sampling a digital waveform . . . 199

ACKNOWLEDGMENT

I would like to dedicate this work to my family, and in particular to my father, who has given me unfailing encouragement and support for many, many years. Also I would like to express my gratitude to my advisor, Prof. John F. McDonald, for providing an environment in which students can aspire to achieve. My group of friends and fellow students at RPI, too numerous to mention here individually, are also in part responsible for my success. Thanks also to Hank Dardy of NRL who provided funding for this project for many years, and to the MOSIS integrated fabrication service, who provided cost-free fabrication for several designs under their MEP program.

ABSTRACT

Serializer and Deserializer circuits (SERDES) are important building blocks for communications systems in both long distance optical networks as well as for high speed connections over distances as small as the length of a circuit board. The focus of this work is in the area of high speed SERDES, particularly in the area of serial transmitter circuits which allow high data rates with respect to fundamental limits inherent in device parameters such as f_T and f_{max}. Using the IBM SiGe 5HP 50GHz f_T/f_{max} BiCMOS process and CML logic, our earlier designs achieved 20Gbps data rates. SiGe was chosen for its low cost to performance ratio, but the design architectures are not process specific. All our designs are monolithic, relying on no external tuning components and feature internal VCOs. The high rate of operation is made possible by using quarter-rate multiphase clocks and novel multiplexer architectures which reduce the need for final retiming. These designs were fabricated and tested and the results compared with simulation. Newer architechtures subsequently developed have shown promise for even higher speed operation in the same process. A completely new transmitter chip has been designed, incorporating an internal pseudo-random number generator, a novel 16-to-4 multiphase multiplexer, and an entirely new final 4-to-1 edge steering multiplexer architecture (ESMA). The 16-to-4 multiplexer addresses difficult design constraints on the synchronization of multiple clock phases along with those phases after having passed through clock dividers. The ESMA 4-to-1 multiplexer allows the generation of a serial stream in a symmetrical manner using four quarter-rate clock phases. Heavy emphasis was placed on optimization for speed and symmetry in order to produce the highest rate output. A new test chip with 1825 HBTs using these designs has been fabricated and tested up to 27Gbps. A second improved prototype has also been designed, and simulation results indicate the potential for greater than 30Gbps operation, with some corresponding reduction in output quality. This improved prototype has not yet been fabricated, but it is planned for in the future.

CHAPTER 1
Introduction & History

Serializer and deserializer (SERDES) circuits are now ubiquitous in high speed digital electronics. They are useful in the transmission of data through optical media for long and short distances, as well as through coaxial and twisted pair cables for networking. SERDES circuits are even useful for short chip-to-chip interconnections in order to reduce pin and trace counts. Increases in data rates can have the direct benefits of either reducing the quantity of required physical resources, (such as pins and traces), or increasing the utility of existing resources, (such as improving the throughput of a large scale fiber optic infrastructure).

Improving data rates can be achieved through several means. One might radically change the circuit technology, such as by moving to compound semiconductors and/or exotic transistor devices such as HEMTs. Performance can also be improved within a technology by improving devices, perhaps by modifying device geometry or impurity concentrations. Lastly, new circuit topologies can be created which increase data rates while keeping the base technology unchanged.

In many cases, it may be more cost effective to stay with a known and proven manufacturing technology over one with superior performance. Yield issues and higher costs make many compound semiconductor processes unsuitable apart from all but the most demanding applications. Given the prevalence of SERDES circuits, improvements in data rates achieved in silicon processes are to be preferred for reasons of integration. Advances through new circuit designs provide the greatest potential for benefit as they can often benefit designs using future advances of devices within the base technology.

This research has focused on circuit design in IBM's well established 5HP SiGe Family of BiCMOS processes. Using innovative designs, circuits have been produced with data rates higher than those reported by others in the literature using similar processes. This has been accomplished by using multiphase clocks produced by ring oscillators rather than a single or quadrature set of phases. Using

this technique, and by designing circuits which eliminate or greatly reduce the need for data retiming, serial transmitter and receiver circuits have been created with the potential to work at a large fraction of the unity-gain cutoff frequency, (f_T), a theoretical upper operating limit of the HBT devices.

1.1 SERDES Overview

1.1.1 Serial Communication Systems

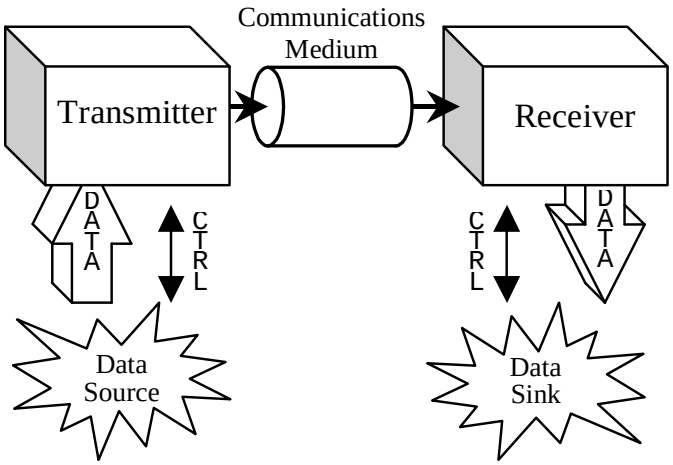

Figure 1.1: Serial Communication System

Most high-speed serial communications systems in their simplest form can be broken down into 3 major components, (see Figure 1.1). These components are the Transmitter, the Communications Medium, and the Receiver.

First we have the Transmitter, which typically takes in parallel data from a data source and then outputs it serially via multiplexing or via a shift register. The transmitter usually requires some form of additional communication with the data source in order to ensure that correct data is read. The serial data, typically with no additional clocking information, is sent out along the communications medium. Because no separate clock signal is sent, the data itself must be used to recover a clock.

The communication channel consists of some material medium, as well as the electronics necessary to convert incoming/outgoing electrical signals. The actual signals transmitted over the medium may use an entirely different data encoding than

the incoming or outgoing electrical signals. Phase Shift Keying (PSK) is a common encoding. Some channel interfaces might require single-ended signals as opposed to differential. Certain types of channels may put constraints on the outgoing data such as the requirement of having a zero average DC component, necessitating the transmission of as many zeros as ones. Channels may require drivers, receivers, and often repeaters, which influence the final signal that arrives at the Receiver. Losses along the channel and SNR ultimately limit performance. A single communication channel might be used by different independent serial streams simultaneously, as in the case of wave division multiplexing (WDM). In the context of this document, the communications medium represents the physical medium, the repeaters, and drivers necessary to communicate with the transmitter and receiver electrically, and will not be further addressed.

At the Receiver the incoming signal is examined, bit values extracted, and the data is output again in a parallel form. As with the transmitter, the transfer of the parallel data usually requires some form of additional communication with a data sink to transfer the information reliably. The Receiver must accurately recover the data after it has passed through the channel, and to do this it usually has to somehow extract clocking information from the serial data stream.

1.1.1.1 Transmitter

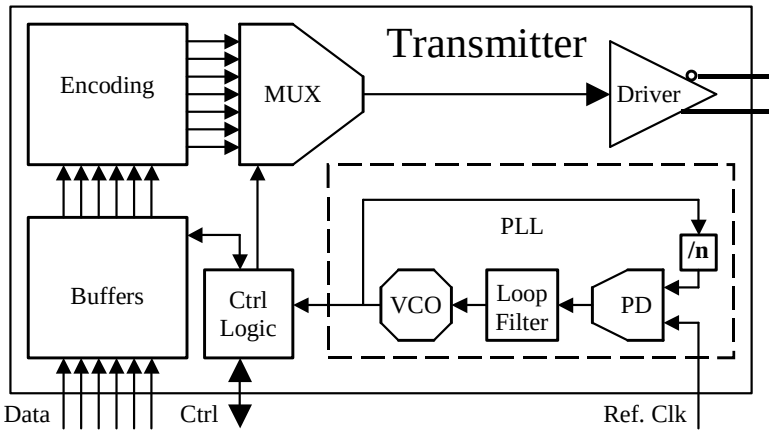

Figure 1.2: SERDES I Transmitter Block Diagram

Figure 1.2 shows a simplified representation of a Transmitter. There must be

some form of communication, (in this case labeled "Ctrl"), between the Transmitter and the source of the parallel data to ensure the data at the inputs is valid when required by the Transmitter. Depending on the application, asynchronous parallel input to the transmitter might be desired and implemented using buffering. Control lines indicating FIFO status are common. The parallel data bits can be from a single source or from multiple sources such as in time division multiplexing (TDM). Encoding before transmission is done on-chip when necessary. In certain cases, the communication at each of the parallel lines is itself implemented as a complete SERDES system at a lower data rate. The transmitter requires drivers to send the signals off chip, meeting signal levels and matching impedance characteristics. The data at the transmitter output may be transmitted continuously over the communication medium or it may be broken up into packets. Packets are often used in a shared channel, along with some mechanism for detecting resource-sharing conflicts.

The transmitter is responsible for the accurate timing of the output signal, without which the receiver on the far side of the communications medium cannot operate. The transmitter thus requires an accurate reference clock, often generated off chip or using off chip components such as crystals. When the reference clock is generated off chip, the transmitter may also generate a clock signal internally which is synchronized with the external clock using a phase locked loop (PLL). This internal clock may run at many multiples of the external reference and may also be multi-phase.

Because all signals contain noise, there must be sufficient noise margins in the output signal to accurately differentiate between a one and a zero bit to a large degree of certainty. The inherent component noise is usually of an unbounded nature, such as gaussian, so the change of error due to signal noise is never identically zero. By making the noise margin five times or more the standard deviation of the noise, bit errors from this effect can be reduced to one in 10^{10} or lower.

Another important characteristic of the output is the rise, (or fall), time of the signal as it transitions between bit values. This slew rate determines the spectral distribution of the output, which has a $sync^2()$ function as its envelope. Perfect generation or transmission of a signal would require infinite bandwidth, but a rough

rule of thumb is that a bandwidth of $0.5/T_{rise}$ can pass enough of the signal to preserve the pulse shape well enough for digital recovery. In a very high-speed design however, the rise time can equal the bit time, which means that the bandwidth can also be related to the fundamental frequency of the serial stream, which is half the bit rate. Thus a 10Gbps serial stream which has a 5GHz fundamental frequency requires at least a 5GHz bandwidth for transmission. If the edges are sharper, as is almost always the case, a greater bandwidth is required. With too little bandwidth, the pulse shape becomes distorted and the chances of detecting the wrong bit value increases. This minimum requirement is generally given for signal transmission. To create a good signal, generally much higher bandwidths are required.

One last important characteristic of the output is called jitter. Jitter is unwanted variation in the signal timing, which affects the placement of the edges. Jitter can arise from several mechanisms which in some cases produce unbounded jitter, while in others produces jitter with a hard limit. Various mechanisms allow the effects of signal noise to manifest themselves as jitter. Because the receiver uses the edges to determine when to sample the data bit, large amounts of jitter can cause the sample to occur at the wrong bit, or during the next transition, again yielding bit errors. The sample placement error need not be large to increase bit error rates in high speed signals as the effective noise margin drops rapidly when the rise time is on the same order as a bit time.

1.1.1.2 Receiver

A typical Receiver, (shown in Figure 1.3), must be able to extract a clock signal from the incoming signal in order to accurately sample the signal's data. To do this, the signal must arrive at the receiver with a sufficiently large signal-to-noise ratio to allow the correct determination between bit values. Limiting amplifiers with a tight bandwidth to reduce noise are often required to buffer the signal before it can be sampled. The samples provide information on both the data, as well as the phase of the input serial stream, by oversampling. A PLL typically adjusts the frequency and phase of a local oscillator to match that of the incoming data stream, so that the data samples are taken in the center of the data bits.

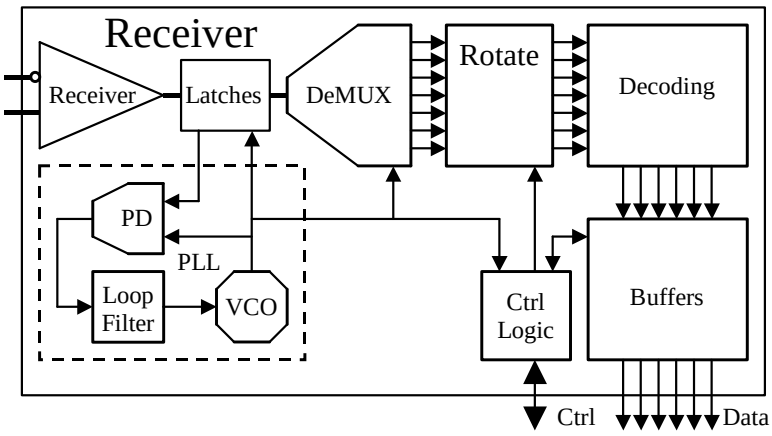

Figure 1.3: SERDES I Receiver Block Diagram

When the transmitter is not sending data, the receiver cannot synchronize itself with the non-existent incoming signal. This condition is called "out of lock". After the transmitter begins sending data, a certain amount of time is required for the receiver to regain synchronization, which is termed the lock acquisition time. In some cases trade-offs must be made between acquisition time and locking range. Jitter can be partially compensated for in the PLL by using a low bandwidth, but that increases the lock acquisition time as well. There are limits as to how much jitter can be tolerated, with the consequences of exceeding those limits being simple bit errors, or even total loss of lock.

Some packetized systems might make use of multiple transmitters that have completely dis-synchronous phases. Once in lock, rotation and decoding circuitry are often used to prepare the data for output. The received data must be presented in a parallel form at the outputs along with a reference signal indicating when the data is valid. Buffering may be used to asynchronously pass data between the receiver and the data sink(s). Various indicators of receiver status are also commonly implemented such as loss of lock(LOL), buffer overruns, etc. When the parallel data is intended for different destinations, it is sometimes advantageous to use another lower data rate SERDES system for each of the parallel output lines.

1.1.1.3 Formatting/Encoding

Typically the transmitter and receiver will use binary voltage levels at their respective inputs and outputs, often sending the signals differentially in the case of high-speed systems. Even in the limited arena of binary voltage signaling, there are multiple ways of representing data bits. As an example, two such methods are return to zero (RZ) and non-return to zero (NRZ) encoding. RZ encoding uses data bits with a zero bit interspersed between each of them. NRZ encoding omits the interspersed zero. The RZ encoding has edges which are closer together for a given bit rate when compared to NRZ. This has obvious negative implications due to the increase in necessary bandwidth, or corresponding reduction in bit rate. The reason RZ encoding is sometimes used is because it allows easier reconstruction of the clock from the transmitted signal as the signal spectrum has a peak at the clock frequency. The NRZ encoding has no such peak, and clock recovery can be more complicated.

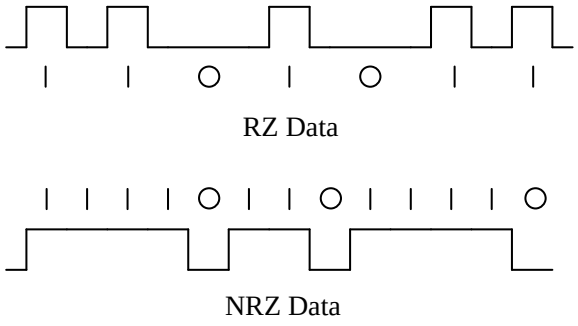

Figure 1.4: Data Formats, NRZ .vs. RZ Data Formats

Data can also be encoded in ways to ensure data integrity, or to alleviate some problematic property of the data channel. Parity bits and cyclic redundancy checks (CRCs) can be generated to reduce the risk of data corruption. As an example of encoding for a specific property, many communication systems require a DC average value of 0V, necessitating an equal number of ones and zeros. Another common requirement is the minimum density of edges in the data stream. This is analogous to requiring a maximum length sequence of ones or zeros. In addition to being necessary to maintain receiver synchronization, frequent edges are necessary in capacitively coupled systems to prevent droop due to RC decay. Methods for

ensuring these properties include Huffman [51] and 8B/10B encoding [52].

1.1.1.4 Framing

The individual bits in a serial stream do not have an inherent property that allows them to be directed to a specific parallel output line. If the receiver can acquire lock at any point in the serial stream, then the parallel output lines can end up "rotated" with respect to the data presented to the transmitter's parallel input lines. Without placing constraints on the data in the serial stream, detection of this rotation is impossible. Depending on the application, there are several ways to correct for this. In a packetized system, a synchronization field can be used with a specific bit pattern that allows the receiver to detect the rotation. A barrel rotation circuit could then automatically correct for this condition before the real data is received. This would require that some framing-related encoding scheme be pre-selected and would reduce the generality of the circuit's application. A synchronization field also has the added utility of allowing the receiver time to acquire lock, before the actual data payload is received.

In a non-packetized system, an encoding scheme for the data can be chosen which would indicate rotation, perhaps by having one of the bits in the "word" dedicated for such a purpose, or encoded in a detectable manner which uses less of the bandwidth. Encoding the data at the transmitter would involve adding auxiliary data and would reduce the available bandwidth.

Different methods might be more suitable for single source parallel data as opposed to multiple. For that reason, the encoding is often left to a higher level circuit in order to preserve generality. A general-purpose receiver may then incorporate a barrel rotation circuit controlled by an external higher level circuit that detects the rotation and can correct for it.

Serial communications research is leading to rapid advances in data transmission rates. The absolute data rates achieved are not necessarily the primary consideration as cost is a major concern for the implementation of these technologies. Circuits using inexpensive high yield manufacturing processes such as SiGe [1] can be much more desirable than those using slightly faster more exotic III-V semi-

conductors. For valid comparison, circuit performance should always qualified by the performance of the devices used in the implementation (f_T,f_{MAX}). SERDES-related research encompasses many different circuit design areas including voltage controlled oscillators (VCOs), phase detectors(PDs), amplifiers, filters, latches, multiplexers, and demultiplexers. They form subsystems of SERDES circuits such as phase locked loops (PLLs), clock multipliers, data retimers, and clock and data recovery units (CDRs). Only when all the individual subsystems have been developed can an integrated monolithic implementation be achieved.

1.2 SERDES Research Project at RPI

The RPI SERDES project has been a joint effort involving several researchers over a large time frame. It was initially started in 1998 by Thomas Krawczyk with the goal of developing a monolithic 20Gbps SiGe serializer/deserializer, (or SERDES), circuit to be used in short haul operations. The SiGe process chosen was IBM's 5HP. This project was funded by the Naval Research Labs, (NRL). One of the project goals was to try to retain flexibility in data rates for interoperability, while simultaneously being capable of achieving high maximum data rates. Professor John F. McDonald was responsible for obtaining the funding for this project and overseeing the work. This researcher joined the project full time in late 1999, and has to date contributed to the design of two complete receiver and transmitter chips, and has also created a third 5HP transmitter chip, of completely original design. Young Yim, (another graduate student involved in this research endeavor), is currently working to move the newer designs to IBM processes such as 8HP to realize even higher data rates.

1.2.1 Research Timeline/Description

This researcher became associated with the SERDES project in a part-time capacity in the fall of 1998 when assistance was provided to Thomas Krawczyk in the areas of layout, simulation, as well as some initial design work. Many current papers were evaluated for ideas and architectures. This work required learning how to use Cadence Integrated Circuit software tools. Cadence layout tools and IBM's

5HP design kit were used to create finished layouts, which were then submitted for manufacture. In due course, this researcher created the layout for a significant percentage of the first RPI SERDES chips, and designed the edge detection logic in the receiver. On page 10, Figure 1.5 shows the history of the entire project.

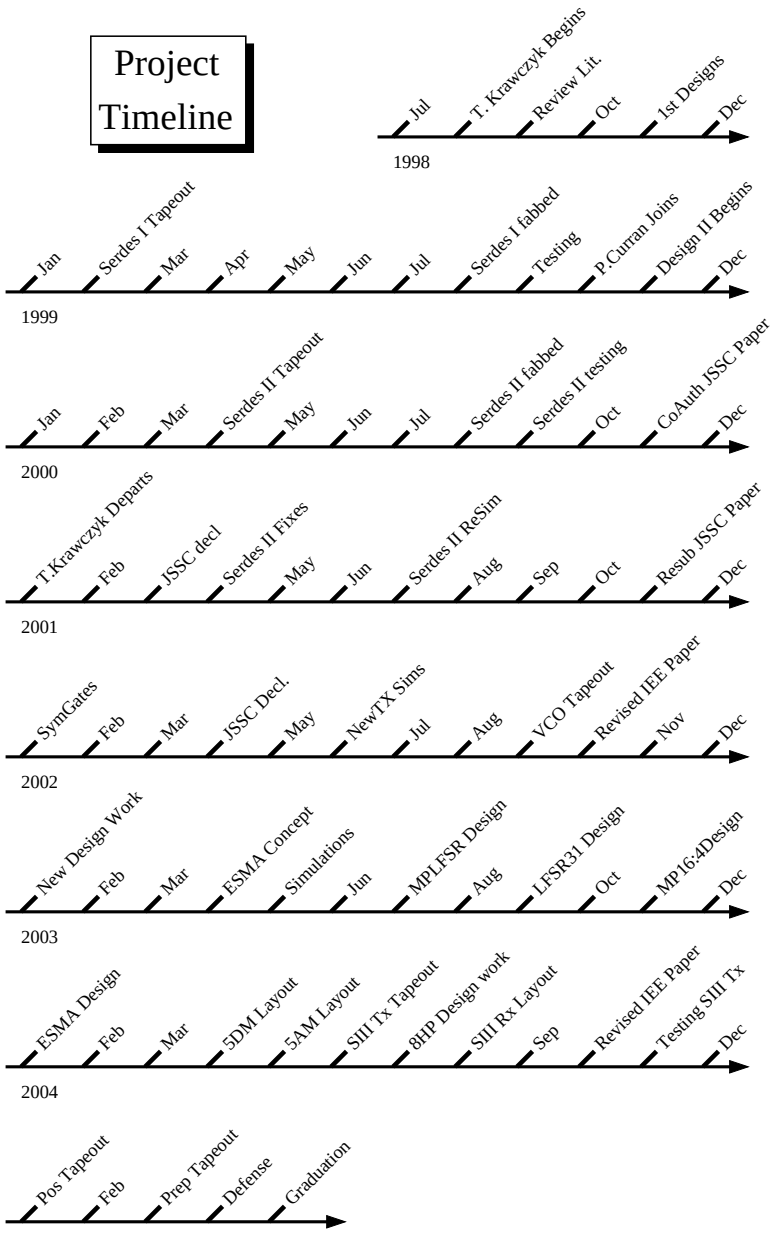

Figure 1.5: Research Timeline of the SERDES Project

In the fall of 1999 this researcher joined the project in a full time capacity, when the initial SERDES designs were being tested. The chips were delivered unpackaged,

in wafers, necessitating the use of our probe-station testing lab. These circuits underperformed to a significant extent, primarily due to relative inexperience in design. Several design errors were identified, and in November of 1999 work began on the SERDES II chip. For this design cycle, there was a need for more flexibility and testability. This researcher contributed a great deal to the design and layout of the on-chip bit error rate, (BER), detecting scheme. The chip design was sent for fabrication in Apr 2000, sponsored by Sierra Monolithics, Inc.

The SERDES II chip was delivered in August 2000. Two quarter-wafers were received for testing purposes. Performance was again tested and deviations from expected behavior analyzed. To facilitate testing, LabVIEW^{TM} software was written to control our test equipment, including a rented spectral analyzer. This researcher co-authored a paper with Thomas Krawczyk that was submitted to the JSSC in Nov 2000. Due to the tight fabrication schedule, several errors were accidentally introduced into the SERDES II design. The paper was declined for publication in March 2001, most likely due to these flaws. These were identified and corrected in layouts, and then simulated. Results were accepted for publication in August 2004 [32].

Work has continued on new architectures with 50GHz f_T devices. Although newer 120GHz f_T 7HP commercial SiGe processes are available they are quite expensive. A question worthy of significant research is the highest rates which may be achieved as a function of the inherent device characteristics. Results in this area would be applicable to faster devices in newer technologies.

In late 2002 this researcher developed several new gates and a new transmitter multiplexer architecture which had the potential to overcome many of the inherent shortcomings in our previous transmitters. In the course of designing a testing chip for the new multiplexer design, several other new circuits were required in order to supply data to the multiplexer correctly. The design and analysis of these new circuits, as well as the discovery and exploration of the factors limiting performance, form the basis for my thesis.

In Sept. 2002, an opportunity for free fabrication on the 5DM process arose with a very small delivery window. In order for work to proceed in the 40Gbps

range, a 10GHz-range VCO would be required for operation and for testing. A 10GHz VCO test chip was submitted at that time, which was delivered in Dec 2002 and tested. Simulation results compared favorably with the measured results.

A full serial transmitter, including the new final multiplexer design and all the required support circuitry was designed for tapeout on 5DM. Unfortunately, 5DM was unavailable, and the designs had to be switched to 5AM. A slightly less capable design than was originally intended was sent out for fabrication in June of 2003. This design was received and tested at data rates of nearly 30Gbps, and clearly demonstrates the feasibility of the new architecture.

In addition to this, a modified version of the chip has been created and it is planned that it will be tested at a later date, perhaps by other members of this research group. Because the limiting factor in the prior design was the VCO, the newer version has a much higher frequency VCO design. This chip will likely not be available for several months. This should allow testing of the new architectures at data rates of above 35Gbps, however the quality of the output is expected to be significantly less at that rate.

Most if not all of the new circuits developed can be implemented in the newer technologies with a corresponding increase in operating rate. As previously mentioned, another graduate student, Mr. Young Yim, has joined our group to work on this and is re-implementing my designs in the experimental 210GHz 8HP process.

1.3 State of the Art

There are many different categories of research papers that can be relevant to SERDES designs. Among these are papers describing multiplexers, demultiplexers, AGC amplifiers, phase locked loops, VCOs, CDRs, as well as complete SERDES implementations. The ratio of data transmission rate to maximum device operating frequency is our primary metric. This research is focused within the regime of monolithic integrated designs. This is dictated primarily by the constraints of the available probing station and testing equipment. External loop filter components or crystals cannot be easily connected to unpackaged chips. In addition we currently lack the multi-gigahertz external clock sources and arbitrary data generators that

testing individual components would require.

Although SiGe is not the fastest technology, it is very competitive from a cost point of view, as well as having benefits due to the ability to integrate several subsystems on one die when using a BiCMOS process. Basic principles surrounding the use of SiGe in HBTs was outlined in [13] in 1982, which considered the introduction of Ge into the emitter. Current SiGe HBTs make use of Ge introduced into the polysilicon base. A comprehensive reference on the SiGe HBT is [41]. Most high-speed serial research in SiGe is focused on 120GHz f_T or higher devices, while this research remains at the 50GHz f_T technology node. Several of the circuits necessary to realize complete monolithic SERDES systems at 40Gbps data rates with 50GHz f_T devices have been fabricated and tested, but many components have not.

The first applicable paper examined when this project began described a complete monolithic 10Gbps Tx/Rx chipset [2] in a 25GHz f_T Si-Bipolar process, published in 1998. It was used as a launching point for much of this work. This paper described a method of using multiphase clocks to extract timing information for a high speed receiver. A "leap frog" ring VCO architecture was described, as well as an "f_T doubler" circuit. The paper dealt with some of the issues regarding the generation of the transmitter output data from multiphase sources. One goal of this research was to investigate flexible data rate mechanisms. A multiple data rate CDR circuit is presented in [3], which detects incoming data rate and adjusts a PLL divider automatically. In [4], a 12.5Gbps Tx/Rx set of PLLs are described using a 45GHz f_T SiGe process. CDRs capable of 19GHz [5] using HEMTs with 50GHz f_T have been reported as early as 1994.

Data rates of 50Gbps for a simple MUX and 46Gbps for a DEMUX were reported [6] in a Silicon Bipolar technology with an f_T of 36GHz, along with a 30GHz static frequency divider, in 1996. These circuits were simple E^2CL elements driven by external clocking. Dynamic frequency dividers with rates as high as 79GHz in an 80GHz f_T technology have been reported [7]. The authors report 53GHz static frequency dividers in the same technology.

In [8], a 40Gbps CDR is described in 50GHz f_T SiGe technology in a non-monolithic approach. A 40Gbps bit stream is split into two 20Gbps streams using

an external VCO and PLL. Only the phase detector is on-chip. The total on-chip component list consists of 3 flip-flops and an XOR. Amplifiers with AGC for 40Gbps receivers have been available [9] since 1998 using 92GHz f_T devices.

Technology is rapidly advancing, and 40Gbps SERDES implementations now exist in SiGe using faster transistors such as 120GHz f_T [10], and a CDR using InP 160GHz f_T devices is described in [11]. The SiGe design uses a CMOS LC quadrature oscillator, while the InP uses a bipolar oscillator using delay lines as the tuning element. A 40Gbps CDR has also been realized in a 72GHz f_T SiGe HBT technology [12]. This uses a single phase 40GHz VCO implemented with delay lines which is then statically divided to obtain quadrature phases.

Basic circuit advances such as using inductors or cross-coupled transformers to improve the performance of individual gates is being investigated with positive results [43], where 40Gbps is achieved in simple multiplexers using 100GHz f_T CMOS devices. Inductors are also useful in improving the frequency of dynamic dividers with 40-GHz operation achieved in [44].

Recently, a report of a 4-to-1 multiplexer with a data rate of 132Gbps was published [45], fabricated in a 210GHz process. This is a fairly high data rate to f_T ratio result, slightly below or comparable to my own desired performance. However, this is a non-monolithic design consisting solely of the 4-to-1 multiplexer, using all external components to provide signals. The clock and data are generated and synchronized off-chip. As stated in the paper, in order to achieve this rate, the input from external clock has to be increased significantly to overdrive the multiplexer.

1.4 Contributions to the Field

This Thesis details the work in creating innovative circuit architectures which facilitate high data rate designs. Initial work involved designing testing LFSRs, (or Pseudorandom number generators), for the SERDES I and II chips which operated at a rate of approximately 2Gbps. In addition, the interleaved sampling system of the SERDES I receiver was developed which operated at 5GHz for 20Gbps operation.

After developing these fairly traditional circuits for prior generations of transmitter/receiver designs, new architectures have been created which allow for higher

clocking rates to be achieved. As architectures, they should be applicable to newer manufacturing processes.

More specifically, in the area of final output and retiming, a new Edge Steering Multiplexer Architecture (ESMA) has been invented. This circuit uses four clock phases and symmetrical design to achieve better non-retimed output characteristics than previous RPI SERDES chips at high data rates.

Providing the input data for this architecture necessitated the design of a new Multiphase LFSR. This circuit combines the functions of a 12-bit random number generator and a 16-to-4 multiplexer with round-robin staggered outputs. For designs beyond the capability of this circuit, a separate count-by-N 31-bit LFSR, and a special Multiphase 16-to-4 Multiplexer have been designed.

A great deal of supporting circuitry such as VCOs, PLLs, and pad driver circuits also were developed in the course of this work, as well as software in the form of simulation HDLs, C code, and LabVIEW^{TM} data acquisition tools.

1.5 Overview of this document

The first chapter in Thesis consists of an overview of the scope of the project, the current state of research in this area, and a list of major contributions. Chapter 2 details the more traditional designs for the earlier SERDES I & II chips. Subsequent chapters elaborate on VCO designs, the Edge Steering Multiplexer Architecture, the 16-to-4 multiphase multiplexer, and the multiphase LFSR. Chapter 8 details the SERDES III Transmitter chip which incorporates many of these elements. Presented in Chapter 9 are some of the limiting factors which impede further increase in performance. There is a final chapter giving conclusions. Lastly, several appendices are included which cover additional relevant aspects of the work, such as the CML logic family, the process technology used for the designs, LFSRs, basic PLL/CMU criteria, and jitter.

CHAPTER 2
Previous Work

This work, funded by the Naval Research Lab (NRL), had the desired goal of achieving a monolithic short-haul serializer/deserializer system with a 20Gbps NRZ data rate using a process with 50GHz f_T HBTs. Another goal was to allow flexibility in the data rates, both for backwards compatibility as well as the possibility of perhaps reaching even a 40Gbps data rate. Current commercial designs using this same technology are at 10Gbps rates, placing our circuits well outside the realm of existing designs in terms of performance. Using innovative circuits, more performance can be squeezed out of a given technology than has been achieved in the past. Previous RPI SERDES II designs produced working chips operating at speeds in the 20Gbps range in an $f_T = 47$GHz process. Commercial SERDES designs operating at nearly half-f_T are either extremely rare or non-existent. In order to reach these data rates under the mostly-monolithic constraint, it was required that new circuit architectures be developed. Two complete generations of full-SERDES designs have been designed fabricated, and tested. The major contributions by this researcher to these two SERDES designs are documented in this chapter. Later chapters describe more recent work on designs for a 40Gbps serializer.

2.1 SERDES I

Researchers at Rensselaer began working on the SERDES project in 1998. Current papers in the field were used as a starting point for the research, with 10Gbps being among the fastest designs reported. The first SERDES designs from our group were submitted to be fabricated in Feb, 1999 and wafers were obtained for testing in August of that same year. The fabrication was funded by DARPA. Two separate main chips were laid out, a transmitter and receiver. In addition, several small VCO test structures were fabricated. These chips were our initial design experience and drew heavily upon [2]. They included a 10-20Gbps transmitter/receiver pair that utilized a multiphase clocking scheme to eliminate the need for

a full frequency clock. The use of a multiphase clock, particularly in transmission, is deprecated due to the difficulty in creating a bit stream without an inherent duty cycle distortion due to asymmetry in the design and layout. However, given our limitations of monolithic design, we must make use of, and optimize, ring VCOs, because 40GHz single phase VCOs are not realizable in a 47GHz process. The receiver needed to use multiple sampling latches for the same reason. Clock and Data Recovery, (CDR), was achieved by oversampling the input data by a factor of two, with an interleave of 8 phases, and using the exclusive-OR between samples to detect bit transitions.

2.1.1 Current Starving VCO

In order to make our SERDES designs operate, a multiphase VCO was required. The VCOs are used as elements in PLLs and provide timing for all the on-chip digital circuitry. The basic ring oscillator method of operation is that when an odd number of digital inverters are placed in a ring, the circuit becomes unstable with a period equal to $2N$ times the propagation delay in one of the inverters, (Figure 2.1).

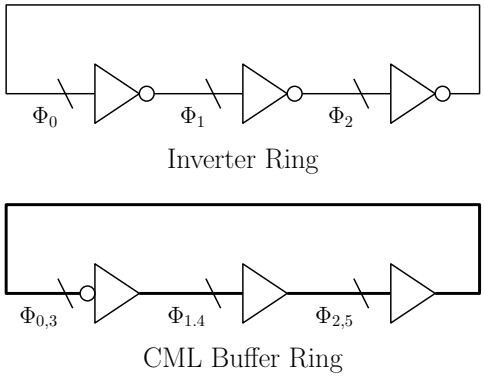

Figure 2.1: CML Ring Oscillator .vs. Inverter Ring Oscillator.

The output of each of the N inverters carries the same signal, but at a phase offset of twice the propagation delay, thus creating an N-phase clock. When using differential CML circuits, inversions are accomplished by wiring, (not shown). This allows us to put an even or odd number of buffers in a ring, with a single wired inversion. The period of the CML ring is $2N$ times the propagation delay of one

of the buffer elements, but because each of the buffer outputs can be used in an inverted or non-inverted sense, you can obtain $2N$ phases separated by a single propagation delay. (An edge travels around the ring, is inverted, then it travels around the ring once again to complete a full cycle.) The frequency of the ring can be modulated if desired by making the delays through the various elements adjustable using some sort of controlling signal. If the duty cycle of the phases isn't critical, a single delay element can be made adjustable, but it is more common to have all the delay elements identical. All our designs have used four phase VCOs, which means we have a total of 8-phases to work with (If you consider the complement of a phase to be a distinct separate phase.) This is a useful interpretation when edge-triggered circuits are in use. We refer to the eight phases as $\{0°, 45°, \ldots, 270°, 315°\}$, or $\{w, x, y, z, \overline{w}, \overline{x}, \overline{y}, \overline{z}\}$, or even $\{\Phi_0, \Phi_1, \ldots, \Phi_6, \Phi_7\}$, depending on the context. Pairs of alternate phases, $\{\Phi_n, \Phi_{n+2}\}$, are separated by 90°, (also known as being in quadrature.)

In our original SERDES, we used a VCO design based on one conceived by Sam Steidl, a fellow researcher at Rensselaer. He came up with a method of speeding up or slowing down a ring VCO by modifying the I_{BIAS} current for the simple buffers, (see Figure A.6), by using a variable voltage reference for the current mirror. As can be seen in Figure A.5, the bias current affects the maximum switching speed. As the modified current mirror reference voltage reduced the bias current to slow the VCO, the VCO architecture was designated a "Current Starving VCO". These VCOs were very fast in general, using as they did only a single buffer delay. Thomas Krawczyk was responsible for the design and layout of the VCO used.

2.1.2 Transmitter

2.1.2.1 Data "Shuffling" Scheme

In this initial design cycle, a transmitter was fabricated which incorporated multiphase VCOs in order to make use of an innovative "shuffling" data scheme that required only a quarter-frequency clock rate for operation. Two 5GHz quadrature clock phases are used to obtain a 20Gbps output. This alternative method allowed us to refrain from dealing with the maximum-rate data signals until the final 4-to-

1 multiplexer/pad driver circuit, requiring only small gains in the final stages of those circuits. Thomas Krawczyk, a fellow researcher, designed this architecture. The information here is presented to serve as a contrast with the newer multiplexer system of more recent design which serves the same function.

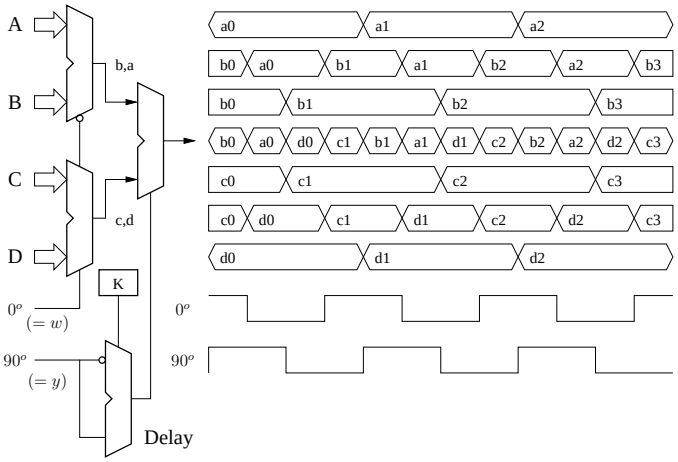

Figure 2.2: SERDES quarter-clock shuffle multiplexing. This allows correct interleaving of four data streams at the output using a quarter rate in-phase and quadrature clock.

This design consists primarily of three 2-to-1 CML multiplexers in the hierarchical configuration shown at the upper left of Figure 2.2. In the illustration, "w" and "y" represent in-phase and quadrature clocks. The input data lines on the left of the figure, $\{A, B, C, D\}$, would remain stable while they were "selected" by the "w" phase, That is, the corresponding shift registers would not be clocked while the data input was active. A state machine driven by one of the phases controls when the four bit registers would be loaded or shifted, and they shift in pairs.

All edges present in the serial data stream are generated by the "w" and "y" phases. If the multiplexers were "ideal", the phase difference of the "w" and "y" clock would determine the positions of adjacent edges in the output stream. However, these *real* multiplexers have a propagation delay. This requires that the "y" phase be delayed by exactly the same amount as the "w" phase in passing through one multiplexer, in order for the timing of the edges generated by the "y" phase to appear centered between those generated by the "w" phase. To create a delay of the same duration, an identical circuit was used to ensure uniformity.

The "y" phase is sent through a an extra multiplexer, (seen at the bottom left of Figure 2.2) This type of serial data generation is referred to as "unretimed", as there is no full data rate clock latching the data after the multiplexer, or even a single clock phase responsible for the final timing.

Figure 2.3 consists of a detailed breakdown of the SERDES transmitter. This design consisted of several functional units, each of which will be described. The details of the various components are given elsewhere, and only a functional overview will be presented. The pad drivers and receivers have been omitted from the diagram for clarity.

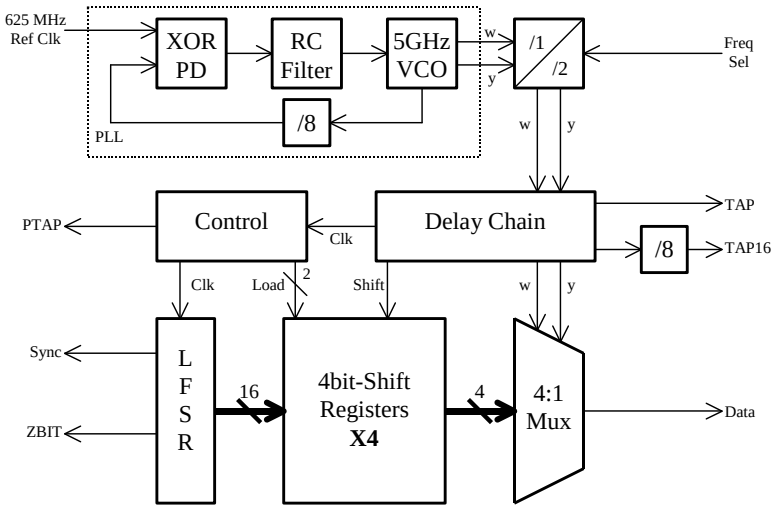

Figure 2.3: Block diagram of the SERDES I Transmitter.

The transmitter is designed to take 16 bits of data from the LFSR(linear feedback shift register), and serialize the data at the data output pad. The LFSR is the only source of data for this chip. Probe constraints limited the number of external inputs which could be supplied. The "$ZBIT$" line signaled when the sequence repeated for triggering purposes. A key design feature of the transmitter was that it would make use of a quarter data rate clock. When operating at a 20Gbps transmission rate, the local voltage controlled oscillator is only running at 5GHz. The main clock for the transmitter is generated by a 4-stage ring buffer VCO, embedded in a divide-by-eight PLL, that is driven by an external 625MHz reference source. For the divide-by-8 blocks, a series of three divide-by-2 circuits were used which each consisted of a master-slave(MS) D-latch with inverted feedback. The

VCO was based on a design by Samuel Steidl, as previously mentioned, and used the "current starving" technique as a control method. The ring buffer nature of the VCO made multiple phases available. For this design two quadrature phases are required, and in the figure they are labeled "w" and "y". These phases are used to multiplex the output data via the 4-to-1 multiplexer using the "shuffling" scheme previously described.

The PLL in the SERDES transmitter was our first attempt, and it was of the most primitive sort. The phase detector was of the XOR type, implemented as a single CML gate. The RC loop filter consisted of two low-pass RC ladder stages on one side of the output of the differential CML XOR. Static phase error was not a concern in this design as the generated clock was unrelated to any other external input signals, instead, all other signals were derived from it. The external clock was expected to be a much more accurate reference, so the loop bandwidth was intended to be fairly large, and was set to 500MHz.

The in-phase(w) and quadrature(y) signals pass through an optional divide-by-two circuit which was controlled by an external pin. This would allow either 10Gbps or 20Gbps operation. The division circuit generated a new "w" and "y" from the original "w" exclusively rather than dividing each of the original signals by two. It was thought that this would produce a better aligned pair. This was accomplished by using a pair of divide-by-2 blocks as described above driven in parallel by the "w" phase, but the input to one of the blocks was inverted.

The clock signals are distributed to various parts of the chip via the "Delay Chain" unit. This consisted of a series of buffers with "taps" at different points that would then pass the suitably delayed signal on to other units. In some cases "taps" were added merely to balance loading between the two phases, and for testing purposes two of these were made available at output pads. The first, labeled "TAP", had the full 5GHz clock signal while the other "TAP16" had a divide-by-8 unit added to make frequency measurements easier.

One output of the delay chain was used to drive the bank of four 4-bit shift registers. That signal was inverted to half of the registers so that they would shift in an alternate pattern as required by the "shuffling" scheme. One bank of registers

could then be reloaded while the other bank was in the process of shifting.

2.1.2.2 LFSR

An internal source of data was required to supply the transmitter, and a linear feedback shift register(LFSR) was designed for that purpose. The SERDES I LFSR was extremely simple, (Figure 2.4.) As it was intended to provide a detectable sequence at the receiver, the sequence was chosen to be short. A Fibonacci-style 4-bit LFSR was created using a maximal sequence of {4,3}, (see Appendix B on LFSRs), with the addition of reset circuitry to detect the all-zeros condition to feed back a forced one in that case. To create a full 16-bit parallel output for the transmitter, the output of the last LFSR bit is shifted into another 12-bit shift register chain. This resulted in a repeating 16-bit pattern which did not necessarily provide ideal characteristics.

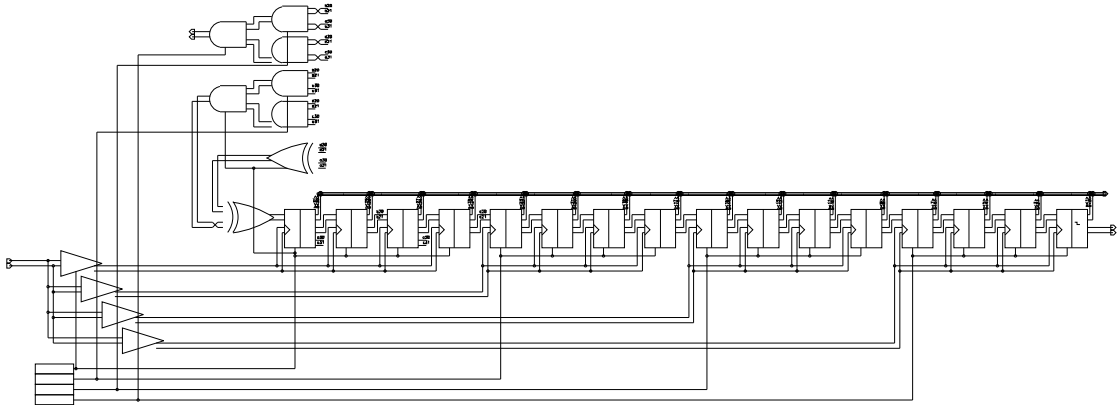

Figure 2.4: SERDES I LFSR schematic diagram.

The circuit, (shown in Figure 2.4), introduced a layer of buffering for the clock, driving each group of four latches with its own buffer. A four input AND gate was used to detect the all-ones condition to drive a signal to the output pin. The propagation delay between the clock and the output change was on the order of 45ps.

Figure 2.5 shows the layout of this block of circuitry. This was one of the first layouts this researcher created for the SERDES project. The left side of the layout contains all the clock buffers, as well as the all-ones-state detection circuitry. The

Figure 2.5: SERDES I LFSR circuit layout.

layout of the four groups of four registers on the right follows a serpentine layout, so that the beginning and end of the chain are close together. This minimizes parasitic loading and increases the maximum clocking rate. However, this kind of attention to detail was not really necessary as this block of circuitry operated at one sixteenth the rate of the highest speed circuitry.

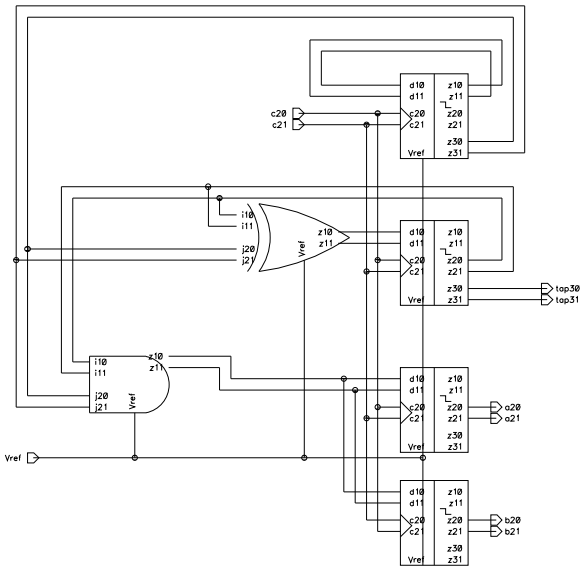

Figure 2.6: SERDES I Transmitter control circuit.

In order to generate the clock for the LFSR and the "load" signals needed by the shift registers, a state machine (CONTROL) unit was necessary. This circuit

generated a 25% duty cycle pulse at one fourth the frequency of the main phases, which was then used to enable a load for one half the shift registers. The other load line was derived from the first, but changed after a delay of one half of a main phase period. The LFSR also advanced to the next state using this signal. An additional signal equivalent to a divide-by-4 of the main clock was incidentally generated, and sent to the "PTAP" output pad to monitor the operation of the control unit.

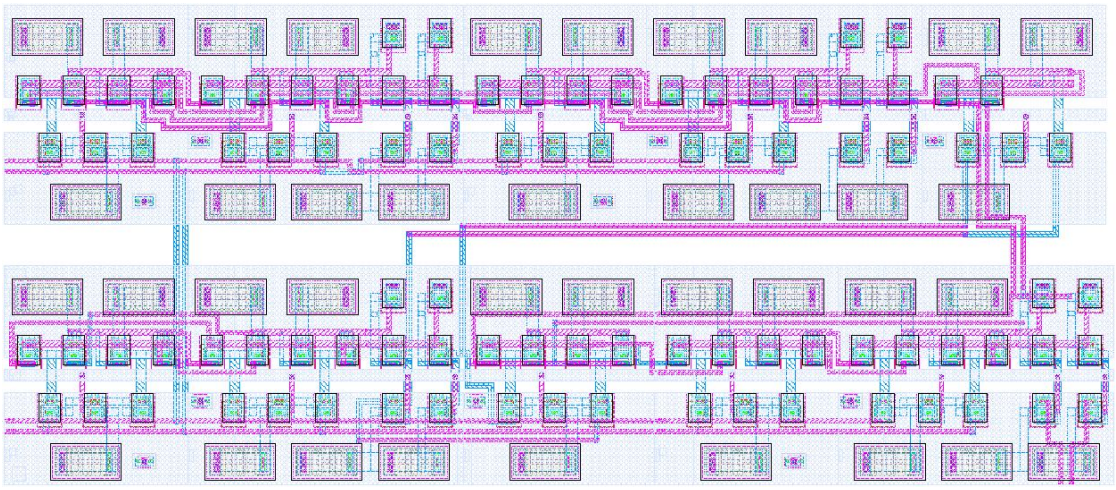

Figure 2.7: SERDES I control circuit layout.

The layout of the control block is shown in Figure 2.7. As this control block needed to work at a high data rate, particular attention was paid to the layout. Minimum spacing between transistors and other components was used to reduce parasitics. In retrospect, the design could have been greatly improved by incorporating many of the individual logic gate functions within the current trees which formed the latches.

In this transmitter design, the shift registers were designed and laid out by Thomas Krawczyk. The output of the shift registers was sent to the final 4-to-1 multiplexer to generate the final serial stream. In simulations without parasitic extractions, the transmitter performed up to 23Gbps.

2.1.3 Receiver

The design of the SERDES receiver is diagrammed in Figure 2.8. The receiver integrated circuit consists of two major sections, concerned with test signal generation and the receiver itself. The upper part of the figure is for test signal generation and features a 5GHz VCO of the simple current-starving architecture. This VCO is controlled by a voltage supplied via an external pad. Three of the VCO phases pass to a frequency modifier block used to create the onboard generated test signals. One phase of the VCO is available at an output pad for monitoring. In the frequency modifier block, one phase is passed through a divide-by-two circuit identical to those used by the transmitter. Two of the phases that are in quadrature are combined using a symmetric XOR to produce a signal at twice the VCO frequency as described in [22]. One of the three phases is passed through unmodified. These signals, at the VCO frequency, half the VCO frequency, and twice the VCO frequency, are fed to the input select unit. The input select unit also receives one signal from an external pad. The selector is just a 4-to-1 multiplexer controlled by two external select signal pads, ("SelA", "SelB").

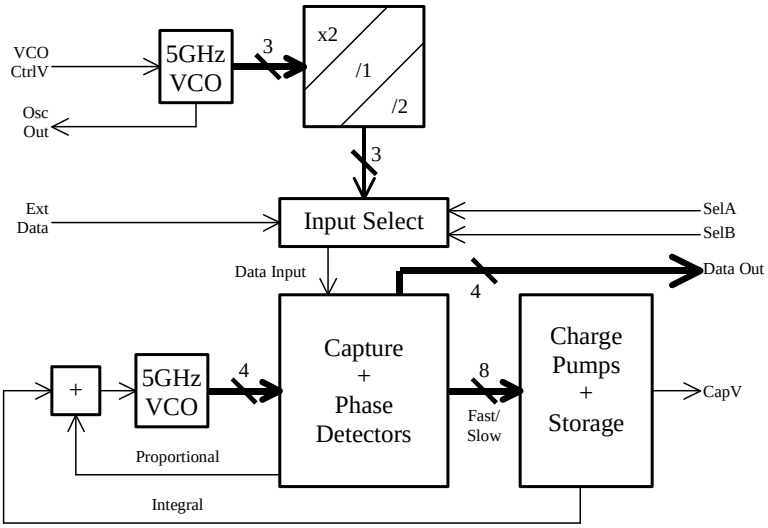

Figure 2.8: Block diagram of the SERDES I receiver.

The output from the input select unit demarcates the boundary between the testing circuitry and the receiver proper. The receiver has a separate VCO that is used to drive the data and phase detector latches. Each of the four VCO phases

and their compliments are used, resulting in eight samples taken per VCO period.

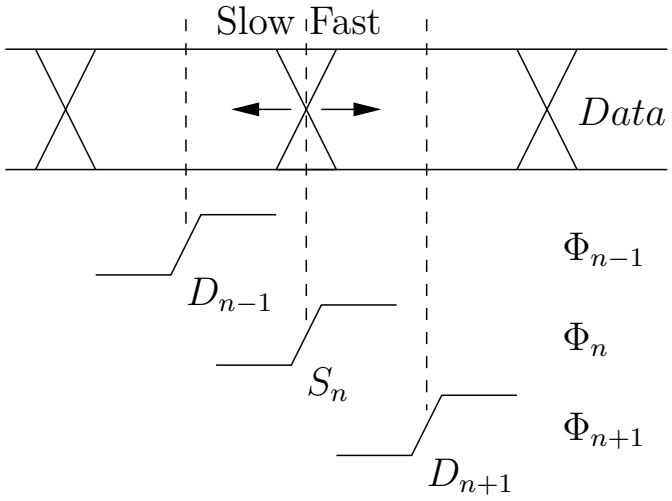

Figure 2.9: Receiver CDR Operation. Each of the eight phases ($\Phi_0 \ldots \Phi_7$) sample the data stream creating four S_i samples and four D_i samples.

Four of these samples(D_i) are intended to occur in the center of data bits while the S_i samples are used to detect the position of the potential edges occurring between the bits. The values captured by the data latches are used directly as the output data from the receiver. The exclusive-OR between each data bit and an adjacent edge detection sample would indicate the presence of an edge. The detection of such edges is necessary to perform the clock-data-recovery(CDR) operation. The local oscillator must adjust to variations in the phase of the input stream.

Edges detected just after a data bit rather than just before would be an indication that the edge was present earlier than in the target crossover point, and a "slow" signal should be generated as a command to the VCO. Likewise, if the edge arrives just before a data bit, a "faster" signal needs to be generated. These "faster" and "slower" signals need an off state so that when no edge is detected, the VCO remains at the same frequency and phase as identical bits are sampled in sequence. Usually, these signals are summed and integrated using a charge pump which supplies the control voltage to the VCO.

This type of CDR is called a "bang-bang" PLL approach, because small discrete adjustments are constantly being made. (The phase detector transfer function

is a step.) In some designs, even those with multiple data capture latches, as few as a single edge capture latch is used. This has the benefit of having a smaller loading on the input amplifier(s), but it places restrictions on edge placement in the serial stream. The edge frequency restriction must be enforced between every pair of bits, rather than as a total edge frequency per bit.

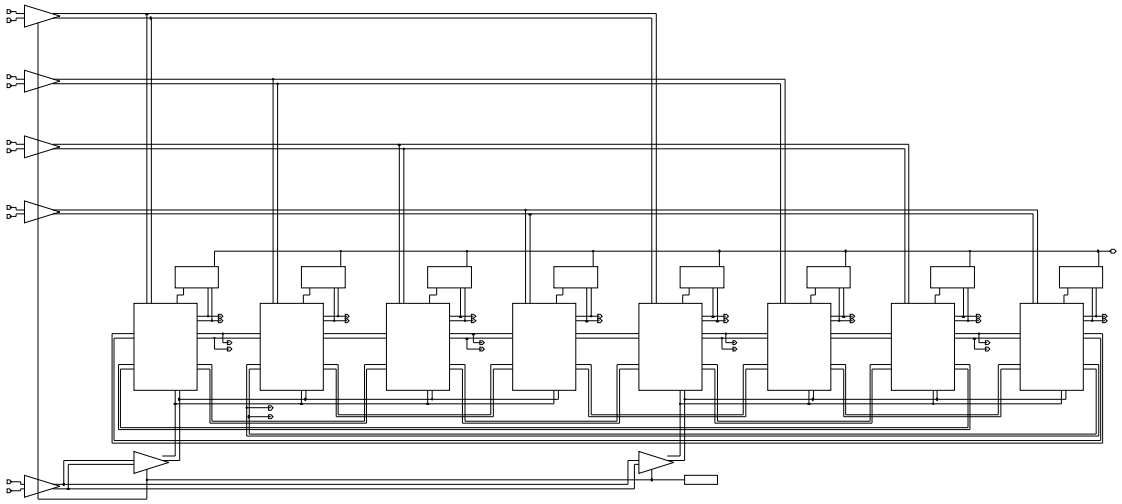

Figure 2.10: Receiver Sample and Detection Schematic. Clock phases enter from above left, the data from the lower left. The larger blocks represent capture latch circuits, while the smaller associated blocks represent circuitry to generate the summed proportional current signal.

For this design a capture architecture was developed that exploited the symmetry of our clocks and full edge detection. The schematic for the architecture is shown in Figure 2.10. The eight sections are each driven by a specific phase, and they latch either a data bit or an edge detection bit. Each block performs an exclusive-or (XOR) with its neighbors, and generates a fast/slow signal. Each capture block is identical, with all differentiating functionality determined by wiring between blocks. The circuitry of each main block is shown in Figure 2.11.

The reason for the circuit's complexity is that it ensures that the XOR data is valid. Consider that if the inputs from the capture latches were directly XORed, then there would be a period between the first and second latching where the data in one of the latches was from a previous 4-bit group. In this design, the phase detector XOR outputs are gated so that they were only active immediately after

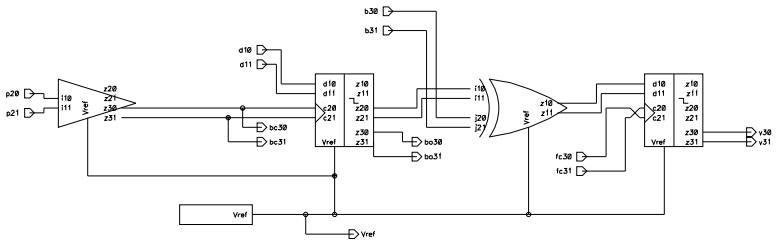

Figure 2.11: Receiver sampling block schematic diagram.

their second input became valid. Data from each block is fed to both previous and later blocks for use in comparison. In addition, the clock is also fed back to latch the good value of the previous pair of latches.

The symmetry of the circuit in Figure 2.10 enabled the creation of the layout in Figure 2.12. The data paths between adjacent phases were sent in channels in the layout for this purpose. The samplers are arranged in an interleaved pattern rather than linearly so that the delays on average would be uniform, and shorter.

The eight potential fast/slow signals are combined to produce a proportional signal that is sent to the VCO. Figure 2.13 shows the circuitry of each of the smaller blocks in Figure 2.10. The open-collector outputs allow them to be connected together for summing. Each can sink 0.7mA. Note that multiple edges can be detected at the same time, and the average is computed, giving many different output values in the range $0 \ldots 2.8$mA. In addition to the proportional signal, the fast/slow outputs are used to control the charge pumps, implemented with FETs, connected to a capacitor to integrate the error. The pFETs and nFETs are $W = 5.59\mu m$, and $W = 3\mu m$ respectively, with $L = 1\mu m$. These drove the 4pF $75\mu m$ x $75\mu m$ MIM capacitor. The layout of this circuit is shown in Figure 2.14, which basically "capped" the layout of Figure 2.12.

The proportional and integral signal were summed using a linearized buffer operated in differential mode to create the reference signal for the VCO, thus closing the loop. Constants needed to be determined for weighting the proportional and integral gains for this PLL, and at the time we were quite inexperienced in PLL design. To quickly find usable values, this researcher wrote a C-program to act as a discrete event simulator to measure things like lock in time, and the same-value run lengths which would result in a bit-slip.

Figure 2.12: Receiver sample and detection circuit layout.

2.1.4 Layout

The transmitter and receiver were manufactured as shown in Figure 2.15. Each of them was approximately $2mm^2$. Both chips were meant to be tested via a pair of 10 pin, (6 data channel), probes. In addition, a slightly modified layout was created which was intended to allow data generated by the transmitter to pass directly into the receiver while still allowing testing using the same pair of 10 pin probes.

2.1.5 Design for Testing

Our limited external signal probe capability prompted us to incorporate an LFSR to generate pseudo-random data for the test transmitter to send. Due to our limited testing facilities, we used on-chip signal generation to provide the chip with data, and the data was only demultiplexed to 4 lines. The final 4-to-16 stage

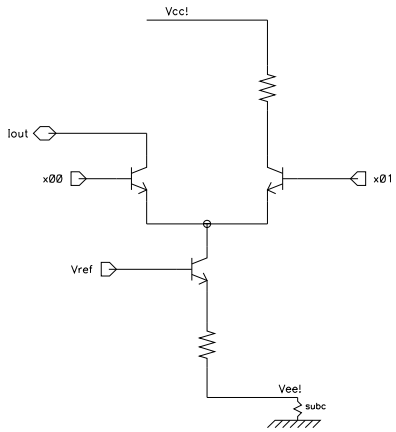

Figure 2.13: Proportional current generator. This circuit uses a single-ended output to draw current through a common collector node.

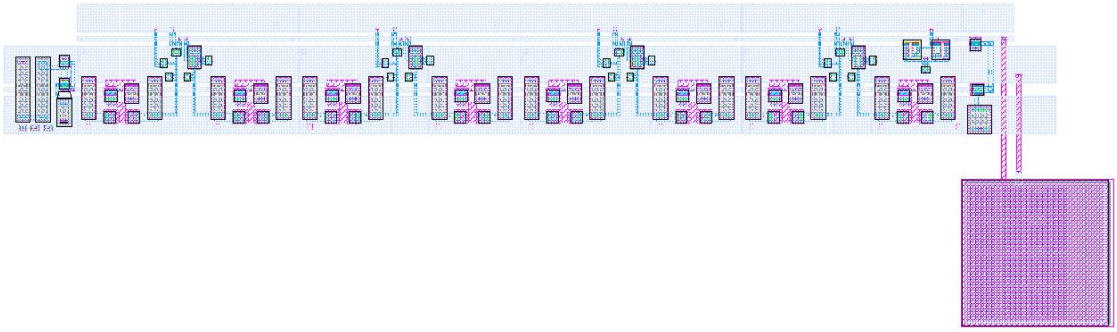

Figure 2.14: SERDES I integrator circuit layout.

would be implemented in a later design. Both chips deviated significantly from expectations based on simulation.

Our results from this initial design cycle showed how far our simulations deviated from the physical circuits, especially in the area of the performance of the VCOs. Analysis led to feedback with IBM's modeling staff to improve accuracies. In testing, the receiver operated at nearly the desired rate, while the transmitter under-performed by around 25%. A paper was prepared and submitted to the International Solid State Circuits Conference (ISSCC), but unfortunately was not accepted for publication, most likely due to the mismatch in operating ranges of the transmitter and receiver.

The transmitter VCO center frequency was intended to be 5GHz, with a control voltage between -.8V to -1.6V, so at -1.2V we should have approximately 5GHz.

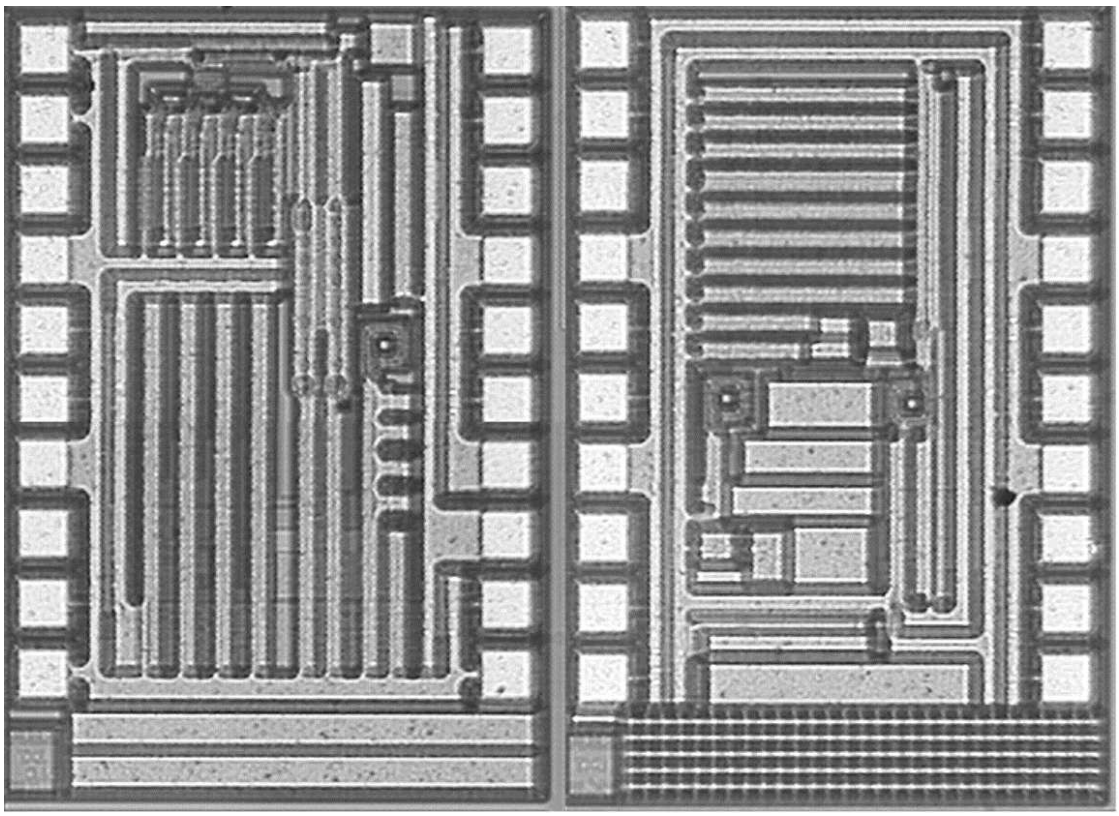

Figure 2.15: Microphotograph of the SERDES I chips. The transmitter is on the left and the receiver is on the right.

The gain even in simulation was non-linear, and actually went negative in testing at high frequencies due to drive beyond peak f_T current. We decided to come up with a more linear and controllable design for the next prototype.

The transmit multiplexer scheme was found to be less than ideal as well, and showed a noticeable duty cycle. This was partly due to asymmetries in layout, and also due to our multiplexer scheme itself. We had intended to delay the second phase by the same amount as the first, but there were additional unplanned for propagation times we had not considered. One of the main problems was that the CML 2-to-1 multiplexers had a different propagation time to output for the data inputs as compared to the select line. We were unable to achieve a perfect balance using these multiplexers. There was always an constant offset between edges produced by the particular phases. This was the impetus for much of the redesign in the next prototype.

2.2 SERDES II

Using the information learned in the first design cycle, our group began to develop a second version of our serializer/deserializer in late 1999, early 2000. The complete new SERDES II design was created and submitted for fabrication in March of 2000. The window of opportunity for fabrication was made suddenly available by a corporate benefactor, Sierra Monolithics Inc. Due to the hurried nature of our preparation for tape-out, and because our allocated chip size changed more than once, we were unable to do the kinds of extensive verification of our final layouts that we normally would have done. Because of this, several small errors were able to make it into the final layout, resulting in small degradations in performance of the physical chip.

The new design used a 3.4V, (down from 4.5V), nominal power supply, and was created using a 5 layer metal process. Most of the advantages of working with 5 layers were lost due to the funding organization's desire to have this design be packageable. Towards that end the chip was designed with both regular bondpads which we could probe using our test station, and C4 solder bump pads for flip-chip packaging. A single chip was created for both transmit and receive operations. The chip contained a full transmitter and receiver in a single layout, which were almost completely independent in operation.

In July 2000, the chips were received and testing began. During the entire period between design revisions, IBM was engaged in refining their process, and the models along with it. This has lead to more small deviations from expected behavior, but these deviations were smaller with this revision than they had been in the previous one.

2.2.1 Symmetric Multiplexer

To address the problems with the first prototype transmitter duty cycle, a new 2-to-1 multiplexer was designed with characteristics designed to provide more uniformity in propagation delay for the data and select inputs. This new multiplexer circuit which was intended to alleviate the imposed duty cycle on the output by the previous design. Prior to this, our multiplexer had used a multi-level approach using

a regular CML tree. This led to asymmetry in output timing due to the different propagation times of the different signal levels through the CML circuit. The circuit below is much more symmetrical.

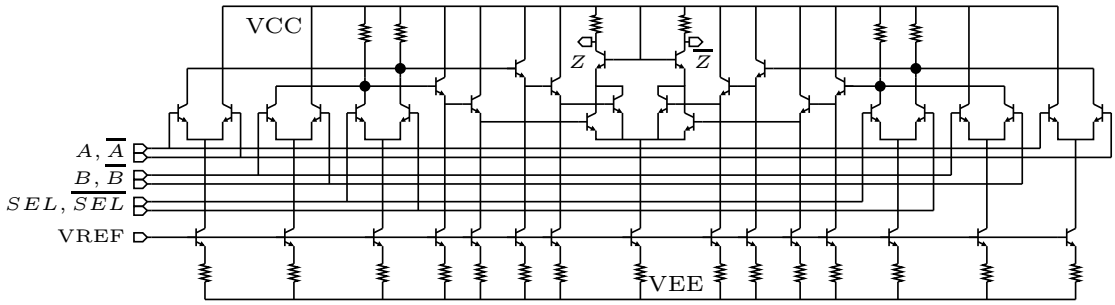

Figure 2.16: Symmux circuit designed to have equal propagation time from each of the inputs to the outputs.

This new multiplexer circuit features a nearly uniform propagation delay between any of its three inputs. It is used in both the final multiplexer, as well as in the dummy delay multiplexer intended to delay clock phases by the exact same amount of time, to ensure symmetric output. A provisional patent was obtained for this circuit, (Figure 2.16.) It should be noted that it was not 100% data-select symmetrical due to loading and state effects in the "interior" of the circuit.

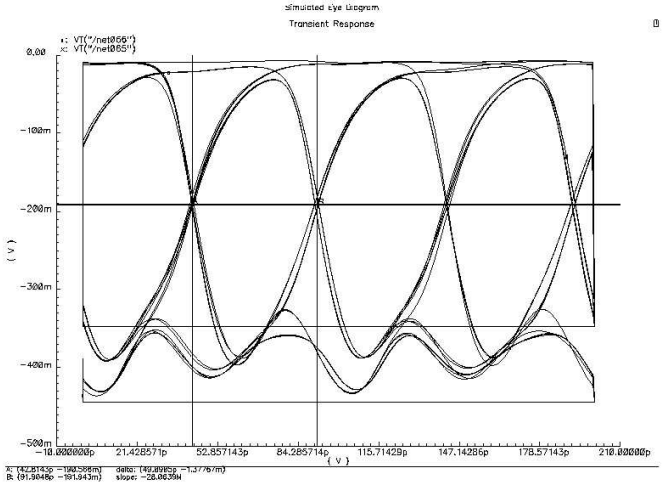

Figure 2.17: Duty cycle plot of the symmux based 4-to-1 multiplexer under ideal conditions at 20Gbps.

Using this circuit and extremely careful and symmetrical layouts, we are able to get reasonable performance. The simulated plot above, Figure 2.17, shows how

much duty cycle the multiplexer scheme imposes when fed with correctly aligned signals. Measurements of the above plot show a very small duty cycle.

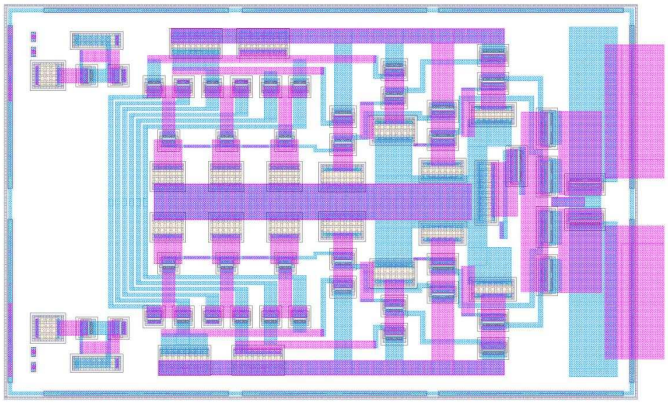

Figure 2.18: Layout of the symmux circuit. The symmetry of the layout ensures that parasitic effects are matched as closely as possible.

Figure 2.18 is a layout of the final 2-to-1 multiplexer in the transmitter, which directly drives the output pads. Similar attention to detail was used throughout the critical areas.

2.2.2 Leap & FFI VCO

For performance reasons we decided to come up with a new VCO architecture for the second prototype. The current starving VCO was too non-linear to safely use in our PLLs. More complex designs seemed to always degrade final performance to unacceptable levels. Eventually, a match of circuit speed-up techniques and control-enhancing slow-downs was achieved which yielded acceptable performance.

In our search for better ring VCO designs, one of the discovered performance increasing techniques was to use the interpolation of the phases at buffer stages to speed up the effective propagation delay at each stage, (See Figure C.6.) Each buffer element would have two inputs, and the output signal consisted of an interpolation of the two. While one of the inputs was tied to the output of the previous element, (as in a conventional ring), the second input is connected to the output of the element preceding the preceding element. In this way, the output of an element "leap-frogs" ahead of the element immediately ahead of it. We chose a four buffer VCO

architecture as being the simplest leap-based design that could provide us with the quadrature phases. Figure 2.19 below shows the basic leap architecture operation. This design utilizes the feed-forward and averaging scheme to allow each buffer to "anticipate" the incoming edge, resulting in a 33% increase in operating frequency. The phases P_i would normally be separated by a full stage delay. However, because P_n is based on the average of the two prior edges, the output can be anticipated. As described above, the frequency of the VCO was controlled by varying the amount of current through the CML buffers, altering the current-dependent value of f_T.

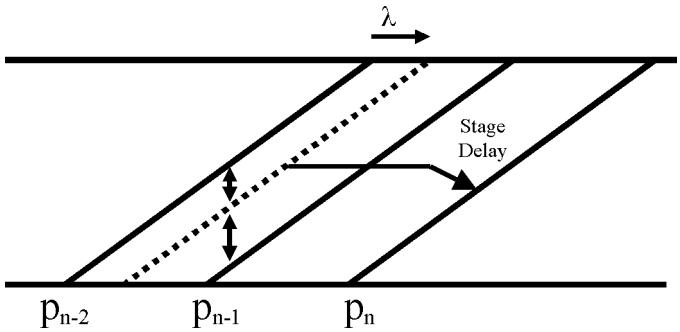

Figure 2.19: "Leapfrog" architecture interpolation plot. By averaging the signals from previous buffers, a 33% increase in frequency can be obtained.

This technique gave a great performance boost but didn't address the controllability issue. What was needed was a way to speed up or delay each phase without sacrificing higher frequency operation. Normal techniques such as using varactors were considered, but abandoned as they either degraded performance to too great a degree or had too small a range.

Eventually, the idea of the Feed Forward Interpolated (FFI) VCO was developed, by Thomas Krawczyk, in which the amount of interpolation between the stages can be varied, (λ in Figure 2.19), effectively "sliding" the position of the interpolated edge. The circuit diagram for one of the buffer stages of the new VCO is shown below in Figure 2.20. It was used in both the transmitter and receiver circuits.

This VCO phase interpolation scheme would exhibit a great deal of jitter as a result of attempting to interpolate between the clock phases if the edges were steep. Because of this, the buffers were intentionally linearized and the gain reduced to an

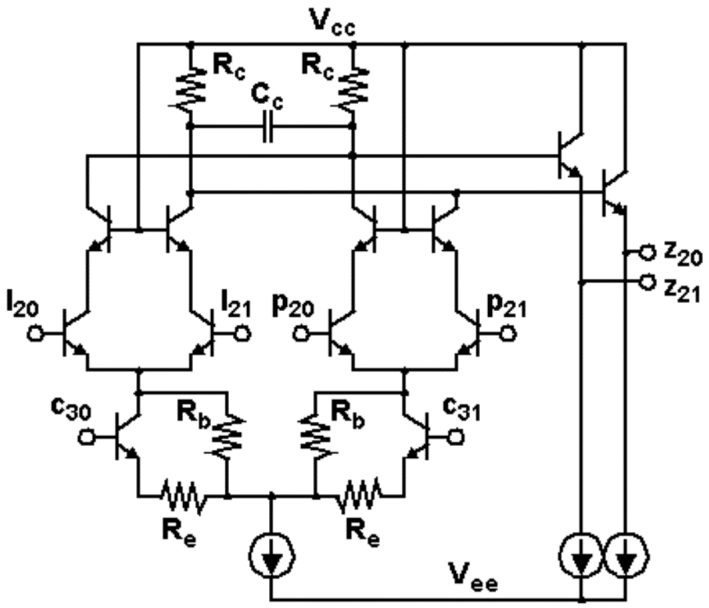

Figure 2.20: Buffer element of the Feed Forward Interpolated(FFI) VCO. By varying the interpolation ratios, the frequency can be controlled.

extent to allow sinusoidal waveforms, and resulting in fairly smooth interpolation. The purpose of the R_e resistors in the CML tree paths in the circuit above are to provide this linearization. In addition, resistors R_b prevent a 100% current swing to one side or the other. The control voltage has to be supplied differentially, and the circuit exhibits a high degree of common mode noise rejection. The optional capacitor at the top of the circuit allows control of the center frequency. In the circuit's highest intended frequency incarnation, the capacitor is left off completely.

Figure 2.21 above shows the large usable linear region of the VCO, as well as it is large range. The phase noise was very good compared to other ring VCO circuits, being approximately -90dBc/Hz at 1MHz. However, this phase noise was somewhat frequency dependent, with the noise increasing as the frequency increased. This was due to the stage decoupling issue, whereby the 4-stage VCO starts to act as two independent 2-stage VCOs when the interpolation is biased too far towards the input from the buffer on the opposite side of the ring.

The phase noise of the FFI VCO is shown in Figure 2.22 for closed and open loop operation. Our PLL was tuned to minimize output noise based on simulations

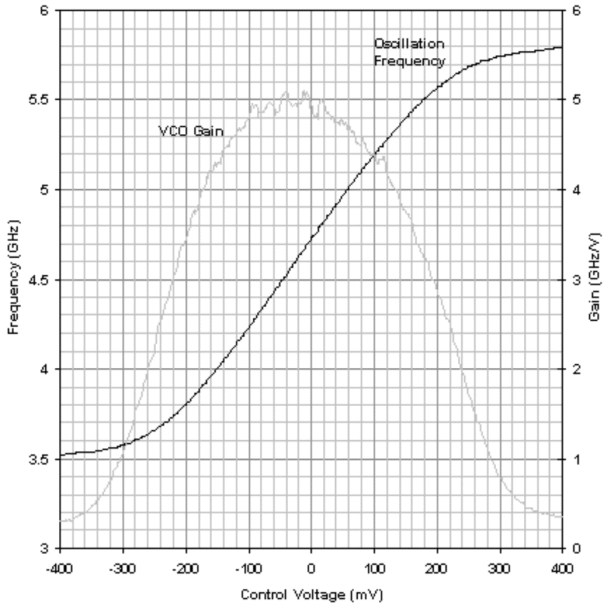

Figure 2.21: FFI VCO control plot. The plot above shows a wide linear range.

of the VCO, and the noise characteristics of our 625MHz reference source [30]. Because of the characteristics of the source were less than ideal, the loop bandwidth was reduced to compensate. Below is a VCO layout showing the degree of symmetry achieved, (Figure 2.23).

With a better reference signal source available, we can increase the loop bandwidth, making the transmitter track the better external source further out. This is planned for future designs as we have purchased a high quality signal source.

2.2.3 SERDES II Testing Scheme

The new testing circuitry for the SERDES II design was planned to allow for self-testing. The LFSR in the transmitter had a complementary LFSR in the receiver, and a means of detecting actual bit errors was devised. This is quite challenging as when a serial transmitter and receiver such as ours operates, as there is no inherent synchronization between the parallel input bits in the transmitter and the parallel output bits in the receiver. The data can be "rotated" at the output, making any simple pattern difficult to test. The scheme devised makes use of the

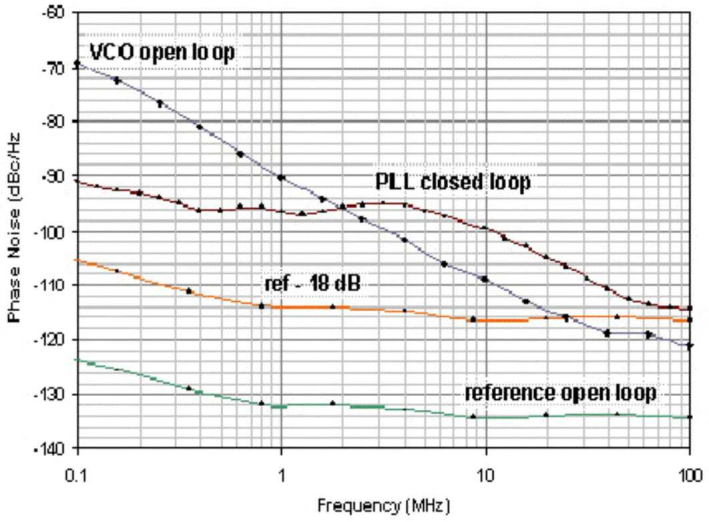

Figure 2.22: Closed loop performance of the FFI VCO in our PLL.

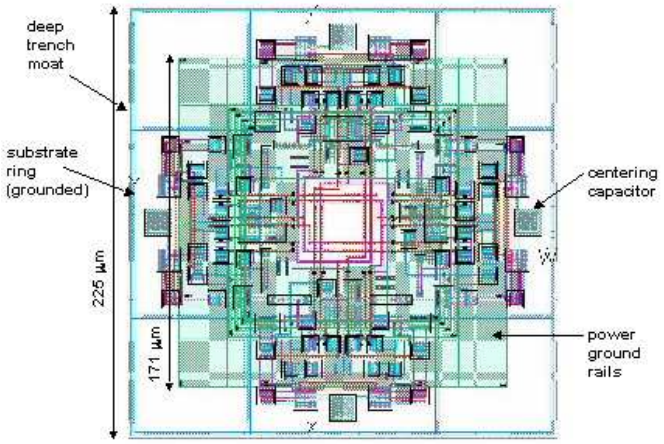

Figure 2.23: Layout of the FFI VCO showing symmetrical loading and isolation.

properties of LFSRs, (see the Appendix B on LFSRs), where each shift register latch produces the same serial stream, only offset in time. If any **one** of the receiver's output parallel bits is examined, it should match the sequence of the LFSR without error. The received bit would be compared against that of a matching LFSR in the receiver, and if different, the LFSR in the receiver would be reset. The absence of reset signals for times longer than the sequence would indicate correct transmission.

The testing scheme can only detect the absence of errors by receiving a com-

plete sequence. If the sequence is too long, it is possible that no full sequences would be detected. It was felt that a full 16-bit LFSR sequence would be overly long,(65535 bits), to safely test this design. Instead, a 12-bit LFSR design was chosen to generate 16 bits.

As described in the Appendix, for the Galois form of the LFSR, if two outputs are from adjacent latches without an intervening exclusive-OR, their output patterns are shifted by only one time-step. If these bits were directly used to form part of a serial stream, and the receiver "slipped a bit", the error would not be detected.

The solution was to carefully choose an ordering of the 12 bits, with duplication, to create the 16 bit parallel input to the transmitter. The offset of each latch in the sequence was calculated, and the bits were selected to provide greater than two bit distances in the sequence position for each pair of adjacent bits in the parallel input.

2.2.4 Transmitter

The major performance changes to the transmitter architecture over the original SERDES design are the introduction of the symmetric multiplexer, the use of newer FFI VCO design, and the testing scheme. The new transmitter design, (Figure 2.24), was given multiple VCOs and a variable divider inside its PLL to allow it to transmit at several different frequencies. By varying VCOs and dividers, 10Gbps, 20Gbps, and even 40Gbps could be selected as desired output rates, given a nominal 625MHz input signal. The 625MHz signal could be varied through a large range with the transmitter still remaining in lock, allowing almost any intermediate data rate. Circuitry for testing now included a 12-bit pseudo-random generator in the transmitter to generate 16 streams of pseudo-random data.

The transmitter PLL also featured a new active op-amp based loop filter, and moved from an XOR phase detector to a 3-state phase detector to more closely track the "RefClk" input. Much more attention was given to optimizing this filter to ensure the best possible performance with regards to lock in range, lock in rate, and minimizing output phase noise. In the original prototype, only four bits were used for testing. In contrast, this design features a full 16-bit architecture. A mul-

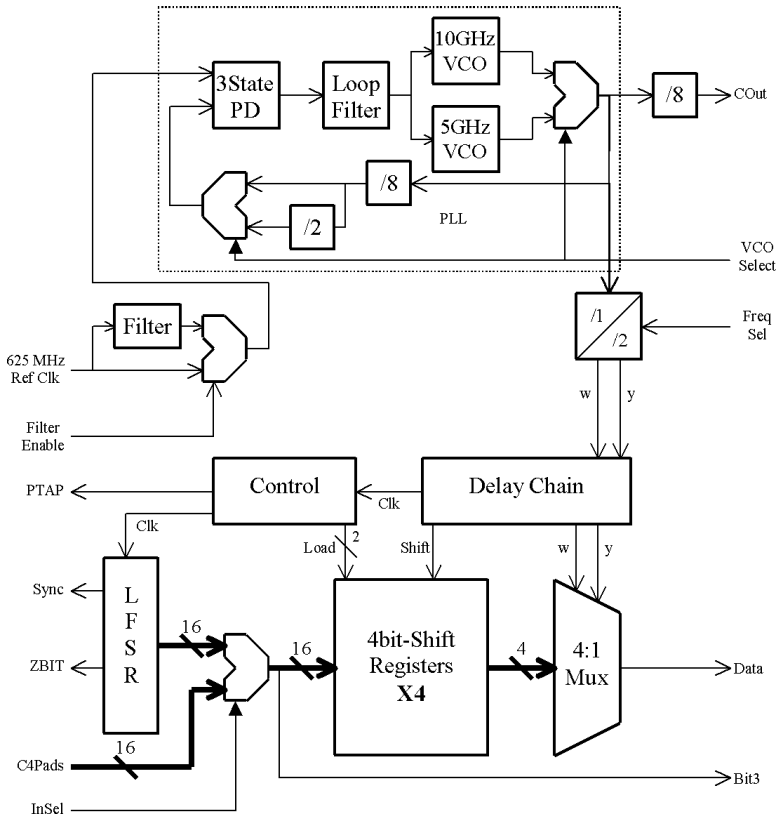

Figure 2.24: SERDES II transmitter block diagram.

tiplexer was added to choose between the LFSR data and external data that would have been delivered via the C4 pads. Although this revision didn't achieve near 40Gbps operation, it did generate outputs well over 20Gbps with greatly reduced quality(jitter). The limiting factor is believed to be in the shift register timing rather than in the main transmitter itself.

2.2.4.1 Transmitter LFSR

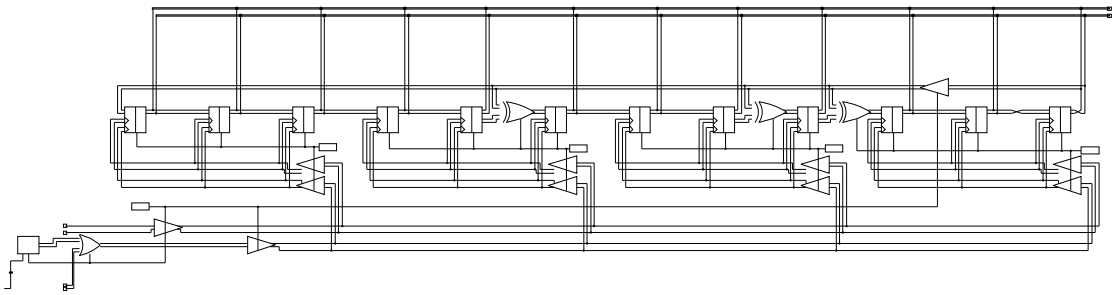

Figure 2.25: SERDES II transmitter LFSR schematic.

Figure 2.25 shows the schematic for the SERDES II LFSR. As discussed earlier, it is a Galois form 12-bit maximal length LFSR. As much of this design is reused in the Receiver LFSR, this design features a global reset. Because of this, the XOR gates cannot easily be incorporated into the current trees of the latches. There are several feedback paths, necessitating the use of the upper buffer. This has a negative impact on the maximum operating rate. One design innovation is that the final latch operates in a complementary manner. The input and outputs use wire inversions so that when the reset is active, it is effectively set to the opposite state of all the other latches. Also, as the LFSR will not operate in the all zeros state, the chances of all the latches being in the same state on power-up will be reduced if one or more of the latches is used as the opposite value, even though the physical layout was the same.

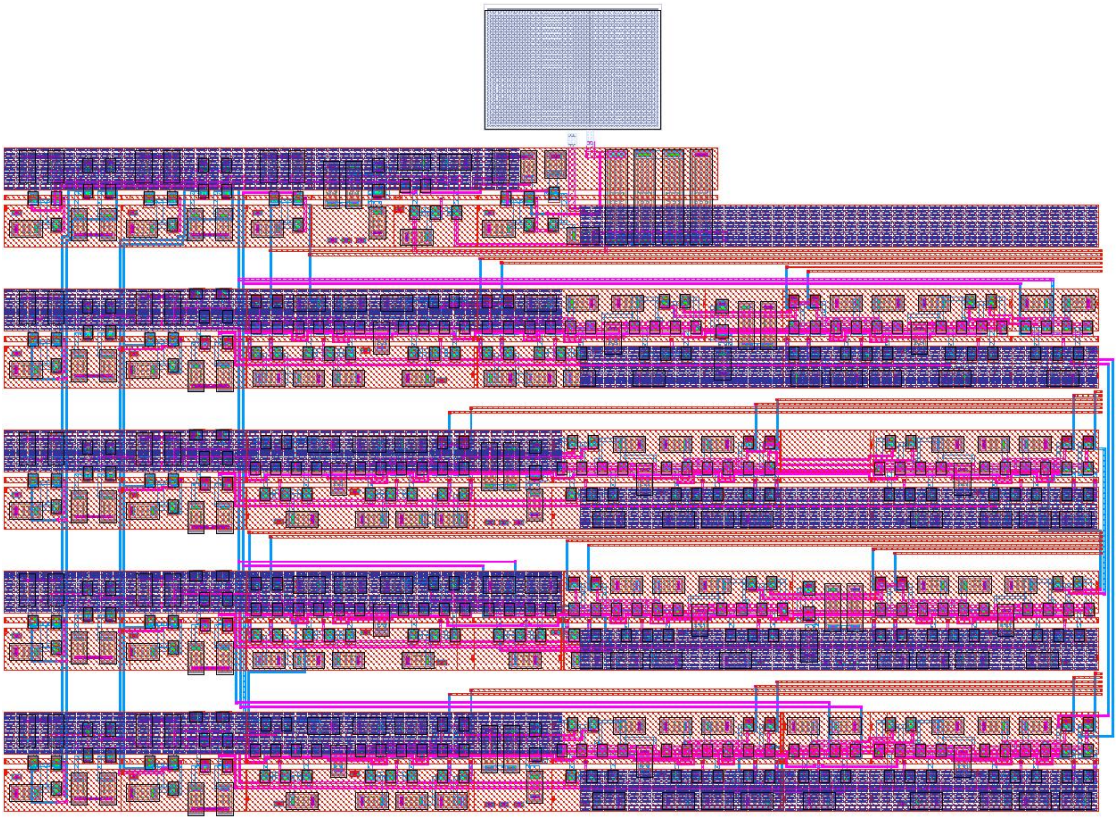

Figure 2.26: SERDES II transmitter LFSR circuit layout.

Even should the circuit power up in the all zeros state, several remedies were designed in. As can be seen in the layout in Figure 2.26, a capacitor was used as part

of an RC delay which automatically strobed the reset on power-up. In addition, the same input select, "InSel", signal which runs the multiplexer to choose the LFSR over the C4 pads as transmitter inputs also triggers this reset. (See Figure 2.24.) All the input buffers for the clock and reset are on the left side. In the layout, all output signals exit to the right where they enter the input selection multiplexer, which consists of 16 individual 2-to-1 multiplexers. The latches which form the LFSR are arranged in a serpentine pattern over four rows so that no feedback path would extend over the whole width of the LFSR layout.

2.2.5 Receiver

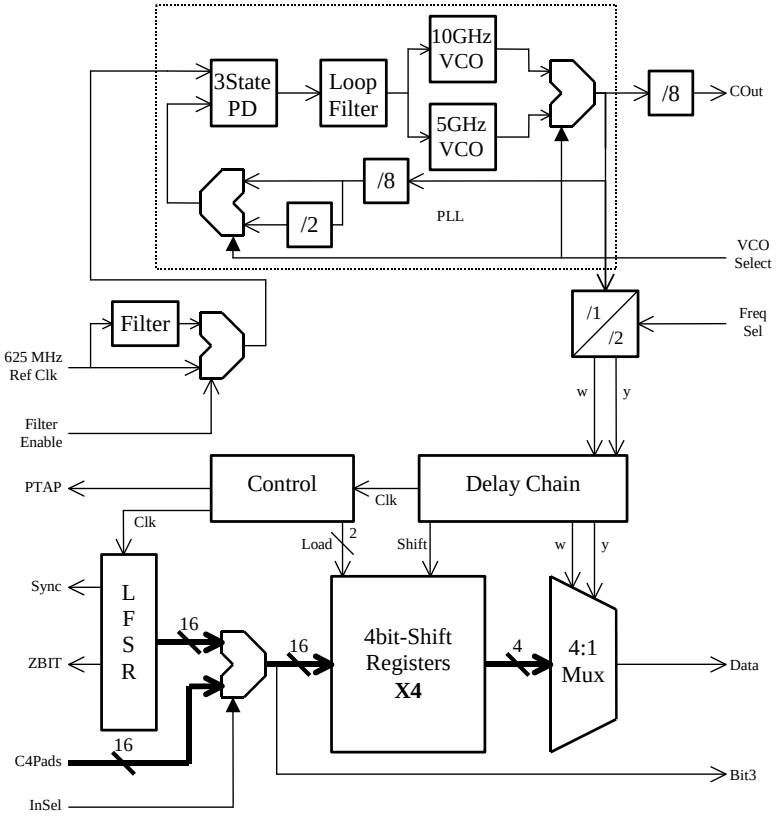

Figure 2.27: SERDES II receiver block diagram.

The first part of the receiver, (Figure 2.27), is largely identical to the design in the first prototype. The input select lines can be used to choose among three internal VCO generated periodic signals, or the external data input. Like the transmitter, the receiver benefited from a new active loop filter, and used an op-amp

based filter/integrator to replace the more problematic charge-pump design. Also, a simplified phase detection scheme was incorporated that reduced the delay between the sampling used for detecting the phase, and the signal passed on to the filter.

2.2.5.1 4-to-16 Demultiplexer

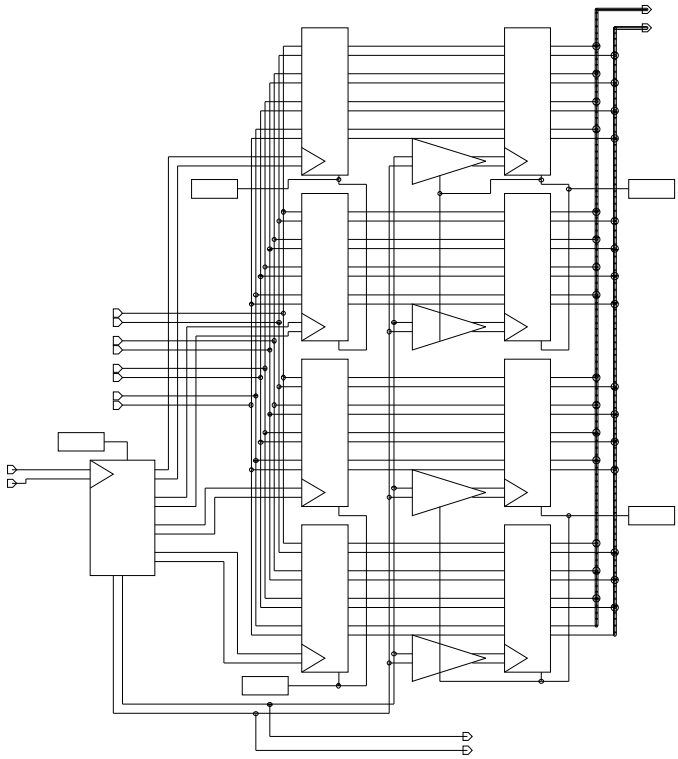

Figure 2.28: SERDES II Receiver 4-to-16 Demultiplexer.

A final 4-to-16 demultiplexer was added, (Figure 2.28), to enable output to the full 16 parallel lines. This consisted of 4 four-bit registers which were latched by in a round robin fashion, followed by a full 16 bit register latched via a single clock.

The control block for the 4-to-16 demultiplexer is seen in Figure 2.29. This basically consists of a two bit counter driving individual state detector gates. A delayed version of the last loading signal is used to trigger the final 16 bit register latch and also provides a signal to the next circuit block that new data is available.

Shown in Figure 2.30 is the full 4-to-16 layout. The 4-bit signals enter from the bottom center, directly from the receiver capture circuitry. The 16 bit signals exit at the upper right. The timing circuitry is located at the bottom, feeding the

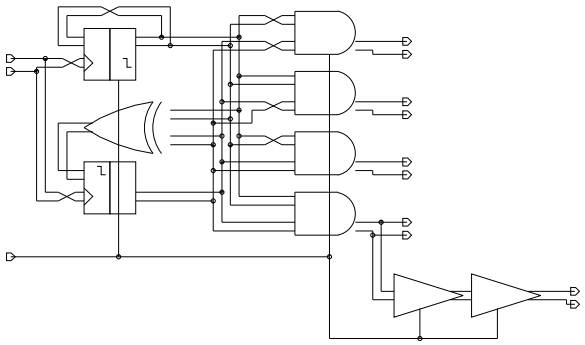

Figure 2.29: SERDES II demultiplexer control circuit.

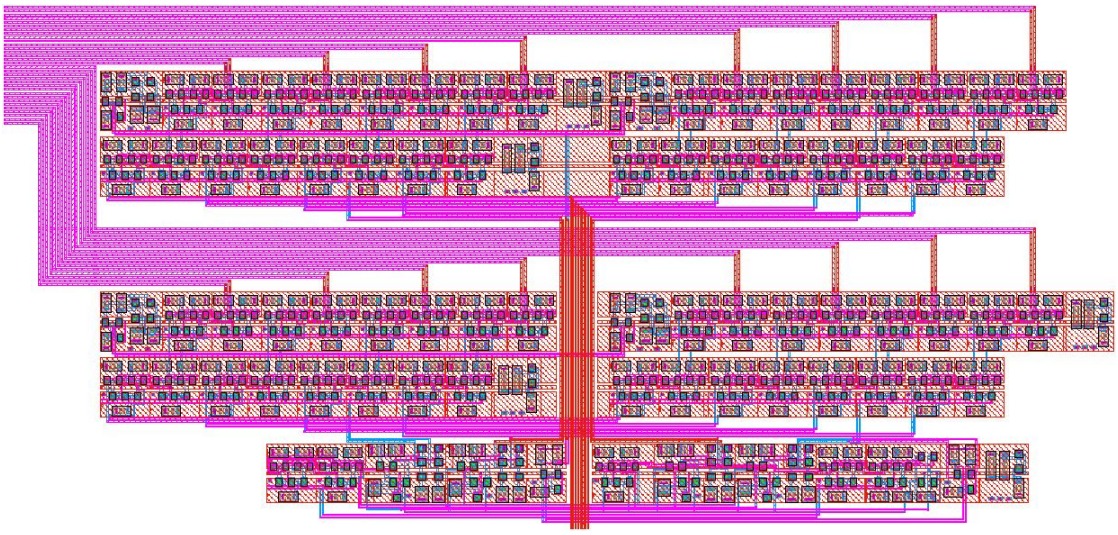

Figure 2.30: SERDES II receiver 4-to-16 demultiplexer layout.

groups of four bit latches above.

2.2.5.2 Receiver LFSR

Also introduced in the SERDES II receiver design was a 12-bit state machine recognizer to match with data from the transmitter's LFSR, intended for the detection of true bit error rates (BER). This LFSR shared many elements with the transmitter LFSR, but only a single output bit was used to compare with incoming data.

The receiver bit error testing circuitry, (see Figure 2.27), allowed one parallel output bit to be selected, (via SelTstA & SelTstB), and sent to the LFSR-based state machine recognizer, Figure 2.31. This selected bit could also be monitored on

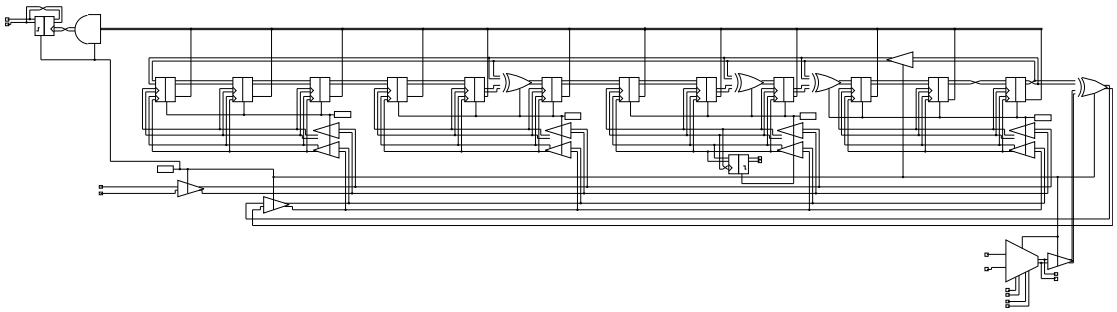

Figure 2.31: SERDES II receiver LFSR schematic. Acts as a state machine recognizer to detect correct bit patterns.

an output pad, "Tstb". On a bit miss-match, the recognizer would reset the LFSR and this would be detected on the "Rst" line, after passing through a divide-by-two circuit. If the recognizer passed through an entire sequence of 4095 bits, the overflow would cause the Good output to be toggled. At 20Gbps data rates, this would generate a square 150KHz output signal, easily detectable/countable by our equipment. The schematic in Figure 2.31 includes the normal Galios form LFSR, and clock and reset buffering as did the LFSR in the transmitter. Also shown on the lower right is the input multiplexer circuitry to select testing bits. On the upper right is the exclusive-OR used to detect bit errors and trigger a reset. The output signal which indicates a reset has occurred is near the center of the LFSR, where the buffered reset signal is tapped and sent to a divide-by-two. Lastly, on the upper left is the circuitry to detect the "roll-over" state to indicate that a full 4095 bits have been detected.

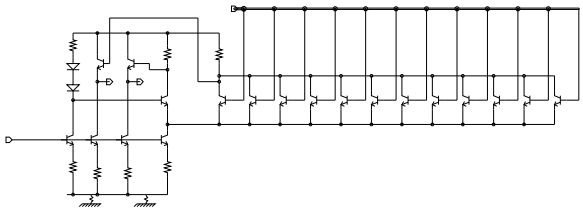

Figure 2.32: SERDES II LFSR high fan-in AND gate.

In order to detect the "roll-over" state, a high fan-in gate would have to be designed. Due to the nature of CML and our power supply, a single gate can normally only have a fan-in of three. To work around this, a large fan-in single-

ended circuit was designed, and is shown in Figure 2.32. This circuit is effectively a large OR, but it is used as a NAND by using complements of the intended inputs. The differential buffer is supplied with a common mode level 3 signal generated on the left, while the individual transistors on the right can divert the current to the other side. The "roll-over" signal passes through a divide-by-2 circuit before going to the output so that a square wave would be produced, rather than a pulse.

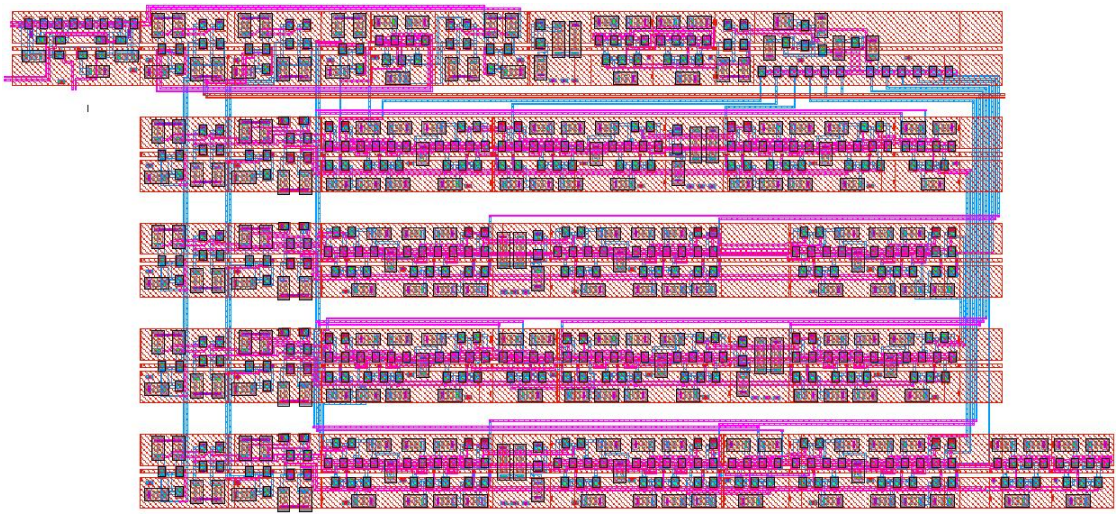

Figure 2.33: SERDES II Receiver LFSR circuit layout.

Figure 2.33 shows the layout of the receiver LFSR. This is quite similar to the design in 2.26, with the major differences being the 4-to-1 multiplexer to select an input on the upper left, the high-fan-in AND at the upper right, and the divider for the reset at the lower left.

2.2.6 SERDES II Receiver/Transmitter Chip Layout

The SERDES II chip seen above in Figure 2.34 has C4 pads as well as wirebond, as it was envisioned that it might eventually be packaged. We added the necessary circuitry to handle a full set of 16 parallel inputs and outputs via the C4 pads, leaving only a few of the signals accessible over our bondpads, which we could probe prior to considering the chip for packaging.

The total chip size was approximately $12.5mm^2$ in size, however much of this large area was unused by the serializer, deserializer, and test circuits. The large

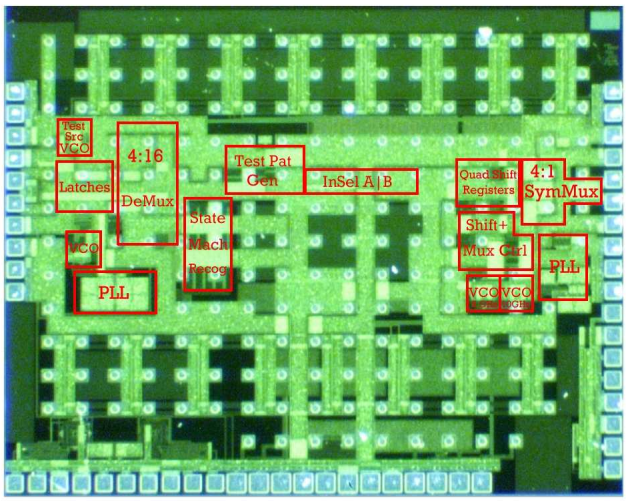

Figure 2.34: Layout of the SERDES II prototype 20Gbps chip.

size was required for the necessary C4 pads, and some of the unused area was appropriated for the test circuits of other group members.

2.2.6.1 Results

Several errors were made in the design of this chip which might have been caught had we had more time for testing and layout. A small mistake in the receiver loop filter greatly reduced its lock in range. (A node in between two capacitors was grounded when it should have remained floating.) Parasitic effects once again had an impact on the VCO center frequency that did not show up in simulation. With these problems we still achieved transmit capability of 14.27Gbps to 21.58Gbps, and the receiver worked from 16.8Gbps to 18.5Gbps.

We were able to rent test equipment and obtain the spectral VCO data presented earlier. In addition, we used our sampling scope to produce numerous eye diagrams. In Figure 2.35, the eye on the left was directly captured output from our transmitter at 20Gbps. The right eye was produced by a commercial group using 7HP technology that has a much, (over $2x$), higher f_T. We obtained a 400mV output swing on each line of a differential pair yielding a difference swing of 800mV.

When a dual-eye plot of the transmitter output is taken, duty cycle is only minimally apparent, as is seen in Figure 2.36. Thus, the new symmetric multiplexer

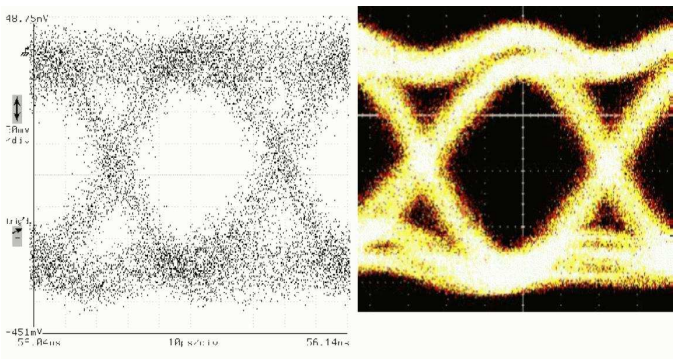

Figure 2.35: SERDES II 20Gbps eye diagram on the left, compared with a commercial 20Gbps eye in a technology with twice the effective performance.

has gone a long way towards alleviating problems with our "shuffling" quarter-frequency transmission scheme. However, we later found that due to small design errors, the performance could have been even better.

Our results regarding the VCO, transmitter multiplexer scheme and circuits were written up and submitted to the IEEE Journal of Solid-State Circuits (JSSC) for publication. As of this date it was not accepted, possibly again due to the mismatch of transmitter and receiver rates. We are seeking the opportunity to publish the bulk of this research elsewhere, after careful review and re-edit.

2.2.7 Second Prototype Corrections

Various problems were identified within the second prototype, almost all of which are correctable in future revisions. Some of these would have been dealt with before fabrication had there been sufficient time for more testing. We were again somewhat disappointed that our VCOs under-performed our simulations, and have adjusted our simulation techniques accordingly.

In the receiver, as mentioned earlier, a loop filter error reduced the lock in range. Also, there were several areas where loading effects caused improper operation at higher data rates. When the designs were re-simulated using more accurate parasitic modeling, interconnect capacitance apparently was such that some line drivers were unable to drive the lines with sufficient speed to ensure correct operation at or near the upper limits of the design specification. This was especially

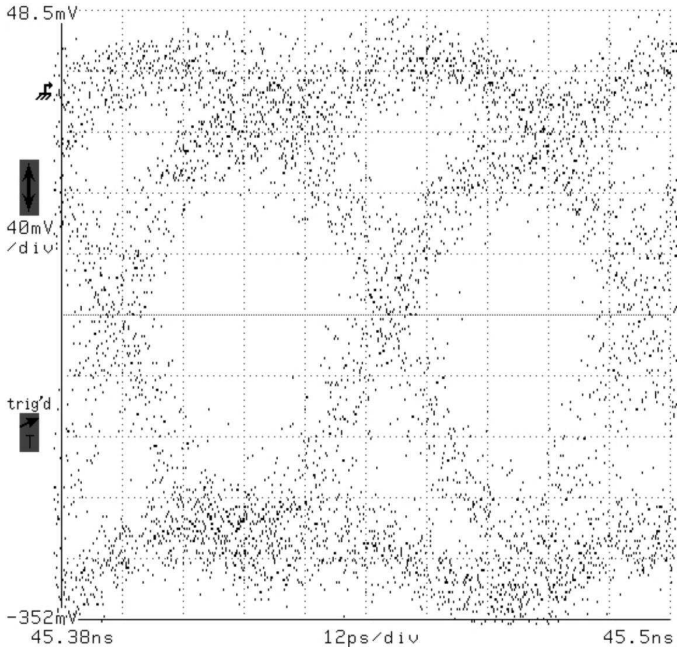

Figure 2.36: Dual eye diagram showing small duty cycle realized with the symmux.

apparent in the 4-to-16 receiver demultiplexer, where each driver had to drive four others as well as all the interconnect, (which turned out to be quite a lot.) The circuitry to handle 16 bits takes up quite a lot of real estate, at least when compared to our previous designs. It may have been mistakenly assumed that there would be intermediate buffers, as these blocks were laid out by different researchers. A similar loading problem also occurred in a delay chain that drove the LFSR in the transmitter. Above a certain frequency, the LFSR produced no output. Simulated at low speeds, or when modeled without parasitics, the circuits performed correctly. In order to fix these problems, the necessary undersized drivers were modified, and the circuits tested using full parasitic simulations at the desired working speed.

Other problems that have been identified include a miss-match in the circuits driving the final transmitter multiplexers. For complete uniformity of operation, all inputs should be driven at the same common-mode level and by emitter follower transistors of the same size. It was found that the outputs of the four 4-bit shift registers, (which feed the final 4- to-1 multiplexer circuit), used $1\mu m$ transistors while the clocking signals to the select lines were $4\mu m$. The clocking signals were

split giving an effective driving capability of $2\mu m$ for each, so the outputs of the shift registers should have been $2\mu m$ as well. In addition, the transmitter also had a hard-wired constant value connected to a "dummy" multiplexer that was intended to only provide a delay. This constant was also at an improper level. As with the other problems, work-arounds have been developed and simulated.

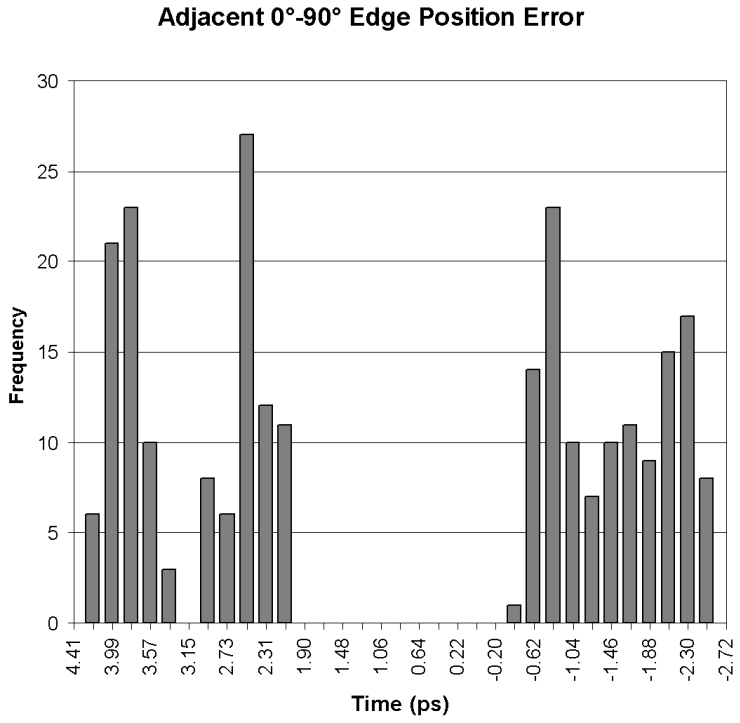

Figure 2.37: Simulated edge position frequency distribution for corrected SERDES II under ideal conditions.

All together, these errors in the symmux led to a visible duty cycle. Figure 2.37 above shows predicted behavior of a "corrected" symmux under ideal conditions. The average error is less than 1ps, which is far superior to that of the original prototype employing regular CML multiplexers. After "correction", the standard deviation of the error has been reduced by 45% to 2.43ps, when compared to the "uncorrected" circuit that was actually fabricated. However, this standard deviation is still much larger than that of the simple CML multiplexers and so is far from ideal. The spread is mostly due to the reduced bandwidth of the symmux giving rise to data-dependent jitter. Other problems with the circuit are its potentially

irregular rise and fall times. Because pairs of independent transistors act on each output line, conditions can arise when they are not in the same state, or not even in a fully differential state when compared to the transistors on the complementary output line. Since some of the preliminary buffers have different internal loading, the propagation time for the select line is still slightly different from that of the data lines.

Some of the most puzzling deviations from ideal behavior were investigated to see if they were due to data-dependent jitter. It was initially found that the "corrected" circuit still exhibited multiple peaks in the distribution for the odd edges. To better visualize what was occurring, the idea was conceived of using a phase plot, distributed over four bit times, so that the edges are automatically partitioned by the hardware that generated them, (Figure 2.38). You can see the output below, which is quite interesting. This is a standard XY cartesian polar figure, with zero degrees on the far right. Ideally, all edges would occur on an axis. Bit A occurs in the upper right quadrant. Bit C is in the lower left, etc. The A-B and C-D odd edges with their two error peaks are clearly visible at the top and bottom.

It is quite interesting that there is a complete absence of almost vertical edges in an X-pattern connecting opposite peaks on top and bottom! This indicates that if an odd edge is early, the odd edge after it is always late. The converse is also true. If an odd edge is late, the odd edge after it is always early. In order to determine if the edge delays were dependent on the bit values immediately preceding them, the same data was plotted, but with the radial distance varied based on the number of identical bit values preceding the edge. This is shown in Figure 2.39 below.

In this plot, the edge placement error is the angle between the point and the nearest axis. If the error grew with longer sequences, the edge marks would tend to curve away from the axis as the radial distance increased. Alternately, if the errors were due to rapid alternation of bit values, the edges closest to the origin would be furthest from the axis, and the edges would curve back towards the axis as radial distance increased. As can be seen, the edges appear along radial lines indicating there are virtually no dependencies on the length of same-bit sequence preceding

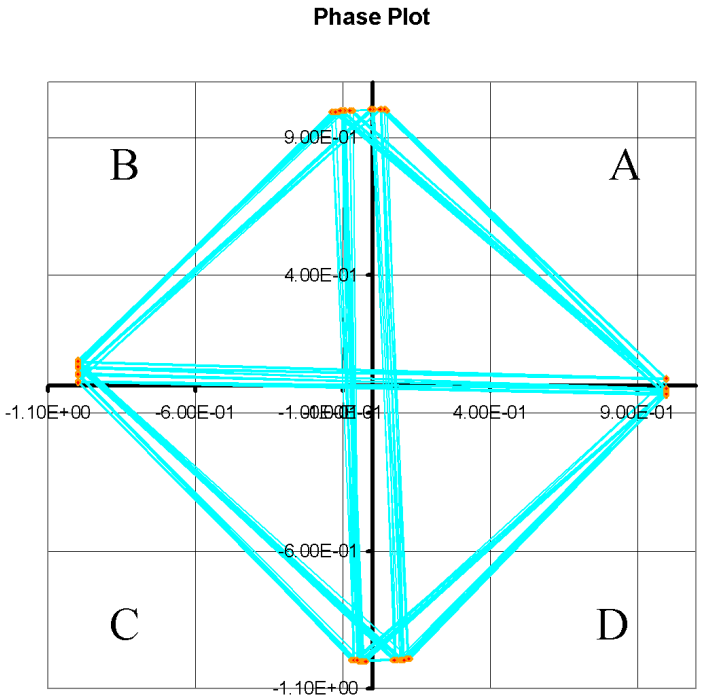

Figure 2.38: Plot showing edge placement errors.

the edge! This was quite unexpected.

The edge data was finally analyzed and explained using a C-program written that correlated each edge error with the hardware that generated it, the sign of the error, and the full n-bit pattern that preceded it. It was found that there was a strong disparity between a bit sequence and an identical except complementary sequence. For a fully differential circuit this should **not** be possible.

The difference in 0/1 behavior was finally tracked down to self-heating effects in the HBTs which make up the symmetric multiplexer [14]. In our multiplexer architecture, we used "dummy" multiplexers to match loading and delay, and these used fixed value "constants" at some of the inputs. The adverse effect of the constant was to bias one side of a differential pair more heavily than the other, increasing its local temperature and altering its behavior with respect to rise and fall times. The self heating model in VBIC is implemented as a thermal capacitance and resistance yielding an effective time constant for dynamic heating behavior. In the IBM kits, the thermal resistance decreases proportionally with area, while the capacitance

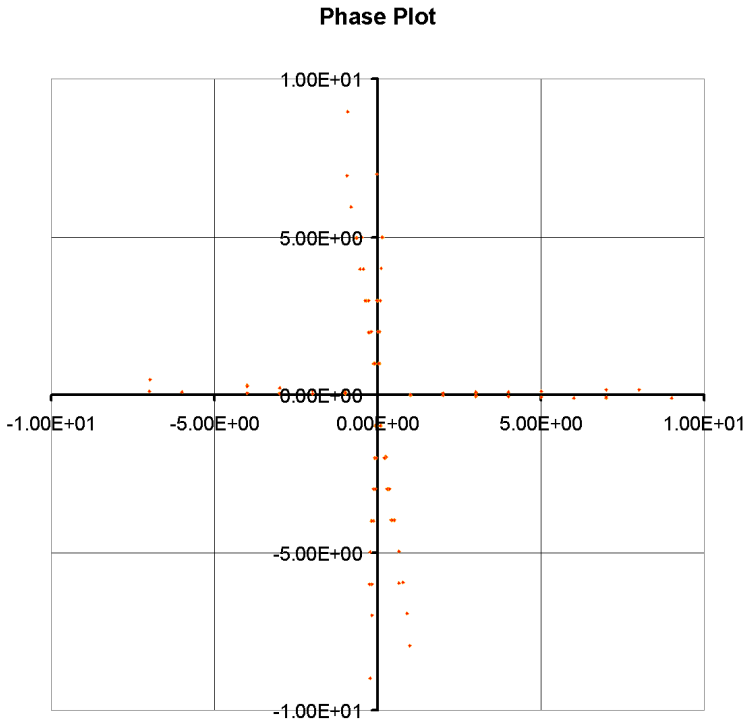

Figure 2.39: Radial phase plot showing dependencies between data and edge placement error.

linearly increases. This leads to a fixed time constant of around 1ns for all the npn HBTs. A 1ns time constant means that after 1ns, a temperature imbalance might decay to 63% of its initial value if the circuit was not otherwise disturbed.

This has important implications for simulation, as the initial HBT temperature is derived from the initial conditions for a transient simulation. To get accurate measurements, the differential test input voltages should be zero at $t = 0$ to prevent an initial temperature bias that could potentially last for many dozens of bit times.

One last area of potential problems in the second prototype to be addressed was in the transmitter state machine, (Control), which initiated the loading verses shifting operations of the four shift registers. This state machine used a single gate delay feedback path with master-slave latches for all the state bits except one, which needed to be a half-period delayed copy of one of the other outputs. This was accomplished by inverting the clock to a master-slave latch driven by the output to be copied, but this created a half-period propagation criteria whereas the rest of

the circuit could make use of the full period. By re-implementing the last bit using combinatorial logic from just the previous state, and retaining the inverted clock, the potential top frequency of this state machine was effectively doubled.

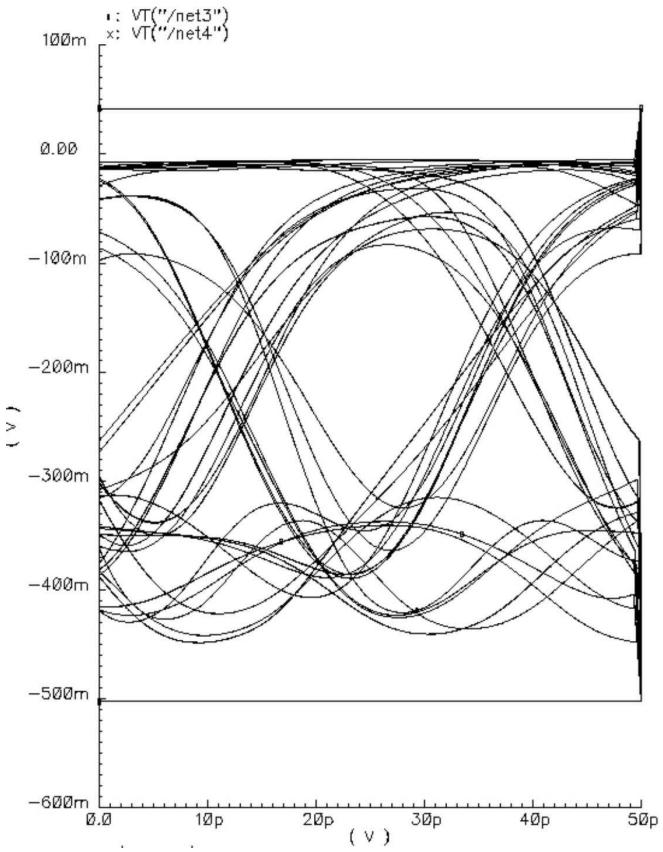

Figure 2.40: Simulated corrected SERDES II multiplexer output at 40Gbps under ideal conditions.

Figure 2.40 is a rough simulation of how our second prototype SERDES output multiplexer might look at 40Gbps, after the corrections, using ideal data and clocking. The output was not physically achieved, but there are no fundamental reasons why we should not be able to reach this goal with future designs. Of course more work does need to be done to "tune" our designs for operation at these frequency extremes. The complex circuitry of the symmux decreases the bandwidth of the output circuitry to the point where it introduces a large data-dependent jitter, which would seem to make reliable operation impossible using that design. Future work will have to incorporate finding a way around this limitation, as well as reducing or completely eliminating the duty-cycle from the transmitter output.

CHAPTER 3
Voltage Controlled Oscillators

This section of the Thesis details the three new voltage controlled oscillator designs which are used in the serializer integrated circuits. These are four-phase ring VCOs with various methods of controlling frequency. The first was created in a 5DM process and fabricated in a test chip. The second was incorporated in the first SIII 5AM serializer chip. The final design was created in 5AM, and was instanced in two different forms for different ranges of operation for the final serializer revisions.

Simple bipolar ring VCO design is covered in [21], where each stage is considered to be a two-pole amplifier which has to meet the Barkhausen criteria of unity gain and 180° phase shift around the loop. This small signal approach ignores the large signal effects which occur as the current is steered in a differential pair. Noise analysis is also presented, but the paper notes that many approximations need to be made, resulting in more qualitative information than quantitative.

3.1 5DM Dual Control VCO

In order to achieve higher data rates to fully test the new edge steering transmitter architecture, a new VCO design was required. In keeping with earlier designs, multiphase clocks would be required to achieve the desired high data rates. A group of variable delay elements would be connected in a ring configuration along with an inversion. Four phases would be sufficient if the clock rate of the phases could reach as high as 10GHz. This was considered to be achievable as several previous test VCOs in the SERDES II series had operated at nearly that rate. With a 10GHz VCO, phases would be available with 12.5ps spacing. Note that this is actually *less* than the delay through a typical lightly loaded buffer. Therefore this VCO would need to make use of a phase interpolation architecture to order achieve this rate, as was used in the previously described SERDES II design.

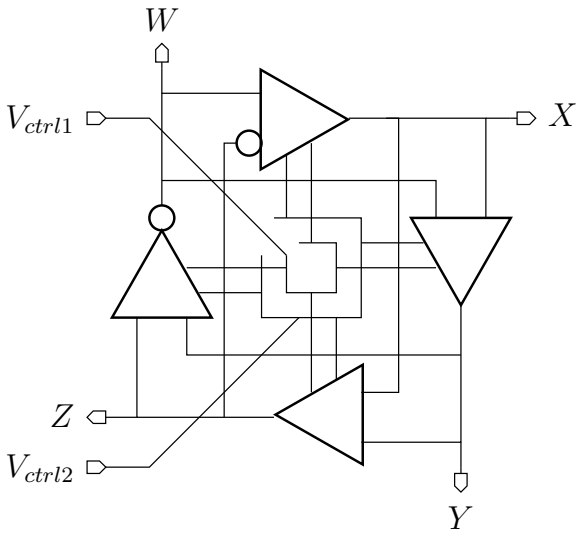

Figure 3.1: 5DM VCO Architecture. Each delay element has two phase inputs and the output phase is a value interpolated between them. V_{ctrl1} is a fine frequency control implemented via varactors while V_{ctrl2} is for coarse frequency control which changes bias currents.

3.1.1 Design

The FFI buffer architecture of the SERDES II,(see Figure 2.19), designs was a promising candidate, but concerns over the operating frequency dependent noise characteristics [30] was the impetus for choosing something a little less radical. A 50/50 **fixed** interpolation VCO, (see Figure 3.1), was designed to use varactor tuning. In the case where true varactors are not available, reversed-biased transistor PN junctions can be used instead. A delay element of this type is shown in Figure 3.2.

The emitter resistors linearize the buffer so that smooth phase interpolation can occur, as well as increasing the buffer bandwidth. The varactors were chosen to be $40\mu m^2$. This value is a tradeoff between highest achievable frequency and frequency range. Because the varactors can only roughly double their smallest value, adding them always decreases the highest available frequency. The inputs a_x and b_x, and the z_x outputs are all at level two. The control voltage V_{fine} should therefore remain below the lower level two value of approximately -1.15V and can go as low as -3V without exceeding the V_{BE} breakdown. In practice, it can go slightly above, as long as the junctions do not become forward biased enough to admit significant

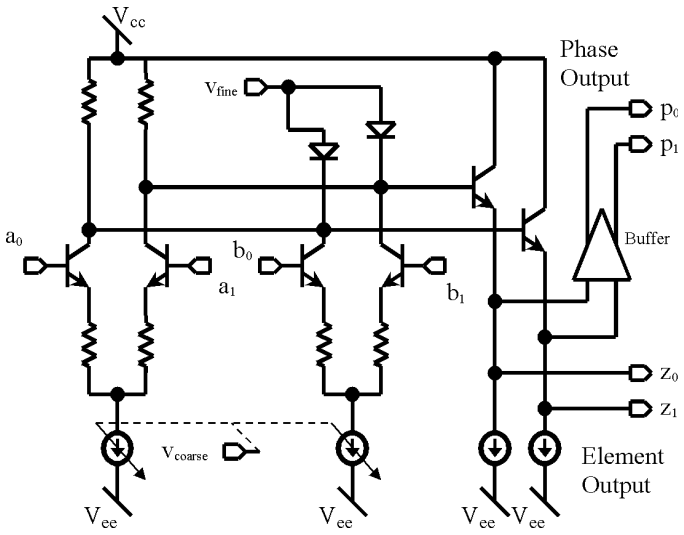

Figure 3.2: 5DM varactor-controlled VCO buffer element. This design uses fixed leap-forward phase interpolation. Fine control is achieved via V_{fine} reverse-biasing the PN junctions of the varactor diodes above the second buffer.

current.

The layout considerations of this VCO are similar to that of previous designs in that a high degree of symmetry is required. This VCO also incorporates the necessary buffers in the proper locations to make connections to the top tier multiplexers in a symmetrical manner.

Normally, the range of a varactor-controlled VCO is fairly small, but with the addition of the coarse biasing control it should allow a fairly wide frequency range to be explored when the circuit is incorporated into a transmitter. This enhancement is shown in Figure 3.2. This design used $4\mu m$ transistors for the interpolating trees, and $6\mu m$ transistors for the emitter followers. This delay element features linearized inputs which are evenly interpolated for a 33% speed increase over a similar non-interpolating buffer. The emitter degenerating resistors were 104Ω. The coarse control input varies the current in the input trees, affecting the bias.

This is effected by varying the voltage supplied to modified right-half current mirrors, one per delay stage. The circuit used to achieve this is shown in Figure 3.3. The circuit is nominally biased as an $8\mu m$ mirror half, and there are also regular $8\mu m$ mirror halves used to bias the non-variable parts of the stage. By decreasing

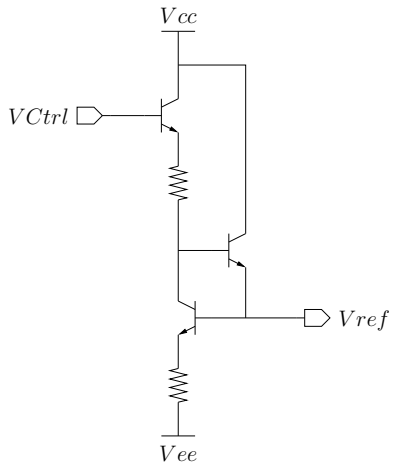

Figure 3.3: Variable voltage reference for coarse frequency adjustments.

the current below that required for maximum f_T, the frequency can be reduced. This is intended to allow us to control the VCO center frequency precisely, and eliminate the center frequency mismatch of previous SERDES designs. With this design, changes of hundreds of millivolts at the control input cause changes of only a few millivolts at the V_{ref} terminal. Note however, that this modifies the magnitude of the voltage swing developed across the input buffer top resistors. This effect is virtually eliminated by the output buffer which also ensures that uneven loading and interconnect doesn't affect the VCO circuit symmetry. The buffer is of the ordinary CML type with emitter followers, giving outputs on level two. The output buffer uses a $4\mu m$ differential pair, followed by $6\mu m$ emitter followers. The inputs to the buffers are also at level two in order to reduce the capacitance from the reverse biased base-collector junctions. Fine control of the VCO is achieved by the use of varactors which load the output of the linearized input buffers. They feed a pair of emitter followers which are biased normally, and these in turn drive other element inputs as well as the output buffers.

In Figure 3.4, the full-parasitic simulated VCO frequency is plotted against both the coarse control voltage and the fine control voltage. The center frequency of approximately 8.8GHz occurs at a coarse control voltage of -130mV and a fine control voltage of -2V.

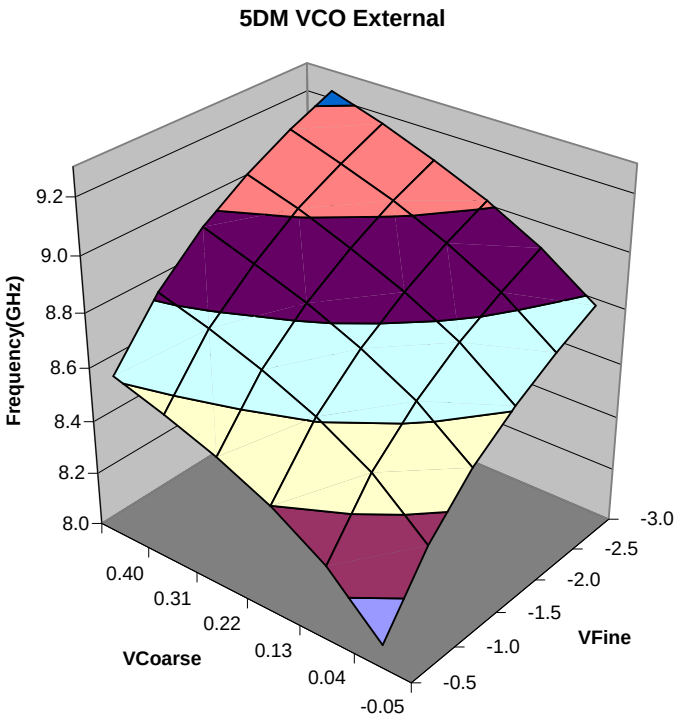

Figure 3.4: Simulated 5DM Voltage control plot. Fine control gain is approximately 210MHz/V near the central region.

3.1.2 Layout

We obtained the opportunity to make use of unused space in a MOSIS 5DM fabrication run. A test VCO as well as supporting circuitry was laid out and submitted within a short deadline period. The 5DM IBM process has MOS-based varactors as well as PN types which are doped to provide a more linear capacitance response to the control voltage. PN varactors were used in this design. The VCO was laid out in a highly symmetrical manner as can be seen in Figure 3.5. The majority of the layout was done using the first three 3 layers of metal. The additional metal layers were primarily used to deliver power evenly via rings and route the fine control line in a symmetrical fashion, leaving through the top center of the VCO. The size of the VCO layout was $300\mu m\ x\ 300\mu m$.

The voltage references for the current sources at the bases of the current trees are distributed around the perimeter, and tied together, to provide the same voltage for each element. Each element contains an output buffer with connections near the perimeter. It is possible that this layout can be reduced in size by an almost factor

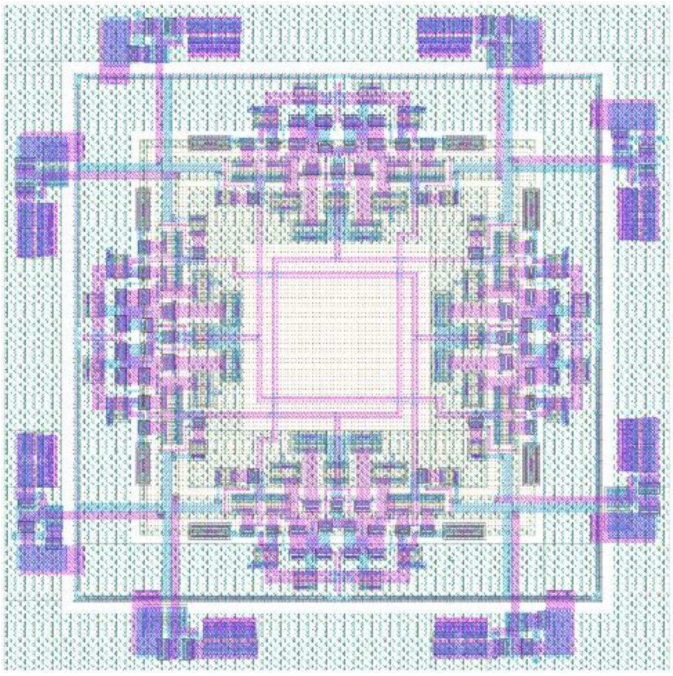

Figure 3.5: 10 GHz test VCO layout. The VCO is located in the upper left corner. Dividers and pad receivers and drivers are located along the lower edge.

a factor of two, but it was though that this might prove detrimental to the circuit's performance. Although not easily seen, there are several wired inversions between the delay elements along the centrally routed lines. These inversions, because they are not present on all lines, present an asymmetry. However, if the parasitic effects of the inversions are miniscule with respect to the common parasitics associated with the lines, they should have no real adverse effect on performance. If the layout were somehow to be drastically reduced in size, the inversion parasitics would grow proportionally larger and eventually would result in a phase mismatch.

The VCO design was embedded in a test chip layout, (Figure 3.6), containing pad drivers and receivers, frequency dividers, and control circuitry for both coarse and fine adjustments. The chip area was $1500\mu m\ x\ 675\mu m$. The chip was designed to be tested using our standard probes, the 10 pin, (6 signal), GGB model as described in Appendix E. To make testing simpler, the frequency dividers were used to reduce the frequency observed by the oscilloscope. Two separate divide-by-8 circuits were used, each consisting of three identical chained divide-by-two circuits.

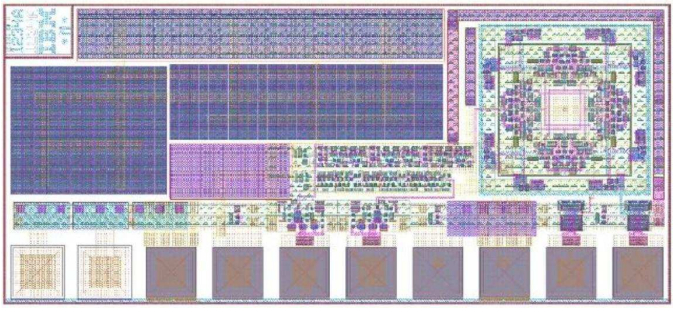

Figure 3.6: 5DM 10GHz VCO test chip microphotograph.

The divide-by-two circuits are also a critical component in most of our designs, so it was beneficial to also investigate whether they would operate at the targeted frequency.

3.1.3 Testing

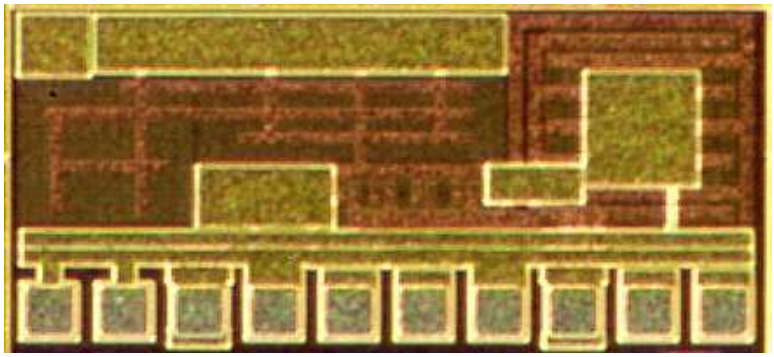

Figure 3.7: Microphotograph of the 5DM 10GHz VCO chip.

The 5DM VCO chip was manufactured as shown in Figure 3.7. It arrived as part of a wafer and we were able to successfully probe the circuit and capture output. It was tested using our probe station and Tektronix 11801C 50GHz bandwidth oscilloscope. A LabVIEW^{TM} program was written to control a GPIB-enabled power supply to adjust control voltages and record the data from our the oscilloscope.

A sample of the captured data is shown in Figure 3.8. The measured frequency values were captured from the divided by eight output, then multipled by eight for this plot. The simulated values were obtained without using extracted parasitics. When parasitics are included, simulated and measured results closely conform within

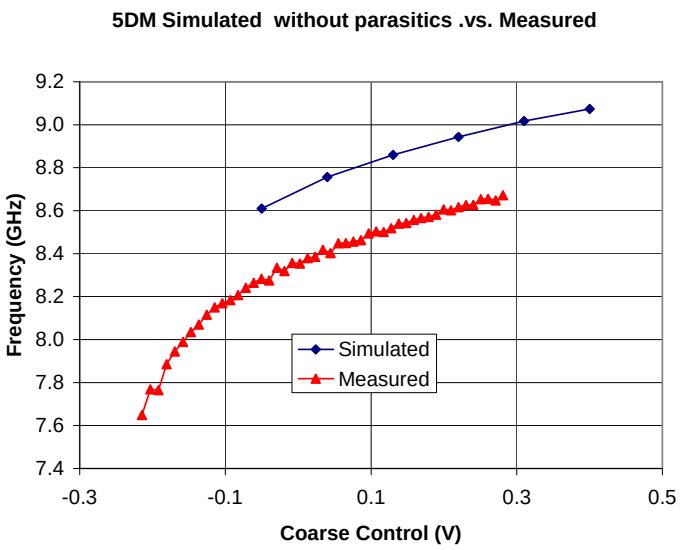

Figure 3.8: 5DM VCO coarse control voltage.vs.frequency, both simulated and measured, with fine control held fixed at -2V.

expected values given uncertainties in operating temperature. The effective slope is approximately 140MHz/V as expected from simulations. The input voltage is the rough control with fine control held fixed at a value of -2.0V.

Figure 3.9 shows data captured from the 5DM 10GHz VCO divided-by-8 output. The period is approximately four 200ps divisions, and has a 200mV peak to peak single ended output. The 800ps period corresponds to a 100ps 10GHz period of the internal VCO.

In the test chip, the coarse control would be used to ensure the PLL is well within lock-in range. The coarse control is extremely coarse, with a slope of over 10GHz/V when operated at lower frequencies. The fine control gives approximately 210MHz/V near the desired target range.

Unfortunately, the design of the test circuitry surrounding the 10GHz VCO used an improper center bias, which prevented the capture of data over the entire usable frequency range. This problem was traced to a last minute design change which involved a tail current source resistance with an incorrect value. Evidently, the desired resistance value was small, and due to layout reliability considerations it was felt that two parallel resistors of twice the target value would do the job better. Unfortunately, although the layout did have the two resistors, the original value was

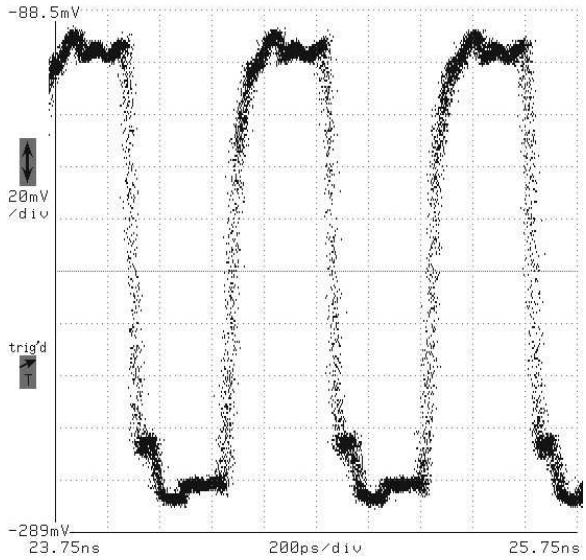

Figure 3.9: 5DM 10GHz VCO divide by eight output.

used in both. This had the effect of limiting the range of the coarse input control voltage that could be applied.

Phase noise was simulated for this oscillator using SpectreRF. The operating frequency was 9GHz, and at an offset of 1MHz the simulation indicated a value of -84dBc/Hz using a fine control voltage of -3V, and a coarse control value of -2.43V(post reference circuit), at 25C.

3.2 SIII 5AM Oscillator

The SERDES III transmitter was originally planned for fabrication in 5DM. However, due to the unavailability of space on MOSIS runs, a choice was made to have the design fabricated in 5AM as that was the next available run after the 5DM. Layouts for the chip were converted to 5AM, which had a smaller number of metal layers. Lack of time precluded the creation of a 5AM version of the 5DM oscillator and a simpler and slower design which had been previously laid out is what was finally included.

3.2.1 Design

The 5AM oscillator is of the traditional "current starving" type first used by Samuel Steidl, a former member of our research group. A simple ring of buffers is

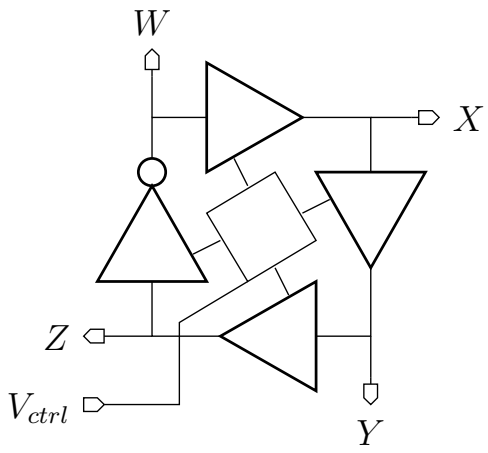

Figure 3.10: 5AM VCO Architecture. This is an extremely simple design using a ring of buffers with the tail current sources controlled by a common voltage.

used to to create four phases and their complements, as shown in Figure 3.10. Like the 5DM oscillator, it features a coarse control circuit which adjusts the voltage passed to the current mirrors.

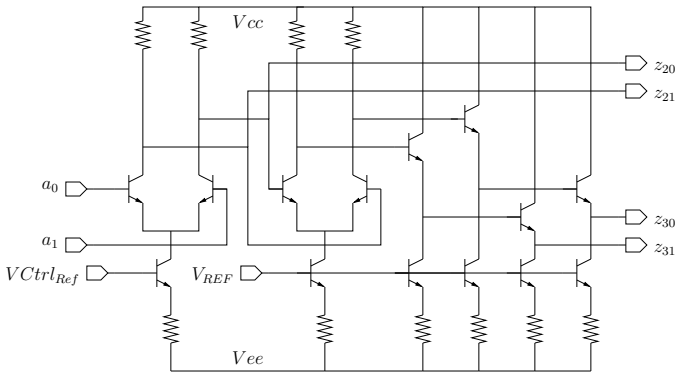

Figure 3.11: 5AM VCO delay element. The second buffer uses a fixed current bias to remove the amplitude variations caused by the varying bias of the first buffer.

As can be seen in Figure 3.11, the schematic of the delay element uses a second buffer to eliminate the problem commonly encountered with current-starving VCOs, in that with a varying bias, the top resistors produce a varying output swing. The second buffer regenerates the normal 250mV swing. The two sets of emitter followers are used to produce the level three output to drive the rest of the chip.

The buffers were designed using $4\mu m$ transistors, with the top resistors having

values of 135Ω and 104Ω respectively. The resistors in the first pair are larger to ensure a full current switch in the second pair, even when the current is reduced to as little as 1/3rd the normal 0.7mA/μm. As described in the CML logic Appendix A, 120mV is needed to swing 99% of the current. The emitter followers were 4μm and 6μm respectively. The output which feeds the next stage is not passed through a V_{swing} regenerating buffer, because to do so would cause twice as much delay. Because of this, the amplitude of the signal to the next stage is affected by the bias current.

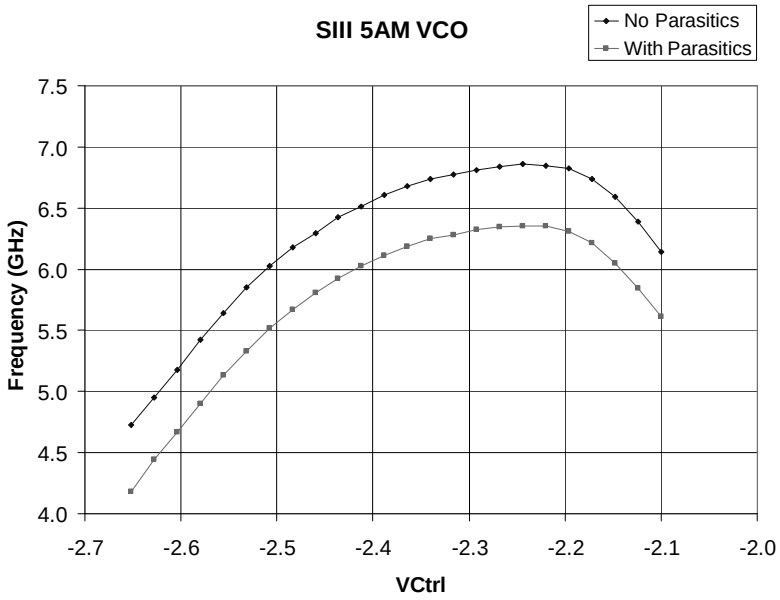

Figure 3.12: 5AM VCO Characteristic. Non-monotonic response typical of the CS VCO design. Simulations with and without parasitics are shown.

This circuit was simulated to yield the control voltage plot in Figure 3.12, (upper curve). As can be seen, the non-parasitic frequency peaks at about 6.9GHz at -2.24V. The non-monotonic nature of the response was the primary reason a PLL was not considered for this design. At the higher region, the characteristic is almost flat. The center of the usable region is around 5.8GHz with a ±1GHz range and a gain of 3.9GHz/V. This is markedly different from the specified gains for the other VCO designs presented because this control is effected via bias variance, while for the other designs the varactor fine control is specified.

Each of the four buffers in the VCO draws approximately 20mA, with a total

power use by the VCO of 290mW, not including the power used by the voltage reference circuits. With those circuits included, the total power is 340mW.

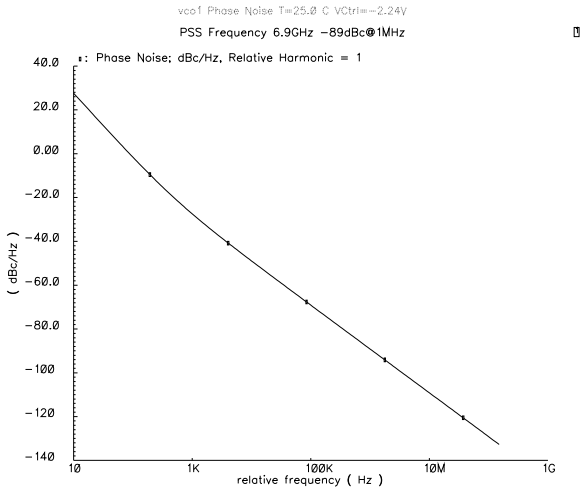

Figure 3.13: Phase noise of 5AM VCO.

Phase noise simulations were carried out using SpectreRF, using the periodic steady state(pss) and phase noise(pnoise) features. As can be seen in Figure 3.13, the phase noise at 1MHz offset is approximately -89dBc/Hz, and decreases at the expected 20dB/dec. The operating frequency was 6.9GHz with the control voltage set to -2.24V. Parametric simulations were performed to determine the relationship between phase noise and frequency. Simulations indicated that the phase noise actually decreased with increasing frequency. (-79dBc/Hz operating at 5.2GHz) This is contrary to expectations as the bias current increases with frequency, thereby increasing current or voltage dependent forms of component noise. The reason for this apparent anomaly is that the The carrier power frequency increases more rapidly than the noise power with increasing frequency, and the phase noise is in dBc/Hz, which is relative to the carrier power. The actual noise power does increase with frequency.

3.2.2 Layout

The reduced number of metal layers in 5AM made a tight layout using a leap architecture very difficult. By switching to a non-leap architecture, it was much easier to fit the VCO design into the power rail layout chosen for the transmitter.

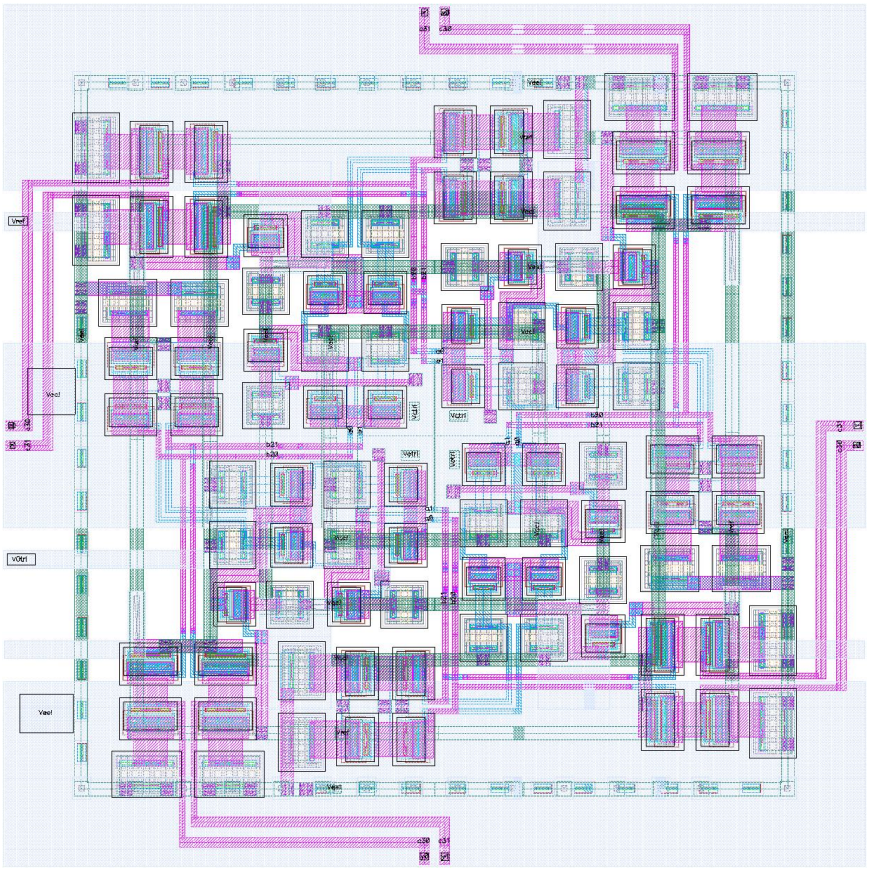

Figure 3.14: 5AM VCO layout.

The circuit is rotationally symmetric apart from the single inversion required, which is achieved through the use of via placement, (Figure 3.14.) The symmetry is desirable to ensure phase uniformity. A ring of substrate contacts surrounds the layout to help insulate the VCO from any substrate noise generated from other circuits. This layout is approximately $191\mu m\ x\ 191\mu m$. As with the previous design, rings of metal distribute power evenly.

3.2.3 Performance

The lower curve in Figure 3.12, represents the post-layout parasitic simulation results. It shows a peak at the same control voltage, but reduced to around 6.4GHz, a decrease of 7.3%. Typically post layout simulations seem to average around 10% decrease in performance. The smaller decrease can mostly be attributed to the extreme tightness of the layout.

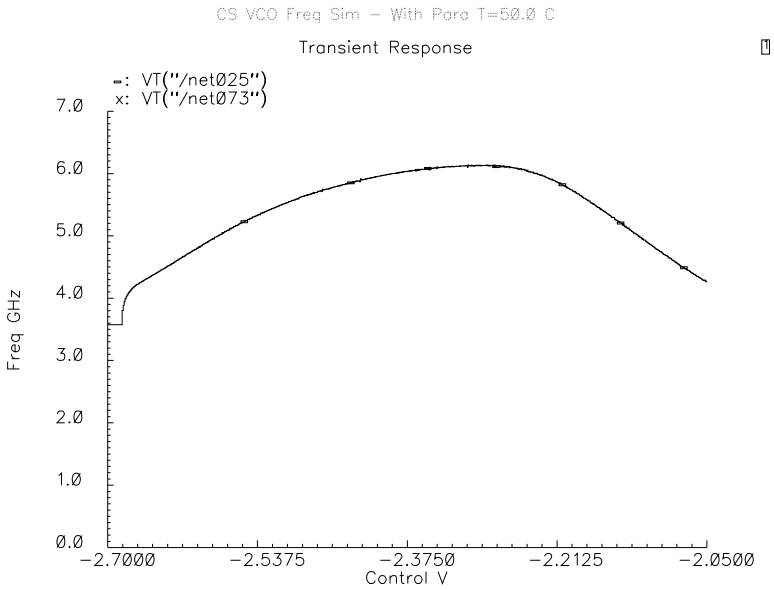

Figure 3.15: 5AM VCO Characteristic with parasitic extractions at 50C.

The plot in Figure 3.15 shows the effect of temperature on the VCO. Previous simulations used the nominal temperature of 25C, while the above was performed at 50C. As a result, the peak frequency dropped to 6.1GHz, and it moved to a control voltage of -2.3V. The gain is sightly decreased. Other factors affect the frequency of operation such as the power supply voltage. In general, varying the supply voltage will move the peak control voltage by the same amount, as the current mirror's input is relative to V_{EE}. This would keep the bias currents the same. Increasing the supply voltage causes the output swings to increase in amplitude, and the C_{BC} junction capacitances to decrease slightly, partially lessening the loading of the buffers. However, the larger voltage swing causes an increase in the base charge storage at each buffer input. This more than offsets the decrease in capacitance at the top of the tree, so the peak frequency tends to slightly decrease as the supply voltage is increased, with the value of $V_{Ctrl} - V_{EE}$ being held constant. This was verified in a simulation using the peak V_{Ctrl} for 3.6V, and the frequency only varied by 200MHz as the supply was changed by $\pm 10\%$.

3.3 Second Revision 5AM Oscillator

A new VCO was designed for the 5AM process which made use of fixed edge interpolation, current starving coarse adjustments, as well as varactor based control. Because this design was meant to be instanced for several different operating ranges, the central tuning components were included in a separate schematic and layout, and adjusted to obtain different frequency ranges. The architecture was identical to that of the 5DM oscillator, shown in Figure 3.1.

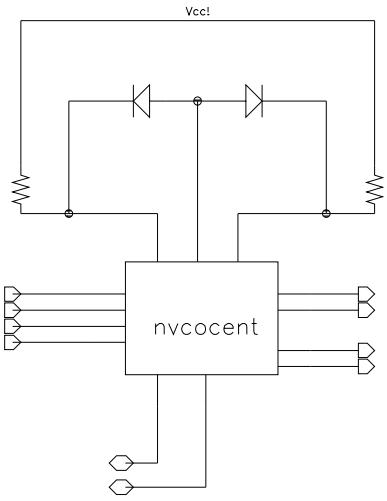

Figure 3.16: 5AM 10GHz VCO Tuning Network. The diodes shown are varactors.

The two VCOs were designed to allow transmitter operation in the 30Gbps and 40Gbps ranges, and they are referred to as the 7.5GHz and 10GHz VCOs respectively in this section. The slower VCO used larger varactors with an area of $210\mu m^2$. The faster one used $48\mu m^2$ varactors. Note that each buffer requires two varactors, and there are four buffers in the VCO. The combined area due to varactors is significant in terms of total layout size. According to [40], large voltage swings across varactors can have a detrimental effect on output signal quality, due to the non-linear capacitance having different values at different points on the output of the waveform. However, as the phase noise of ring VCOs is already quite poor, this effect will be negligible. The resistors used for the buffers were 102Ω.

The main buffer circuit is pictured in Figure 3.17. In this circuit, the three central terminals connect to the tuning network as shown in Figure 3.16. Two

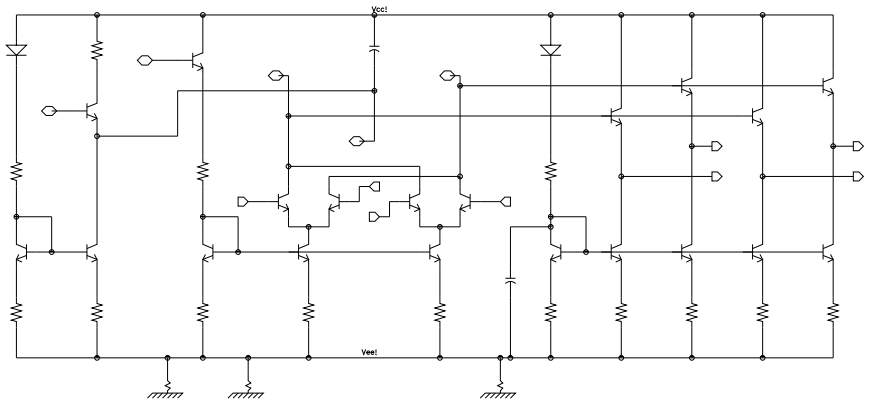

Figure 3.17: 5AM 10GHz dual control buffer. Four of these make up the ring oscillator.

control signals enter on the left, with the leftmost being the fine control. The fine control uses what is essentially an emitter follower to set the voltage at the varactor node. The coarse control affects the half mirror which provides a reference for the interpolating buffers at the center. The outputs are via the emitter followers on the right. The first set provides the inputs to the subsequent stages in the ring. The second, the rightmost, supplies the output signal to the driving buffers for the rest of the chip. In order to supply a strong enough signal, the VCOs feature very large output drivers with $16\mu m$ emitter followers consuming 11mA each.

3.3.1 Layout

Layout of the new VCO was carefully done to minimize size, and thus minimize parasitics. The symmetry of the design extends up to the last metal layer where the VCO lies centered on a power rail. Lower layer rings distribute power and control signals evenly to the four quadrants. Emitter degenerating resistors were not used in this design as the high speed operation precluded a normal linear transfer characteristic and the resistors served to increase the capacitance at the base of the pair and degraded the operation of the biasing current source. This layout, which included the large drivers necessary for adequate delivery to other circuits of the serial transmitter, is quite large. The total area is $325\mu m\ x\ 325\mu m$.

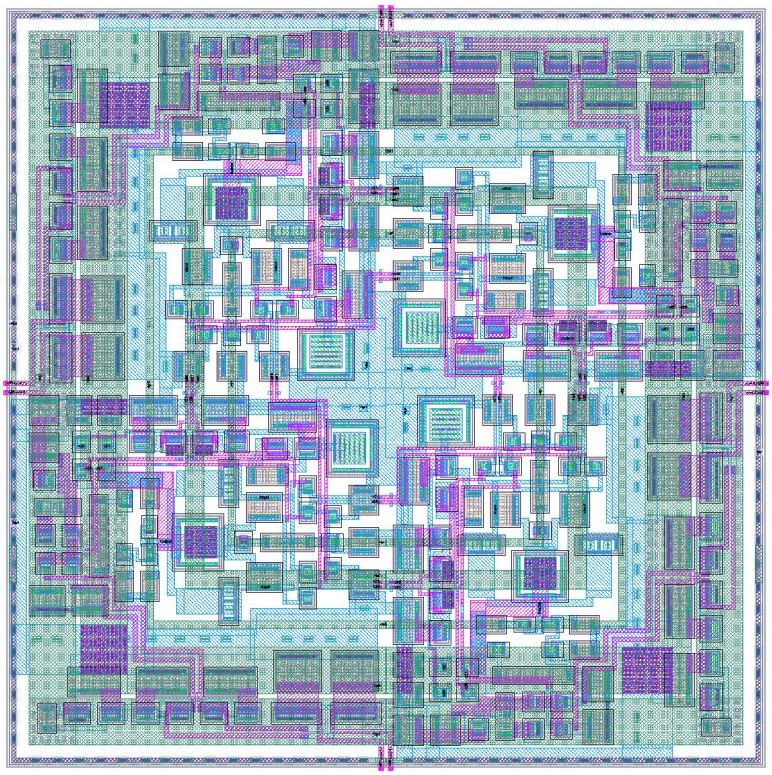

Figure 3.18: 5AM 10GHz dual control VCO Layout. The top layer of metal is not shown.

3.3.2 Performance

Spectre was used to extract the control voltage .vs. frequency characteristics of the two VCO versions under loaded conditions, both with and without parasitic extractions from the layout.

Post-layout simulations of the two versions of the dual control 5AM VCO showed a larger decrease in performance over the simulations without parasitics than did previous VCO designs, indicating that the parasitic effects are more critical at higher frequencies in determining performance. The 10GHz VCO post-layout control plot is shown in Figure 3.19. The highest frequency achieved is around 8.5GHz, while the peak in the simulations without parasitics is 11.2GHz. Likewise, the plot of Figure 3.20 shows a peak of 5.8GHz, while the value without parasitics is 6.9GHz. This is a 25% and a 16% reduction, respectively.

SpectreRF was used to obtain phase noise measurements for these designs. The 10GHz VCO showing -82dBc/Hz at a 1MHz offset, while the 7.5GHz produced

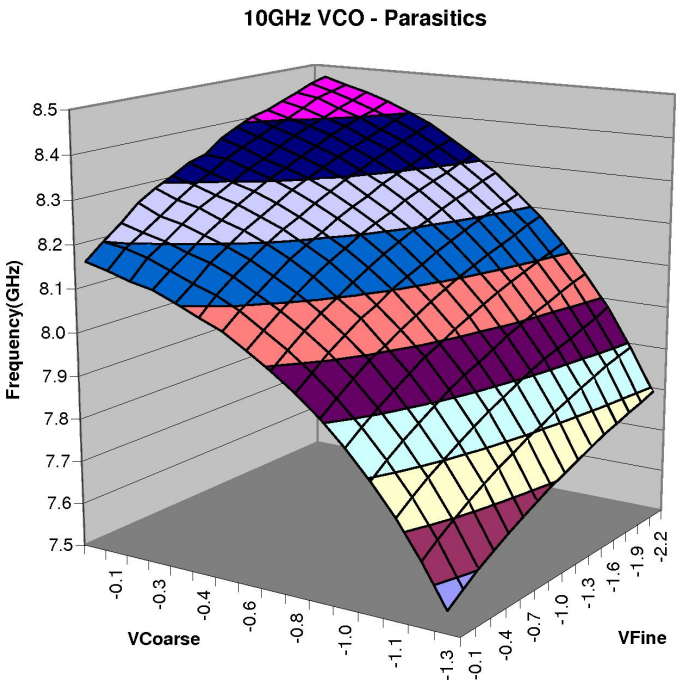

Figure 3.19: 10GHz 5AM VCO control characteristic with parasitic extractions, and full loading.

-83dBc/Hz, (at the same offset).

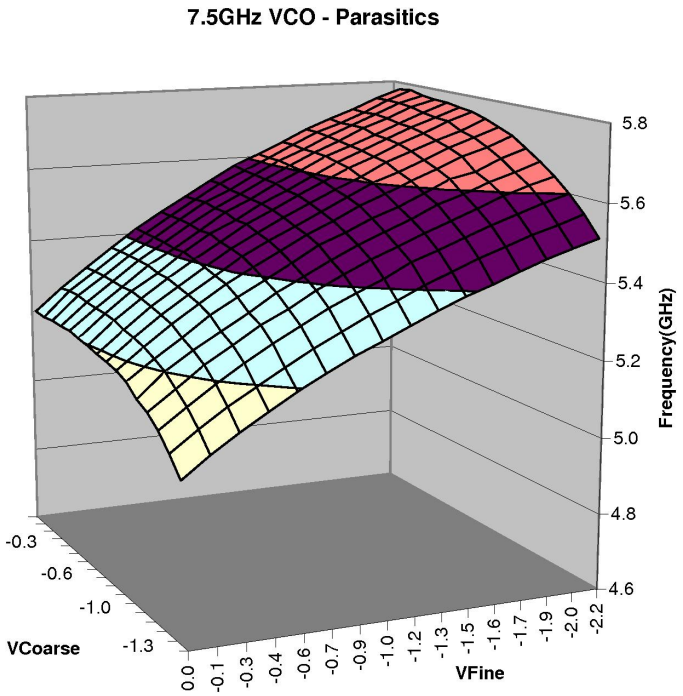

Figure 3.20: 7.5GHz 5AM VCO control characteristic with parasitic extractions, and full loading.

3.4 Conclusions

Three separate VCO designs have been presented in this Chapter. Two of these designs, (the 5DM & SIII 5AM VCOs), have been fabricated, while the third has been submitted in two different forms, (the 5AM 10GHz & 5AM 7.5GHz VCOs). These ring VCO designs make use of both bias current frequency control, and varactor control. Simulations were performed both for the schematic level designs and for the designs including resistive and capacitive parasitics extracted from the circuit layouts. In all cases the simulations were performed with the VCO under the expected load conditions, so the simulations should produce the most accurate results. The effects of the parasitics on output frequency increased as the target frequency was increased. In addition, the effect of temperature and power supply variation on the output was considered. These VCOs were designed with a primary emphasis on speed, driving capability, and range, rather than for low noise.

Table 3.1: VCO Characteristics

Design	Min Freq (GHz)	Max Freq (GHz)	Gain (MHz/V)	Phase Noise (dBc/Hz)	Power (mW)
5DM	8.1	9.2	210	-84	520
SIII-5AM	4.2	6.4	3900	-89	340
5AM10GHz	7.6	8.5	110	-82	820
5AM7.5GHz	5.0	5.8	225	-83	820

As can be seen in Table 3.1, the phase noise of the designs increase monotonically with the target frequency. This is because the desired wide frequency variability goal is in direct conflict with the desire for low noise. LC-tank VCO circuits provide far superior noise characteristics over relaxation oscillators, but they do not have as much variability. As the target bit rates increase for future designs, the VCO phase noise characteristics need to **improve** to maintain the same eye quality. For this reason, it is recommended that LC-ring VCOs should be utilized in future monolithic designs.

CHAPTER 4
Edge Steering Multiplexer Architecture

The primary metrics of a good serial transmitter output are embodied in the so-called "eye diagram". A large open height relative to the noise magnitude is necessary for useful communication. The width of the eye relative to the bit time must also be large in order to extract an accurate clock so that the eye can be sampled in its open region, rather than at the corners. (See Appendix C which covers noise and eye diagrams.)

Most commercial designs do not attempt to generate a serial data stream which is as fast as ours relative to the fundamental limits of the technology. Most designs have a latch at the output of the transmitter which *retimes* the output ensuring that the eye width is as large as possible. In the absence of a full rate clock, a half rate clock can be used to switch between two half rate latches to fully retime the output signal.

Seeking to pursue data rates beyond those available at twice our maximum on-chip clocking rate, we have explored quarter-rate clock architectures which do not use a half-rate 2-to-1 multiplexer or full rate latch to generate the final output signal.

In our previous designs, we have been fairly successful at creating 20Gbps in IBM's 5HP technologies. We have used a variety of novel architectures to produce fairly clean outputs using quarter-rate clocks. Eye width has been difficult to achieve at higher data rates.

The basis of this thesis is a new architecture for quarter-rate retiming. This architecture is superior to the previous designs which our research has produced. It produces better eyes at the data rate of the last generation of SERDES, plus it is capable of producing eyes of the same level of quality at higher data rates.

4.1 SERDES III Symmux

The final output architecture of the 20Gbps SERDES III transmitter utilized a new 2-to-1 multiplexer architecture developed by fellow researchers here at Rensselaer Polytechnic Institute for which they received a patent [31]. Although this was a great improvement over previous designs, there are still problems with the design which make it less suitable at much higher clock rates.

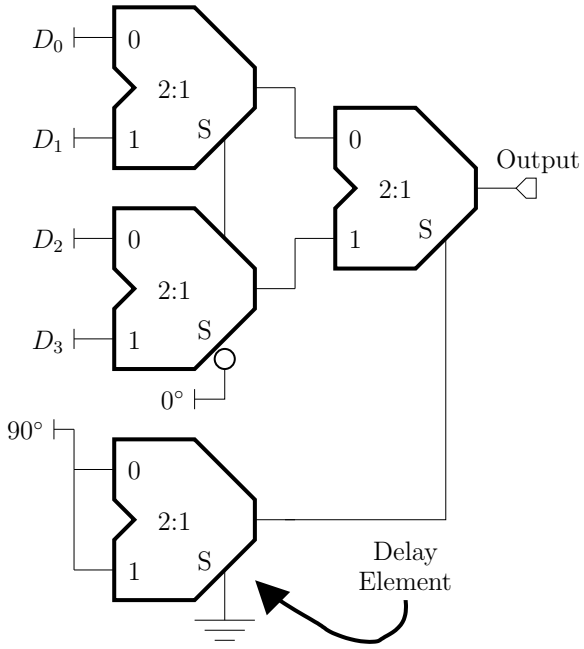

Figure 4.1: Four to one architecture used in our previous quarter clock rate designs composed of four 2-to-1 multiplexers.

As can be seen in Figure 4.1, the previous two designs realized the 4-to-1 multiplexer function using three 2-to-1 multiplexers. An additional 2-to-1 multiplexer identical to the others was used as a delay element to offset the 90° phase by the same amount of time as the input data before reaching the final 2-to-1 output multiplexer. The SERDES I design used standard CML 2-to-1 multiplexers and suffered from duty cycle generation due to the non-identical propagation delays for signals present at the data and select inputs. The 90° degree propagation delay from the data input to the output of the delay element was different from that of the 0° propagation delay from the select input to the output. This caused the signals to arrive at the final 2-to-1 slightly out of synchronization. In addition, the final 2-to-1

in the SERDES I design also introduced non-uniform delays.

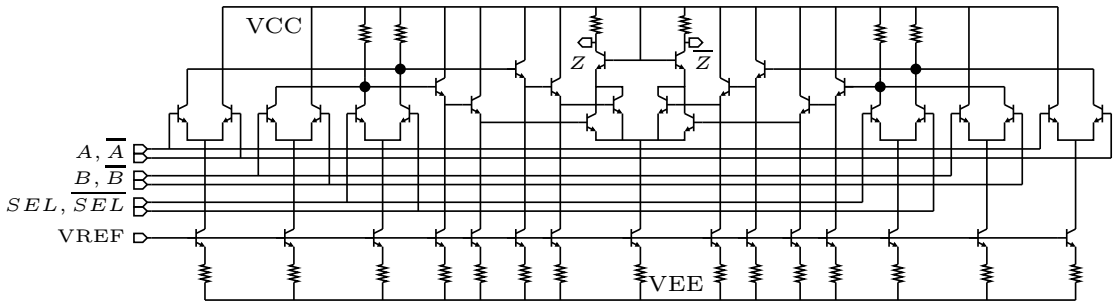

Figure 4.2: SERDES III symmetric 2-to-1 multiplexer.

The SERDES III design used the new 2-to-1 symmux multiplexers shown in Figure 4.2, duplicated from Figure 2.16 for convenience. These new multiplexers were designed to have uniform propagation delays from either type of input to the output. The right side generates A and B signals at the dotted nodes, while the left side generates their complements. The circuit operates by using the select, "SEL", pairs on either side of the circuit to disable, (bring low), either the A or B signal lines. Only one data input signal, (A or B), or its complement, ($\overline{A}$ or $\overline{B}$), will reach the central current tree with a "high" signal, thus steering the current in the correct direction. The other lines reaching the central tree will be in one of two "low" states. To ensure full current switching, the resistors above the dotted nodes should be sized to give twice the normal CML voltage swing so that under all steady-state conditions the current in the central tree is completely steered to one side or the other.

The three "first tier" 2-to-1 multiplexers of the SERDES III design as depicted on the left side in Figure 4.1 have smaller transistor sizes and correspondingly smaller currents than that of the final 2-to-1 multiplexer. This is due to the fact that the final multiplexer is used to directly drive a 50Ω output line. It has no resistors at the outputs of the cascode stage, relying on the characteristic impedance of the line to supply the voltage drop.

The first tier multiplexers have $1\mu m$ input transistors, $2\mu m$ and $4\mu m$ emitter follower (EF) stages, and an $8\mu m$ central tree. For the final 2-to-1, the input transistors are $4\mu m$, the EF stages are $6\mu m$ and $8\mu m$, and the central tree is $12\mu m$. All

circuits are biased at 0.7mA/μm, so the central tree swing is nominally 420mV.

The symmux circuit approximates uniform delays, but beyond 20Gbps, asymmetries in the circuit become apparent in the output. While a 5 picosecond difference might be acceptable in a U.I. of 100ps, at a U.I. of 25ps it becomes intolerable. While the circuit is definitely symmetric in the steady-state sense, during an input transition different numbers of transistors can change state, (and by different amounts), depending on which input is changed. This has a direct impact on the rise and fall time of the outputs, which affects the overall propagation delay.

The primary cause of asymmetry can be seen by examining the states of the transistors in the central tree during different types of transitions which are propagated to the output, (either SELECT, or DATA). During a DATA transition, two transistors change state, with one changing from "half on" to "fully on" while the other does the opposite. In a SEL transition, all four transistors change state, slightly delaying the output transition.

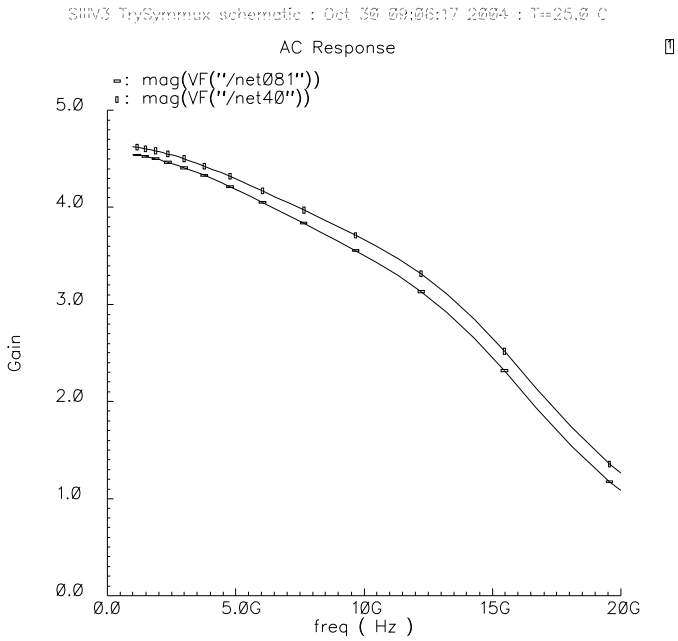

Figure 4.3: Midpoint gain of top tier symmux for different types of transitions. DATA transitions (upper) have higher gain while SELECT (lower) are slower.

This lack of transition symmetry can be seen in Figure 4.3 which indicates the frequency dependent gains at the midpoint of transitions. This causes a measurable

difference in the slew rate of symmux circuits. For example -17GV/s for a downward DATA driven transition at a top tier output as opposed to -15GV/s for a SELECT driven one.

Also of concern is that a given output state is realized by different non-equivalent transistor states depending on the inputs. At all constant output conditions, the symmux central tree has one transistor in the strongly-on state, one transistor in the strongly-off state, and two transistors in the half-way state. For some logic "low" output states the strongly-on transistor is on the same side of the current tree as the strongly-off one, (when $A = B$). In other cases the strongly-off and on signals are on opposite sides, ($A \neq B$). This has the effect of creating small differences in the output levels for both high and low states, which causes data dependent jitter due to the effective pre-advancement or retardation of the signal towards its final value for a transition. The slower the edge transition, the greater the effective delay.

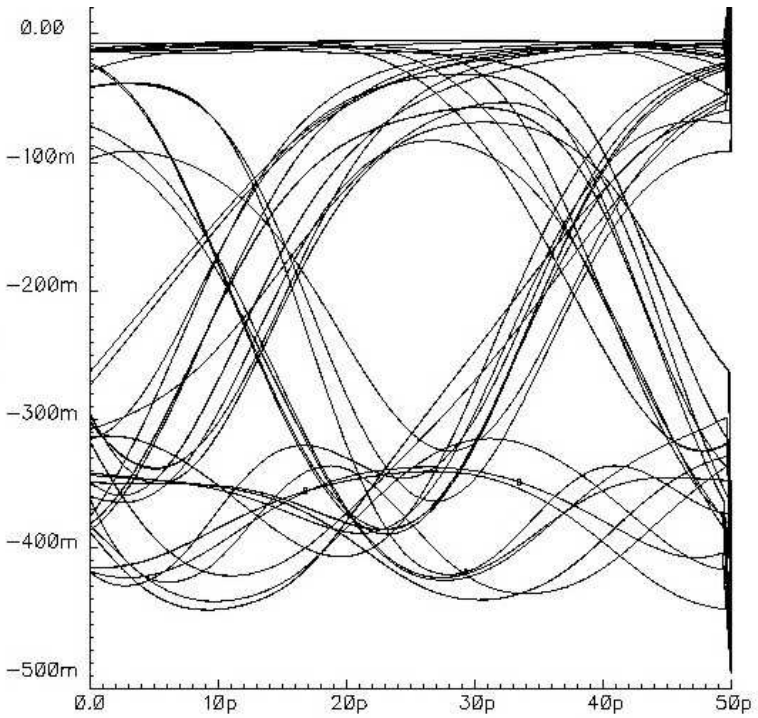

Figure 4.4: Idealized 40Gbps output using SERDES III symmux architecture.

Each symmux used in the final architecture contributes to the overall jitter and

non-idealities in the final serial transmitter output. Even when used with **ideal** data and clocking, the SERDES II architecture produces an unacceptable eye diagram at data rates of 40Gbps, as shown in Figure 4.4. This eye plot does not take into account any VCO noise, phase mis-match, or any parasitic loading.

4.2 Edge Steering Multiplexer

Thinking about the symmetrical nature of a four phase VCO in conjunction with the retiming problem lead to the conception of the new Edge Steering Multiplexer Architecture, (ESMA). It was attempted to envision a symmetrical circuit surrounding the VCO which would generate the output data. It was thought that perhaps some sort of circular shifting register could be created which would simultaneously serve the purposes of a VCO and a multiplexer.

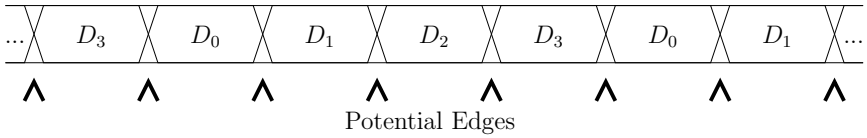

Figure 4.5: Potential edges in the serial output stream would occur between adjacent bits.

It was while examining a list of potential edges in the output serial stream, (and the corresponding clock phases that would be used to generate them), that the notion of the edge-steering architecture was hit upon. Figure 4.5 shows these potential edges that can occur between bit pairs. It was decided to investigate the possibility of a circuit which would produce just those edges.

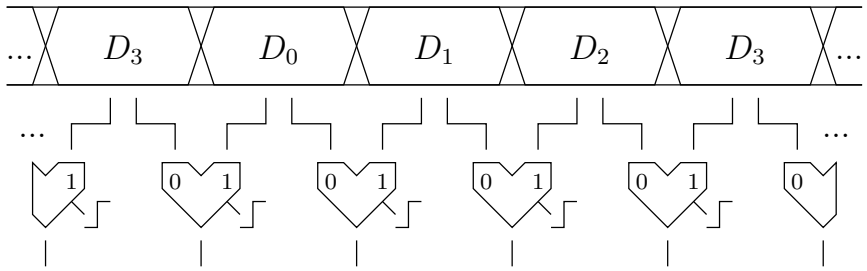

Figure 4.6: Edges in the serial output stream generated by 2-to-1 multiplexers clocked by phases.

As can be seen in Figure 4.6, a 2-to-1 multiplexer, switched at the right time, can generate the required output edges when they are needed. Our multiphase clocks are well suited for generating the clocks for these multiplexers.

For this to work, each input bit needs to be sent to two multiplexers so that the edges on either side can be generated. After the edges which are required for

output are generated, they must be combined to produce the final output. A simple wired-OR multiplexer can not adequately combine the edges as the high loading at the output nodes preclude high speed operation. Instead, a multiplexer tree of discrete 2-to-1 multiplexers is used.

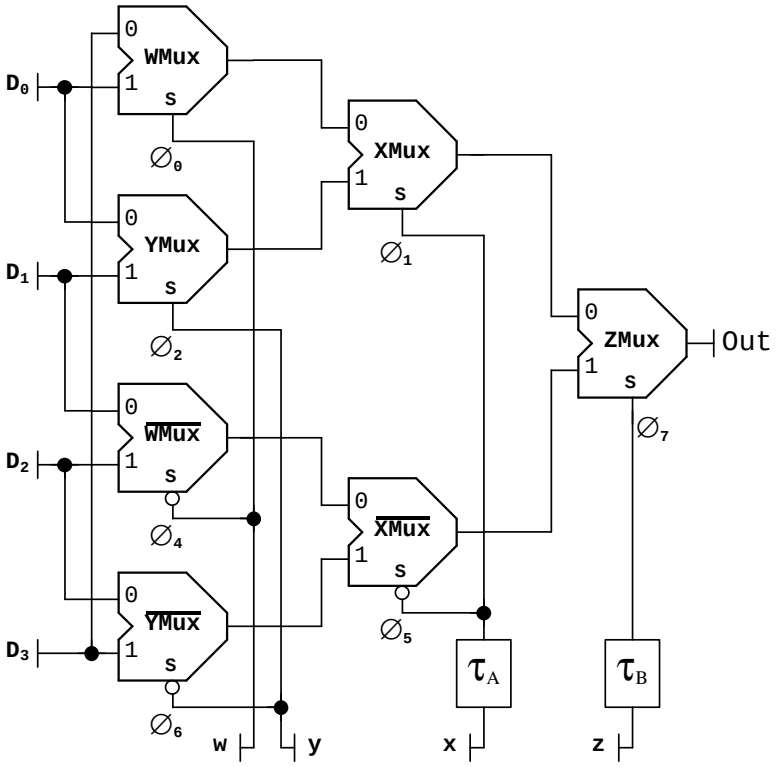

Figure 4.7: Proposed new multiplexer architecture.

This edge-steering multiplexer creates all edges that will appear in the output serial stream at the first tier. These edges then propagate with a uniform delay through the rest of the circuit resulting in an output stream with no inherent duty cycle distortion. The multiplexers after the first tier "steer" the edges to the final output while introducing no new edges into the output stream themselves. For example, the **XMux** multiplexer switches between passing the outputs of **WMux** and **YMux** while they both select D_0. Because of this no new edge is introduced.

The multiplexer architecture shown in Figure 4.7 has several improvements over previous generations. In the first SERDES chip, the propagation delay was addressed to only the first order. High bandwidth multiplexers were used, but the output possessed a duty cycle. The SERDES II chip improved on that using the

symmux architecture, but suffered from poor bandwidth and lack of complete symmetry. The new "edge steering" multiplexer presented here uses the high bandwidth characteristics of standard CML multiplexers while addressing the symmetry issue. Its operation can be understood by referring to Figure 4.8 and Figure 4.9 which show timing diagrams of the clock phases and the data at the multiplexer outputs respectively.

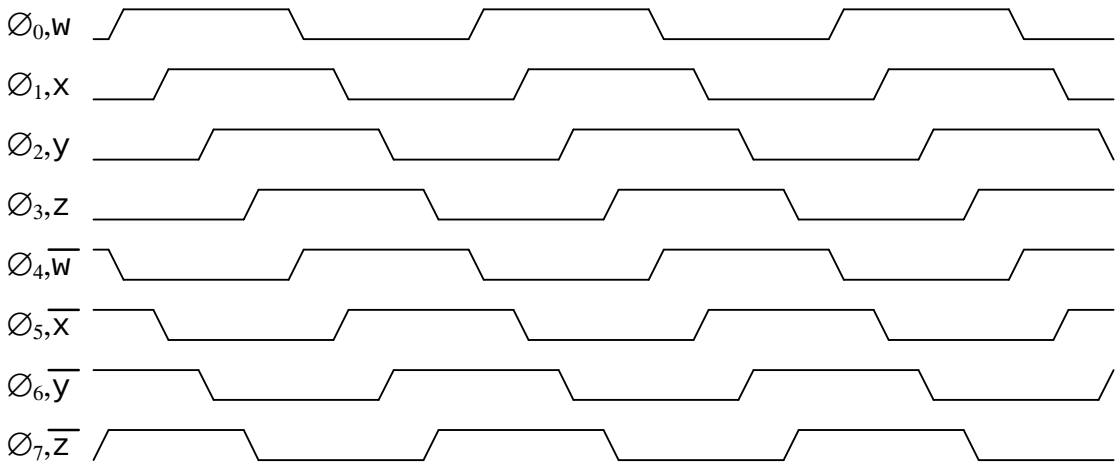

Figure 4.8: Eight phase clocks as generated by a differential four buffer ring VCO. Both phase numbers and letter names are used to refer to specific phases.

Using a differential four buffer VCO, we obtain eight effective clock phases, as diagrammed in Figure 4.8. As can be seen, each phase has a matching complementary phase and we refer to them at different times by their letter or number designations.

The multiplexer operates by using the top(leftmost) tier of 2-to-1 multiplexers to generate the edges that will be present in the final output. This topmost tier is driven by the "w" and "y" phases and their complements. For example, a potential edge exists between data bits zero and one, (D_0 and D_1 in Figure 4.7). Therefore, both these bits are present at the input to a multiplexer, in this case the one labeled "YMux", which is driven by the "y" phase. On the rising edge of the "y" phase, if the values of bits D_0 and D_1 differ, there will be a rising or falling edge at the output of the "YMux" multiplexer. Similar edges are generated for each pair of potential edges in the final output sequence. These edges are indicated in Figure 4.9 by the

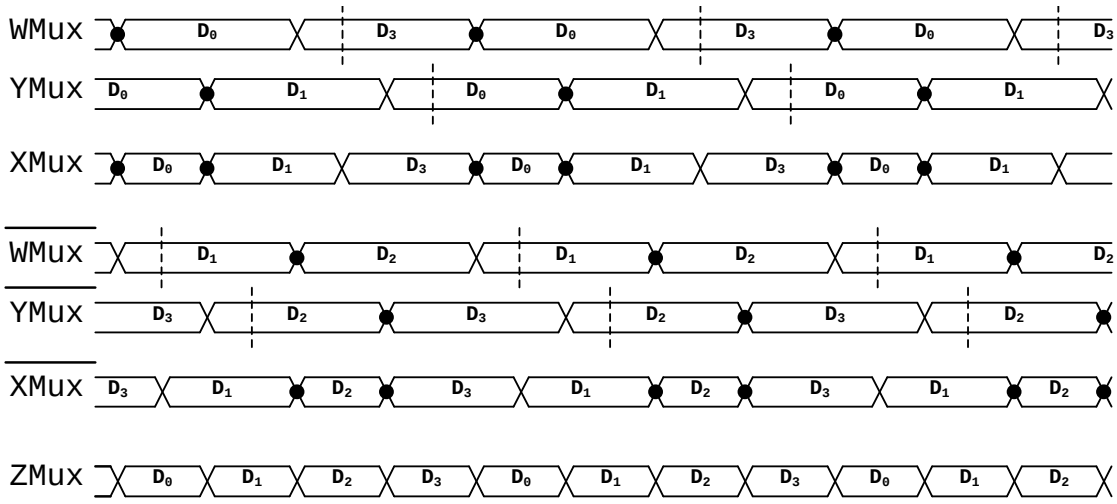

Figure 4.9: Steered edge architecture multiplexer outputs. Each signal represents the output of a 2-to-1 CML multiplexer, and is labeled on the left with the phase line used to drive the multiplexer select.

presence of a dot.

Once all the critical edges are generated by the top tier, the rest of the multiplexers are used to "steer" these edges to the output. In the second tier, the multiplexers are driven by the "x" phase and its compliment. When the rising edge of the "x" phase arrives at the "XMux" multiplexer, it switches between identical D_0 values. Note that the "x" phase should still be delayed by the same amount of time as the propagation of signals through the tier one multiplexers. However, unlike the previous multiplexer schemes this delay isn't critical. The scheme will still work as long as the rising edge arrives at some time between the edges framing the D_0 bit, and thus the rising edge will not introduce any edges into the next tier. The falling edge of "x" can introduce new edges, but these will not propagate to the final output.

At the final tier, the "ZMux" is used to combine the outputs of the second tier multiplexers. At all times when the select signal is changing, the data at the inputs of the "ZMux" are identical. Therefore, no new edges are introduced. In essence all the dotted potential edges from the first tier appear in the final output, and nothing else.

This new architecture does have the drawback of requiring more phases than

the previous schemes, however, the quadrature requirements of the previous architectures and the characteristics of the 4 buffer ring VCO make these available for free. In addition, it should be possible to implement this scheme using just quadrature phases, as the timing for the latter stage multiplexers isn't as critical and may be derived via delays from the existing phases, for a nearly fixed frequency implementation.

Another drawback to this scheme is that generated edges have to pass through 2 levels of multiplexers instead of one. While it is true that this will have negative implications for the signal bandwidth as compared to the simple multiplexer in the original SERDES chip, it is necessary as it alleviates the duty cycle problem to the greatest extent possible. When compared to the SERDES II multiplexer, it has much better bandwidth and subsequent response.

Each data bit is sent to two of the top-tier multiplexers. The ideal time to update that bit is during a time when neither multiplexer is selecting it. The dotted vertical lines in Figure 4.9 indicate the best times to change input bits. As can be seen, there is no real overlap so the data inputs must be each updated in a "round-robin" fashion, with each bit update staggered by one quarter of the main VCO period. The difficulties in producing data for testing this architecture are addressed in the next two chapters.

CHAPTER 5
Multiphase LFSR

The creation of the new Edge Steering Multiplexer Architecture(ESMA) led to new requirements regarding the generation of input data for testing. The new ESMA, (similar to previous designs), multiplexes four serial bit streams into a single stream at four times the data rate. However, this new architecture differs in that each bit must be offset at one quarter the clock period. At the intended clock rate of 10GHz to yield a 40Gbps output stream, the time between each bit change would have to be as little as 25ps. This presents significant design challenges as many of our previous circuits have barely worked at 5GHz, much less at 10GHz.

The combination of a linear feedback shift register(LFSR) and four 4-bit shift registers to implement the 16-to-4 multiplexer functionality would likely not be possible at those rates. Each shift register grouping of four latches would need an update with fresh data that would have to arrive inside a very tiny 25ps window, and ensuring that this occurred while shifting at a rate nearly as fast as the fastest possible clock divider did not appear at all feasible.

5.1 Multiphase Design

It was realized that the four 4-bit shift registers of the SERDES-II design *themselves* were of a size and complexity approximating that of the LFSR itself, (see Figure 2.25). Examining the 12-bit LFSR used in the SERDES II design, the clock buffering already leads one to look at the circuit as groups of individual registers rather than as a single whole. Given this, I designed a circuit that would incorporate the functionality of both the 12-bit LFSR as well as that of the 16-to-4 multiplexer.

This new test generator, (Figure 5.1), essentially would generate four 10Gbps streams at the correct offsets, using four clock phases. To do this, the original LFSR was converted into a multiphase state machine whereby groups of three latches at a time are advanced into the next state by a specific clock phase. Unfortunately, the

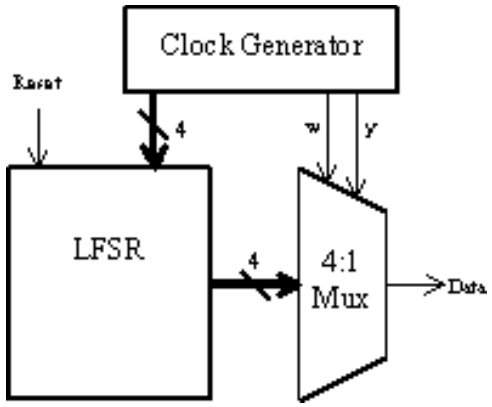

Figure 5.1: Simplified test generation architecture and output multiplexer.

design needed other modifications in order for this to work properly.

In a simple state machine, the next state is computed via combinatorial logic and latched on the next clock cycle. For this new LFSR design, each group of three latches have a next state depending on the values of latches which would have been changed by previously occurring phases. Various adjustments were made to the numbering of the latches to ensure minimal dependencies. It was found that four of the bit values needed to be buffered, so as to allow them to enjoy the same set-up margins as the other latches.

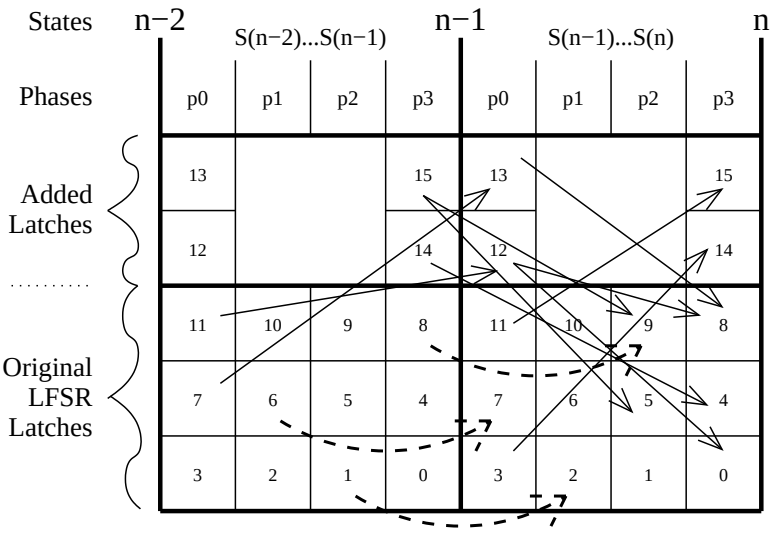

Figure 5.2: State dependency diagram for transmitter multi-phase LFSR. The dashed arrows represent examples of normal shifting transitions, not all of which are shown.

This is illustrated in Figure 5.2. The lower three rows represent the latches present in the original design. Each clocking phase is represented by a column, and two complete cycles are shown. The upper two rows represent added latches 12..15 which are required to buffer the values to prevent their early loss. Since the majority of the latches capture data from three phases prior to their own clock, this would give a minimum 75ps setup time at the target data rate. This design improvement results in a total of 16 master/slave latches, the same number as was present in the SERDES-II shift register design alone, yet it eliminates the need for loading from external sources.

The SERDES II LFSR needed to operate at one sixteenth the output data rate, and the 16-to-4 shift registers had to operate at 1/4th the output data rate. In this new design, the LFSR has the functionality of both and runs at 1/4th the output rate, however since the output rate can be as much as doubled, this LFSR is required to run at 8 times the speed of the SERDES II LFSR, and at twice the speed of the SERDES-II shift registers. Several other changes were need to make the circuit operate at those rates.

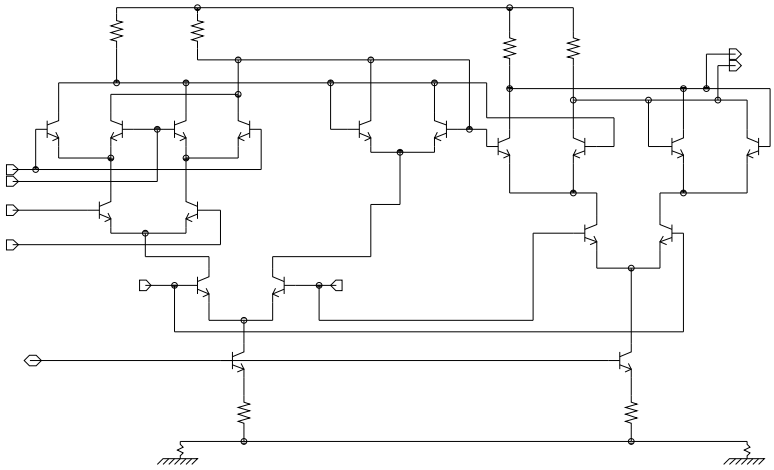

Figure 5.3: Combined XOR and MS latch schematic.

For this LFSR design, the exclusive-or function was incorporated into the master latch tree which results in significant savings in power, area, and an increase in maximum clock rate. This is shown in Figure 5.3. This could not be done in the SERDES II design because of the need for reset circuitry. The combination of reset

and XOR would have required additional voltage levels.

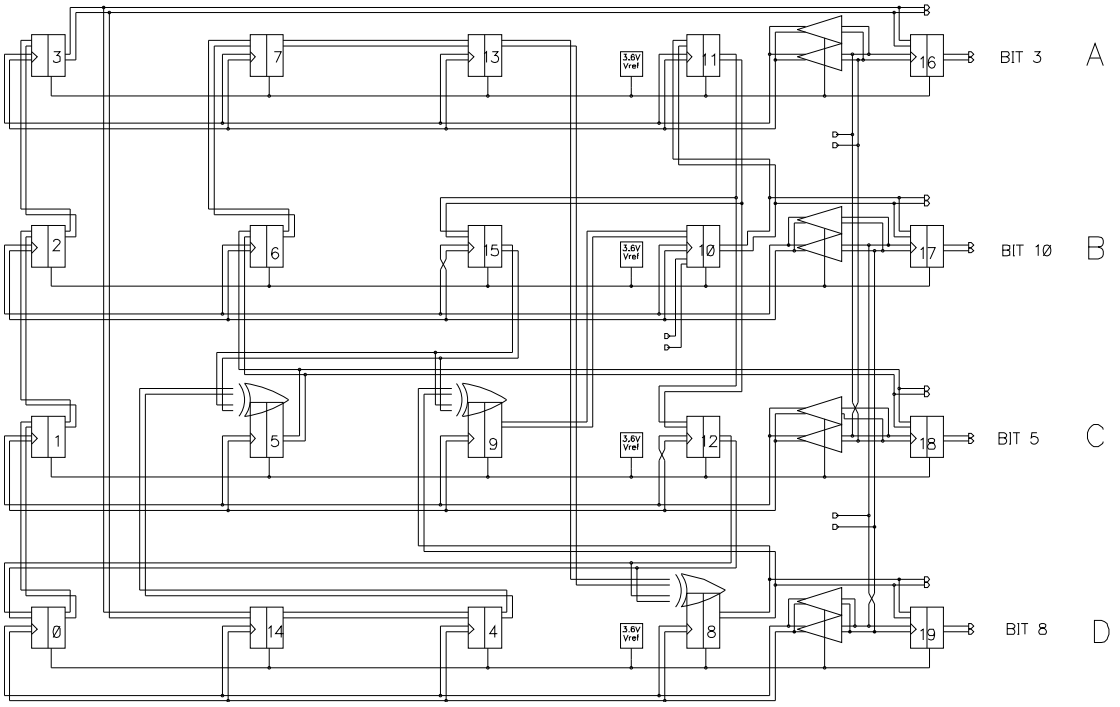

Figure 5.4: Multiphase LFSR schematic diagram.

In Figure 5.4, the full structure of the multiphase LFSR can be seen. The rows correspond to a particular phase. Heavy buffering was required to allow the latches to operate quickly. To ensure a small delay between input clocking signals and output data change, the output data bits were fed to a final set of latches on the far right, which are fed directly from the input clocks.

The SERDES-II LFSR design included reset functionality by incorporating a reset input into each D-Latch. The "sense" of one of the latches was altered by complementing both its inputs and outputs, thereby effectively changing its reset condition to its set condition. This would ensure that all latches would be set to a zero state, apart from the one with its sense altered, which would be set to one.

Using a multiphase clocking scheme for the new LFSR made such a simple reset scheme unrealizable. Luckily, the reset does not have to be synchronized at the transmitter, so instead a reset method was developed that used multiple clock cycles to force the LFSR into a known state. It was found that latch number 10, being present in all the feedback loops, could force all the latches to a known state

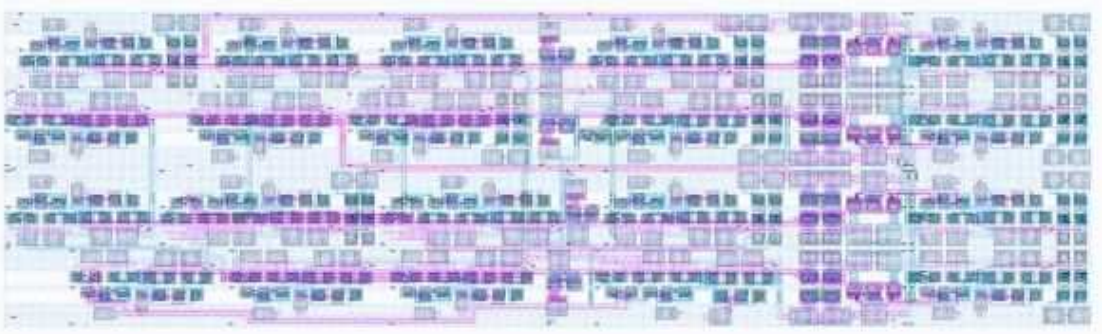

Figure 5.5: Layout of the test generation circuitry.

by forcing its output to one and holding that value for twelve or more of the 10GHz clock cycles.

5.2 Test Clock

In order to be able to capture a standing pattern from the transmitter, a clock must be generated which is synchronized with the LFSR cycle. Normally, a state detection AND gate would need a fan-in of 12 to accomplish this, and this would be difficult to achieve in a 3 level CML logic. Because of this, in the SERDES-II design a single-ended wired- or type of gate was used. At effectively eight times the rate of the previous design, the single ended approach did not operate, and given the multiple phase nature of the LFSR, a single gate solution simply could not be used.

Initially, it was thought that a specific pattern which occurs only once per 4095 bits could be detected as it occurred, and a signal could be generated from that. Since one such pattern consists of a single sequence of 12 consecutive ones, a "ones-counter" was developed which would reset each time a zero was encountered

and send out a pulse on encountering the twelfth one. Unfortunately, this circuit failed to operate correctly approaching 10GHz. The feedback paths in the 4 bit state machine had too many gate delays and loading to operate at that rate. The LFSR has essentially zero propagation delay as all feedback gates are built into the same current tree as the latch.

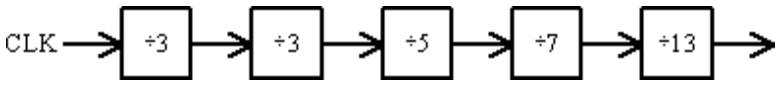

Figure 5.6: Divide-by-4095 block diagram.

As a solution, The idea was conceived that since the pattern length was 4095 bits, and this factorable as $4095{=}3^2 \times 5 \times 7 \times 13$, a simple divider could be used on a clock phase to generate the required 2.44MHz clock. Luckily, the number 4095 is divisible by small factors, such as 3, which allows an initial fast divider to bring the frequency down to an acceptable level so that the latter division stages could have longer feedback propagation delays.

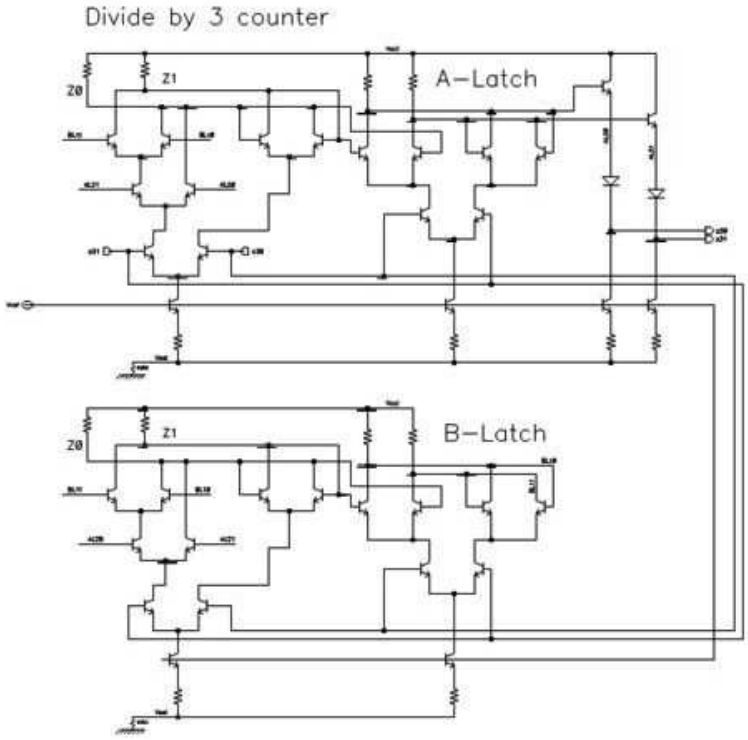

Figure 5.7: The divide-by-three counter circuit schematic.

The divide-by-three circuit merits attention as it must operate at the highest rate possible. Shown in Figure 5.7 is the schematic for this circuit. It uses the absolute smallest possible feedback path to divide at the highest rate. The logic in the feedback path is incorporated into the master latches. Nominal simulations with no parasitics showed this circuit operating at up to 17GHz. The fast divide-by-three circuit can thus operate faster than the LFSR, so it should not introduce limitations to the testing process.

5.3 Problems

In the course of simulating the LFSR test generation circuits, it was found that in simulations with no parasitics, the correct bits were produced at up to 15Gbps. When post-layout parasitics were added, however, the usable range was reduced to approximately 10Gbps. Further investigation into the limiting factors showed that prior simulations even at 10Gbps were unrealistically optimistic, and that those numbers represented an unachievable upper bound on performance. The major reason for this is that the multiphase clocking reduces the available setup and hold times from 100ps to 75ps, which is equivalent to clocking at 13.3GHz. Secondly, the input clock loading reduces the rate at which the latches can operate. Whereas the SERDES II LFSR operated at 1.25GHz, this circuit needed to operate at effectively 13.3GHz. As 10GHz is nearly the limit of dividers operating in this technology, it was decided to investigate other test pattern generation circuits.

CHAPTER 6
LFSR

Generating data to test the Edge Steering Multiplexer Architecture proved to be a non-trivial design task. To fully test the highest speed parts of the design, the lower speed circuits should not be the bottleneck. In the previous chapter, one proposed method was presented and the reasons for rejecting it explained. It was decided to handle the generation of test data independently from the 16-to-4 multiplexer functionality in order to find the best solution for each task. The requirements of the LFSR were thus altered to producing 16 parallel bits of data with a single clock at one sixteenth the target output rate, or 2.5GHz. This is still quite fast for a complex state machine in this technology.

6.1 Bit Pattern

There are many industry standards associated with serializer design, and several of these involve standard testing data patterns. It was hoped that by creating a serial data stream conforming to one of these standards, existing test equipment could be used to evaluate this design. A widely supported standard for SONET/SDH networking consists of a $2^{31} - 1$ length sequence realized by a specific LFSR characteristic equation, (See Appendix B on LFSRs.)

$$G(Z) = Z^{31} + Z^{28} + 1 \tag{6.1}$$

This **{31,28}** equation represents the serial stream which must be present at the *output* of the serializer. Thus, either a Galois or Fibonocci form based on this equation will produce the desired serial stream at the output of every register.

This is not useful as the entire LFSR, (Figure 6.1), would need to be run at the target output frequency and would act as the transmitter itself. Instead, what is required is a state machine that produces the required sequence distributed across 16 separate outputs. No simple combination of multiple 31-bit LFSRs of this sequence can realize this functionality as on a given bit line, the output is required

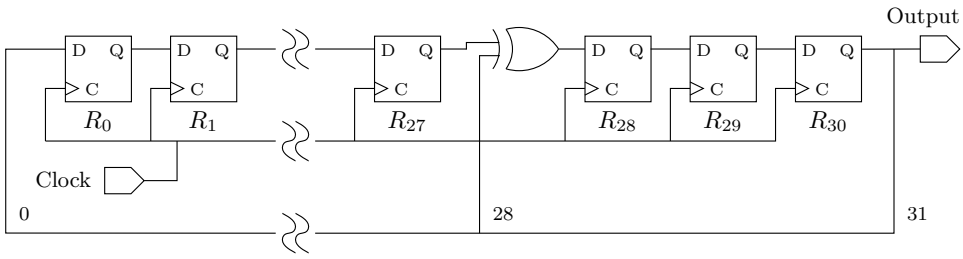

Figure 6.1: Schematic of Galois 2^{31}-1 LFSR circuit.

to be every sixteenth bit of the sequence.

In order to make a state machine that fulfills the above requirements, it was realized that if the simple LFSR in Figure 6.1 could be advanced by sixteen states with a single clock, then the outputs of registers 0 through 15 would produce the desired output. This is because the adjacent registers produce the sequence offset by one from each other, as long as no exclusive-or appears between them.

6.2 State Machine Generation

A general method was devised to generate next-Nth-state functions for an existing state machine. This method is exceedingly simple and can be implemented as a simple computer program. For M bits, a set of M string variables is initialized with text representing the initial value of the register that contains them. For example, if the strings are XSTR[0..30], they are initialized to contain the string values "Reg[0]"..."Reg[30]". A *next state* function is then created to compute the next states using boolean operations expressed as strings. For example, if the new value of register 28 should be the exclusive-OR of registers 27 and 30, then the new value of XSTR[28] would be set to the concatenation of the strings in XSTR[27], the string for the XOR operation, "^", and the string in XSTR[30]. The strings will need to be buffered, so that all the operations which generate new values do not destroy the old values still needed by the rest of the registers.

To calculate the advance-by-N functions of the state of the machine, all that is needed is to call the *next state* function N-times. The strings held in each register at the end of this operation are the functions necessary to realize the operation. Figure 6.2 shows the output of the program "bits31.c" listed in the source code area

```
Reg[0]=(Reg[12]^Reg[15])     Reg[16]=Reg[0]
Reg[1]=(Reg[13]^Reg[16])     Reg[17]=Reg[1]
Reg[2]=(Reg[14]^Reg[17])     Reg[18]=Reg[2]
Reg[3]=(Reg[15]^Reg[18])     Reg[19]=Reg[3]
Reg[4]=(Reg[16]^Reg[19])     Reg[20]=Reg[4]
Reg[5]=(Reg[17]^Reg[20])     Reg[21]=Reg[5]
Reg[6]=(Reg[18]^Reg[21])     Reg[22]=Reg[6]
Reg[7]=(Reg[19]^Reg[22])     Reg[23]=Reg[7]
Reg[8]=(Reg[20]^Reg[23])     Reg[24]=Reg[8]
Reg[9]=(Reg[21]^Reg[24])     Reg[25]=Reg[9]
Reg[10]=(Reg[22]^Reg[25])    Reg[26]=Reg[10]
Reg[11]=(Reg[23]^Reg[26])    Reg[27]=Reg[11]
Reg[12]=(Reg[24]^Reg[27])    Reg[28]=Reg[12]
Reg[13]=(Reg[25]^Reg[28])    Reg[29]=Reg[13]
Reg[14]=(Reg[26]^Reg[29])    Reg[30]=Reg[14]
Reg[15]=(Reg[27]^Reg[30])
```

Figure 6.2: Next state functions for the advance-by-16 31-bit LFSR.

of Appendix D. Note that in this case the Fibonocci form is calculated because it is simpler. The output of the program has been placed in two columns in Figure 6.2 for reasons described later. This method works on any state machine and can advance by any desired number of states. In many cases it will be necessary to simplify the resulting functions to minimize the feedback propagation delay.

6.3 Advance-by-16 implementation

The circuit diagram for the major building block of the 31-bit LFSR is shown in Figure 6.3. This "half-row" block makes use of the symmetry in the feedback functions which can be seen in Figure 6.2. Looking at each row in Figure 6.2, we see that the output of the first column function is used as the value of the second column function on the next clock. The circuit block here is referred to as a "half-row" because the final design was roughly decomposed into eight rows of four registers each. Thus the two register circuit is a "half-row".

Registers pairs listed in each row of Figure 6.2 such as 0 and 16, or 1 and 17, are implemented using one of the "half-row" blocks. The left block is a master/slave latch with an integral XOR at the input, while the right is a regular master/slave latch. The combination XOR-MSLatch is identical to that used in the multiphase LFSR design as shown in Figure 5.3. All latches used in this layout are designed

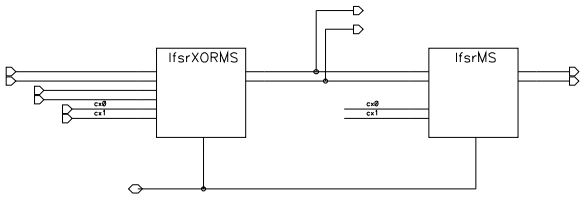

Figure 6.3: A basic building block of the 31 bit LFSR. This "half-row" provides half of the registers that appear in one row of the LFSR layout.

around $1\mu m$ HBTs biased at 0.7mA. Because the outputs of the first 16 latches need to go be used as the parallel 16 bit output, there is a tap off the line which goes between the registers.

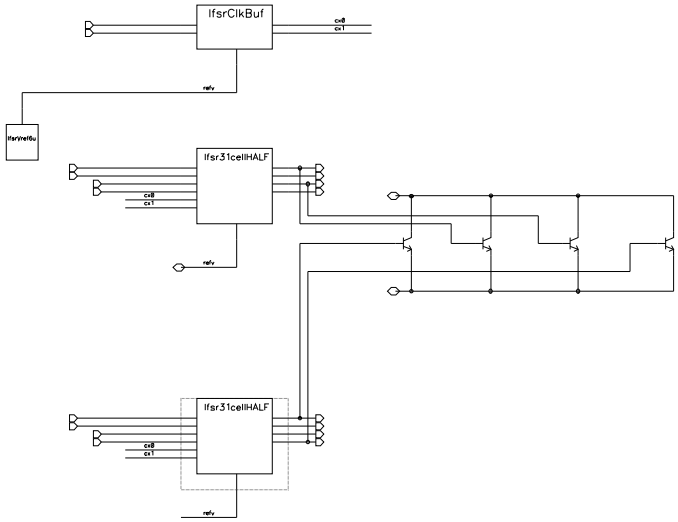

Figure 6.4: LFSR full row schematic.

Each "half-row" is paired with another as shown in Figure 6.4, forming a full-row. This arrangement of building blocks allows the circuit to be formed out of eight nearly-identical row blocks with two output bits supplied by each row. This was a fortuitous design as it enabled changes to the physical output positions of the bits in the layout by changing the wiring between large highly functional blocks. Each row is equipped with a clock buffer and part of a high fan-in AND gate to detect one particular state for testing purposes. The clock buffers use $2\mu m$ buffer, followed $4\mu m$ and $8\mu m$ emitter followers which drive the eight latch loads.

A full schematic of the 31 bit LFSR is presented in Figure 6.5. As can be

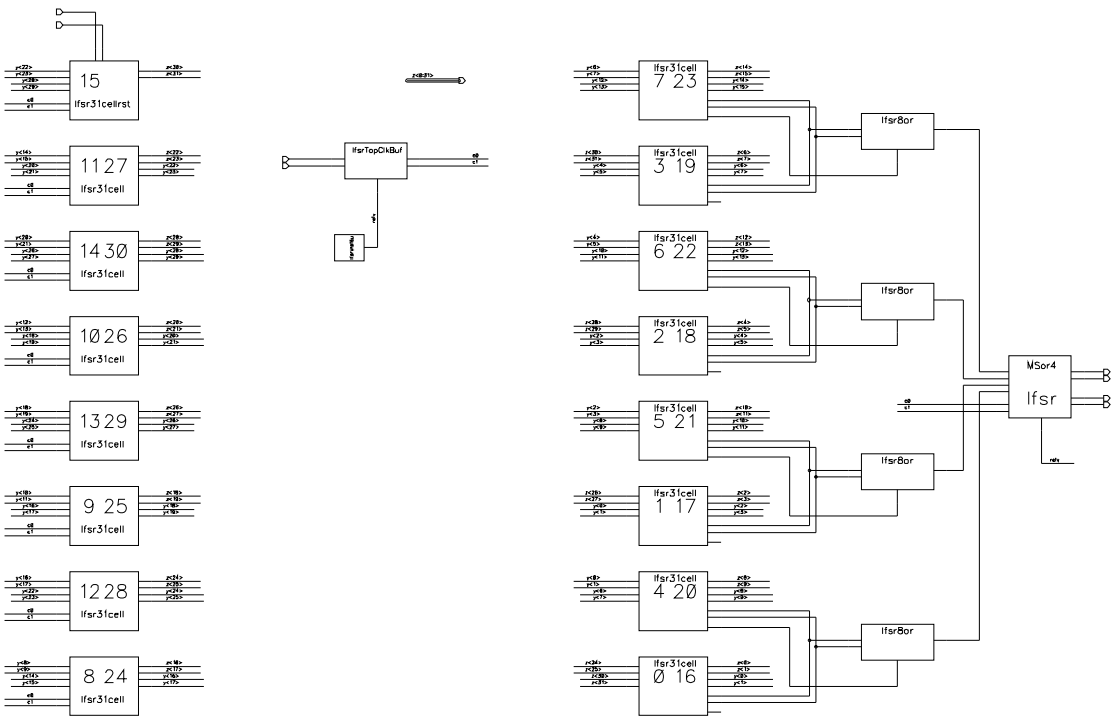

Figure 6.5: SIII LFSR full schematic. This circuit consists of eight interconnected rows, a high fan-in state detector, and a top level buffer.

seen, each row implements four of the necessary registers, except for the top row. Since there are only 31 bits and not 32, register 15 is left isolated. As this register is already asymmetric with respect to the others, the register was also modified so that a reset signal can be applied, which basically just generates a "1" output, capable of forcing the circuit out of the undesired all-zeros state.

The ordering of the registers is determined by the requirements of the next circuit block, (the 16-to-4 multiphase multiplexer), which is described in the next chapter. It required the 16 output bits to be grouped together into sets of 4 with the same modulo-4 remainder, {0,4,8,12}, {1,5,9,13}, {2,6,10,14}, and {3,7,11,15}. This ordering is visible in the register numbers on the left of each block, looking from bottom to top. Each group of two rows produces the four bits in a modulo-4 group. A pre-buffer for the row clocks is shown in the middle. It is identical to the type used in the rows.

It was at first planned to use a series of dividers to generate a test clock signal

from the LFSR. However, because the number 2^{31}-1 is a Mersenne Prime, there are no smaller factors which can divide the clock. A whole 31-bit counter would be the only type of circuit which could implement this, and the propagation delays related to the carry signal would make the circuit difficult to implement correctly in the targeted operating rate. Instead, a two stage single-sided high fan-in AND gate is used to detect the all ones state and to output a pulse when it is encountered. The circuitry to the left of Figure 6.5 shows the combination of the ANDs from each of the rows to generate the final output signal. Each pair of rows use a single ended high fan-in gate architecture similar to that used in the SERDES II LFSR. The four outputs of these ANDs are combined in the next level and latched and divided by two so it can be easily read from an external pin as a square wave.

6.4 Layout

The layout of the 31 bit LFSR is shown in Figure 6.6. The circuit was laid out using a reflective power rail arrangement which shows mirror symmetry every two rows. This same power rail scheme was used for the rest of the chip as well. The circuit uses 320mA of current at 3.6V, giving a power consumption of 1.1W. There are 752 HBTs. The circuit is intended to operate up to 2.5GHz, and the two stage clock buffering gives a propagation delay from clock input to data output of approximately 110ps.

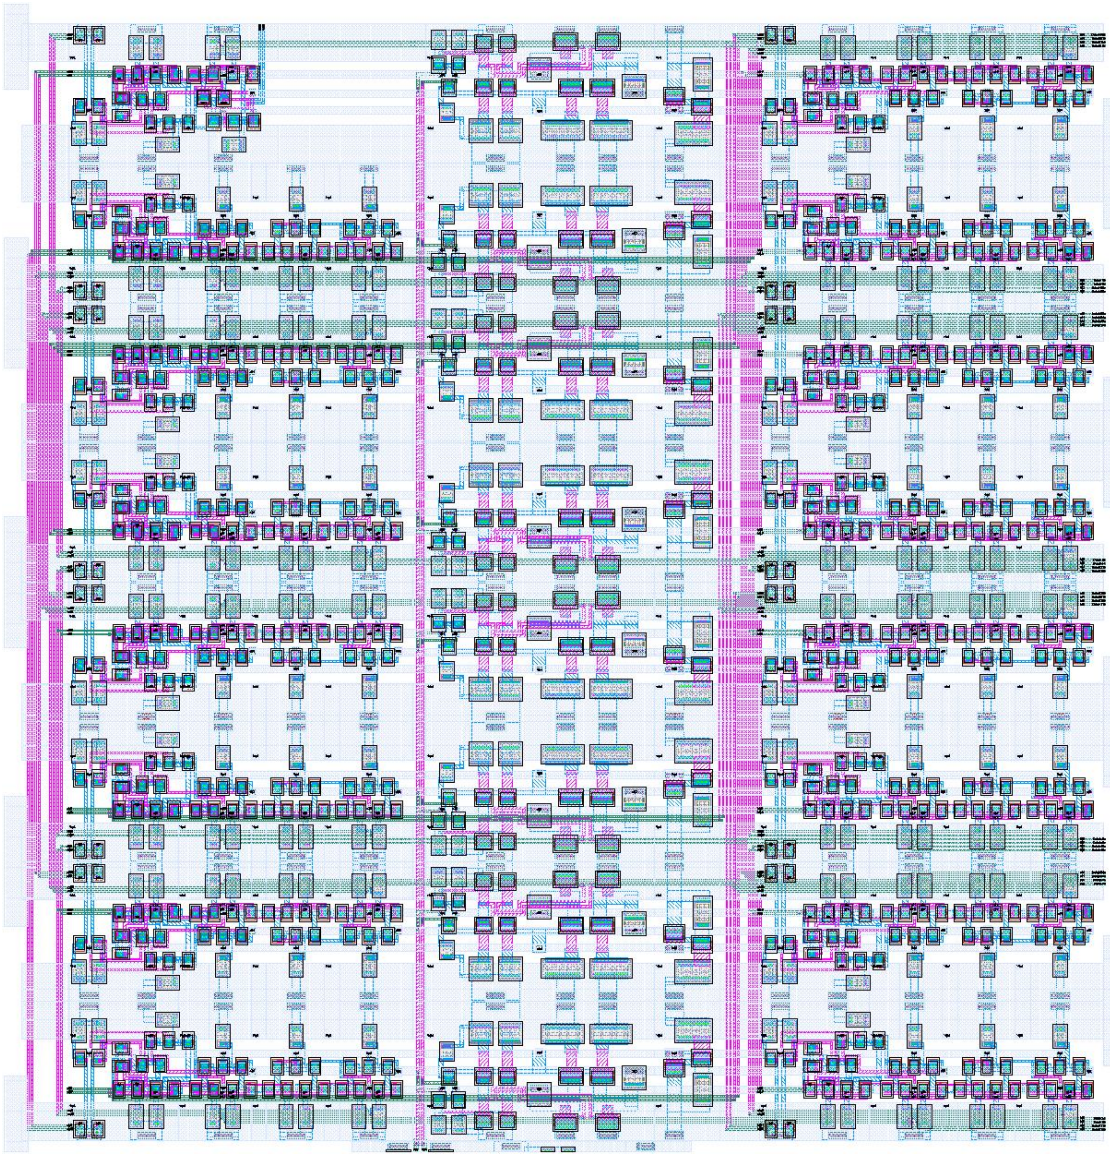

Figure 6.6: 31-bit LFSR circuit layout.

CHAPTER 7
16-to-4 Round-Robin Multiplexer

In order to supply the Edge Steering Multiplexer with four bits of parallel data from a 16 bit parallel source, some form of pre-multiplexer is required. As previously discussed in the Chapter 4 on the ESMA, the Edge Channeled Multiplexer Architecture requires four data inputs to be clocked at offsets of 0, 90, 180, and 270 degrees at the frequency of the quarter-rate global system clock.

For this reason the outputs of the 16-to-4 multiplexer must be similarly staggered by 90 degrees. Due to the availability of complementary signals, two physical differential clock signals at 90 degrees from each other are used. The complements of these two quadrature signals provides the other required clocks. During the design phase I used the designations $\overline{X}$, $\overline{Z}$, X, and Z for the phases to match up with the phase names chosen for the Edge Channel Multiplexer. The 16-to-4 multiplexer should also generate a signal usable by the upstream parallel data generator to indicate when it is desirable for the input data to the 16-to-4 multiplexer to change.

The required high data rate made this a very difficult design challenge. A single propagation delay in this technology can easily generate a 90 degree phase shift at our target operating clock rate of 10GHz. This can result in significant difficulty in designing a circuit which can operate at the highest speeds, as well as over a wide frequency range.

Although largely independent of clock rates, I have found the propagation delays through CML gates can vary to a proportionally higher degree when the rise and fall times make up much of the bit period. Also important is the ratio of the fixed delays to the time difference between clock phases. If the fixed delay causes a signal to cross a fixed clock edge which later affects it, (at the target rate), then slowing the operating rate may very well cause the signal to be clocked before it has the correct value.(see Figure 7.1.)

Ideally, the propagation through any given circuit should be less than the difference in the phases of the two clocks which latch the values before and after the

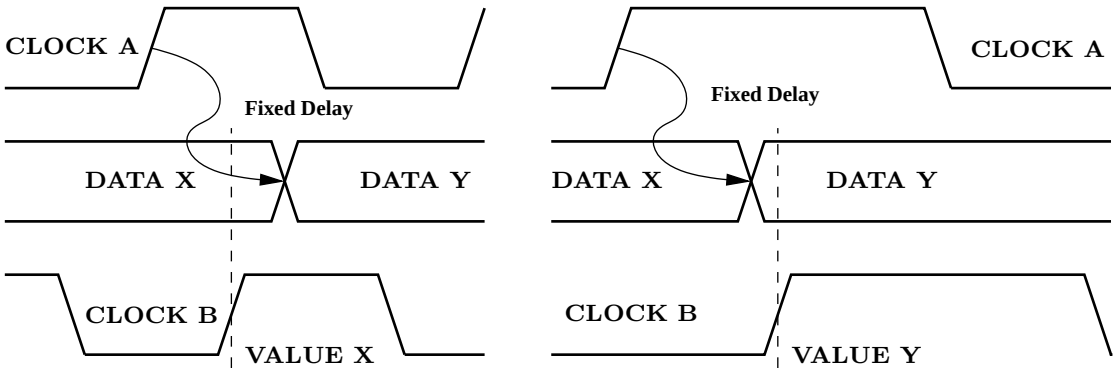

Figure 7.1: Slow Clocking Errors. On the left, CLOCK B causes DATA X to be sampled, which changes a fixed time after CLOCK A. Due to the fixed propagation delay, when the clock rate is decreased, the value DATA Y is sampled instead.

signal has passed through. This would allow the clocking rate to be reduced to an arbitrarily low value without violations. In order to achieve the highest data rates however, this desirable condition must be violated. The delays between our phases are on the same order as the delays through the simplest of gates. In this kind of high performance circuit, if the clock rate is reduced then there is the chance that the delayed edges will leave their safe-designated time intervals.

Several types of multiplexers were considered for this design, however the unique requirements of the final 4-to-1 Edge channel architecture, the high rise time to period ratio, and the quarter rate clocking scheme necessitated a novel design which has not previously been seen in the literature.

7.1 Quad 4-bit Shift Register

A scheme very similar to that used in SERDES II was initially considered for this design and it is shown in Figure 7.2. Four individual parallel loading four bit shift registers could be triggered by the four individual clock phases, (as opposed to the two phases used in SERDES II). Each of the registers requires its own loading signal, generated presumably from its own clock signal. Because pairs of differential phases are their own complements, each physical clock line needs to drive at least 8 master-slave latches and the load signal generators for two registers. This is difficult to achieve without buffering, which would add to the delay between input clock and

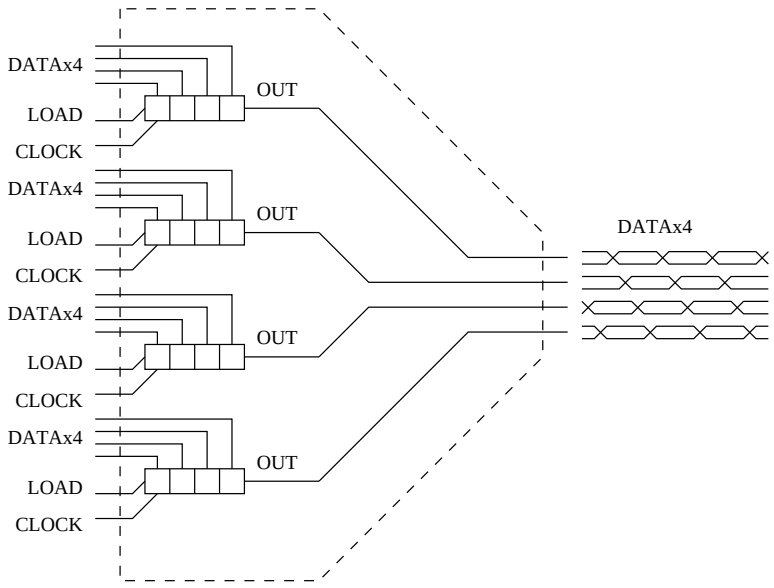

Figure 7.2: Quad 4-bit Shift Register. This is the basic design used in the SERDES II serial transmitter.

output data.

The details of a representative register can be seen in Figure 7.3. After each register is emptied, a multiplexer at the input of each master-slave D-Latch is activated to load data from the parallel inputs rather than from the previous latch in the register chain. The multiplexers can actually be encorporated into the master latch current tree, ensuring the smallest delay. In order to make this system work, the load signals would need to be timed to rise and fall *within* one individual bit-shift time, on every fourth bit. Any divider and buffer circuits used to generate the load signal would introduce a significant propagation delay relative to the original clock, and would make accurate timing over a varied frequency range extremely difficult.

The timing for a single register is shown in Figure 7.4. The output data is delayed by approximately a quarter period(A), while the clock÷4 is delayed by nearly half a clock period(B). These delays are expressed as fractions of a clock period at the target clock rate, 10GHz, but would remain fixed at those durations of 25ps and 50ps even if the clock period were to be increased or decreased significantly.

As the load pulse must be of a shorter duration than the high time of the clock÷4 signal, we must use an MS-latch to generate a load signal from it using the main clock(D). It should be noted that although each register operates on

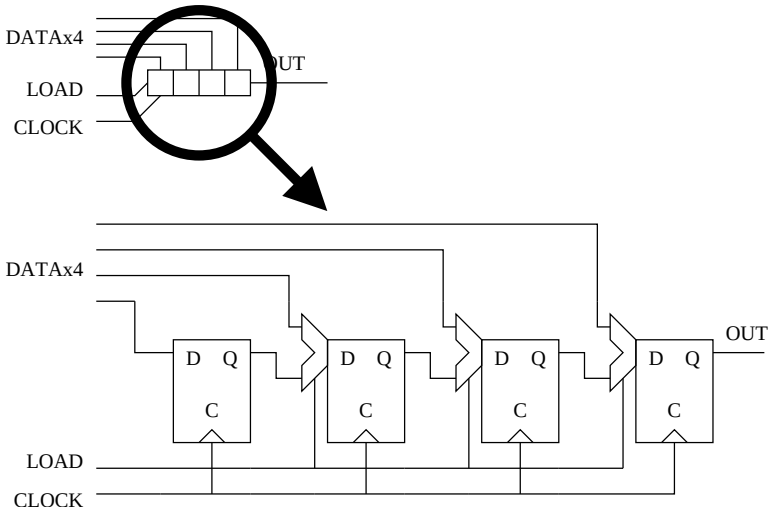

Figure 7.3: Shift Register Details. Each of the four registers in the previous figure is equivalent to the one diagrammed above.

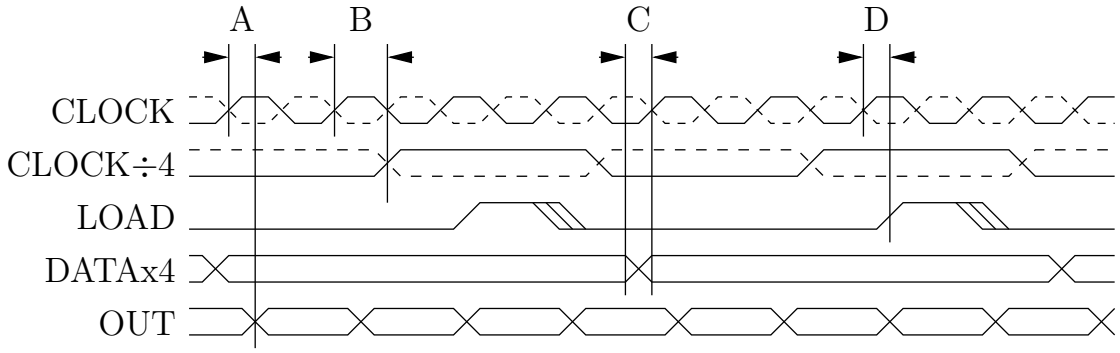

Figure 7.4: Shift Register Timing Diagram. Shown here are the labels for the descriptions contained in the text for the four shift registers.

a different phase, the input DATA, (x4), to each register set will need to change simultaneously. Therefore, there is a very small window(C) during which the input data must change, between the fourth bit being shifted out and the first of a new group being clocked in through the parallel line via the multiplexer. Because the load enable signal should be finished while the data is stable, for all registers, it may be necessary to shorten the load pulse even further, by using an AND to combine it with the complement of the clock.

At first it was thought that identical load generation circuitry could be created for each of the four registers independently, but this was not the case as will be

explained in the next section. A single state machine still might be used to create all the loading signals, but that would spoil the symmetry and create more possible points of failure. Considering the very similar design in the SERDES II transmitter featured an extremely tight layout, (using shared deep trench regions), and the fact that simulations indicated problems well below our target rate of 10GHz, another solution was sought.

7.2 16-to-4 Multiphase Multiplexer

Another basic type of multiplexer in use in serializer circuits is the latched tree architecture [42]. Tree architectures are in general more complex than simple shift register multiplexers, growing in area and power consumption with the square of the number of input bits, as opposed to the linear growth of shift register designs. However, tree architectures have a number of advantages which compensate for this. Typically, the input clock is only used to drive only a small number of loads directly, usually the final output latch and a clock divider. This does not change as the number of inputs is increased and the delay between input clock and output change is minimal.

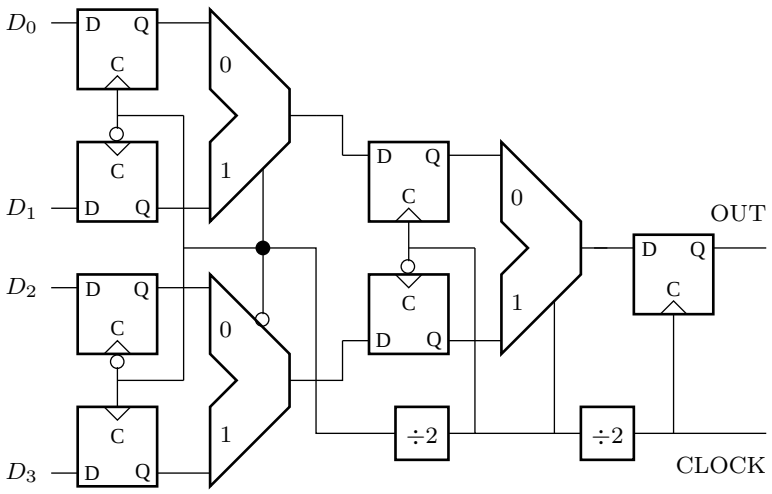

Figure 7.5: Standard Latched 4-to-1 Multiplexer.

In the above 4-to-1 circuit, (Figure 7.5), a secondary and tertiary clock is generated from the original using ÷2 circuits. Master/slave latches hold the inputs to the individual 2-to-1 multiplexers stable except for when transitions are absolutely

necessary. In the presence of fast enough rise times, edges which are present on the inputs can still generate significant noise in the output, (even when not selected.)

Due to the nature of the logic family used in our designs, the circuit is implemented slightly differently. CML trees can handle many inputs, but generate only a single output. It is possible to combine the 2-to-1 multiplexer with the MS-latch immediately following it, reducing the number of current trees needed as well as propagation delays. The blocks shown in Figure 7.6 thus have two separate clock signals going in, one at half the rate of the other. The latch operates at the higher rate, sampling the output of the multiplexer at the rising edge.

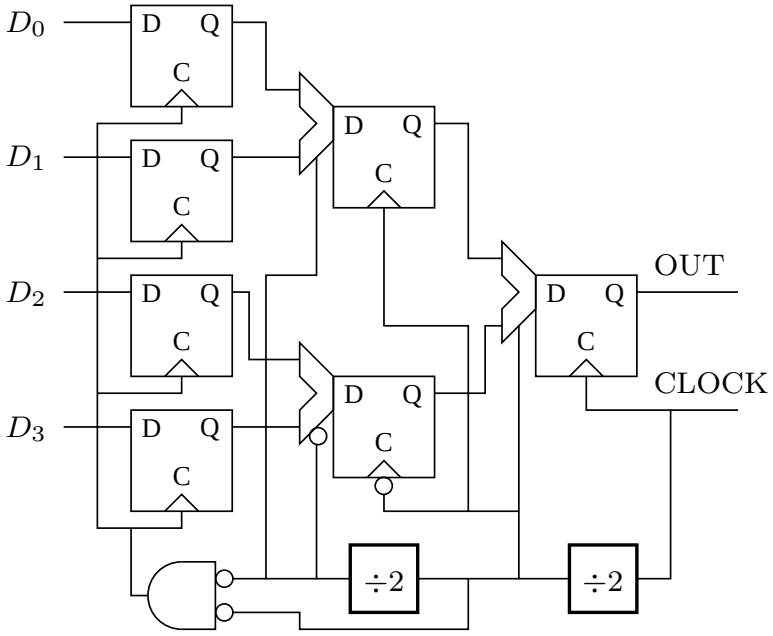

Figure 7.6: CML Implementation of the Latched 4-to-1 Multiplexer.

Three input levels are required for the CML Mux-Latches, Level 1 for data, Level 2 for the multiplexer, and Level 3 for the clock. Because the latches require a Level 3 clock, they require an emitter follower to generate that clock from the select signal from the multiplexer in the following stage. Because of this, the multiplexer switches to the new input just **before** the old input changes state. This allows maximum settling time at the latch output before the multiplexer re-selects it. Separate input latches are left at the beginning. It was initially thought that four identical 4-to-1 multiplexers, including their dividers could be used to realize

the needed 16-to-4 functionality.

7.2.1 Hidden state

During simulations it became apparent that there was some sort of problem with the design. The output bits were all clean and the internal edges appeared where expected, however the output bit stream failed to correctly match the input. The problem was quickly traced to the divide-by-two circuits, shown in various forms and with timing diagrams in Figure 7.7.

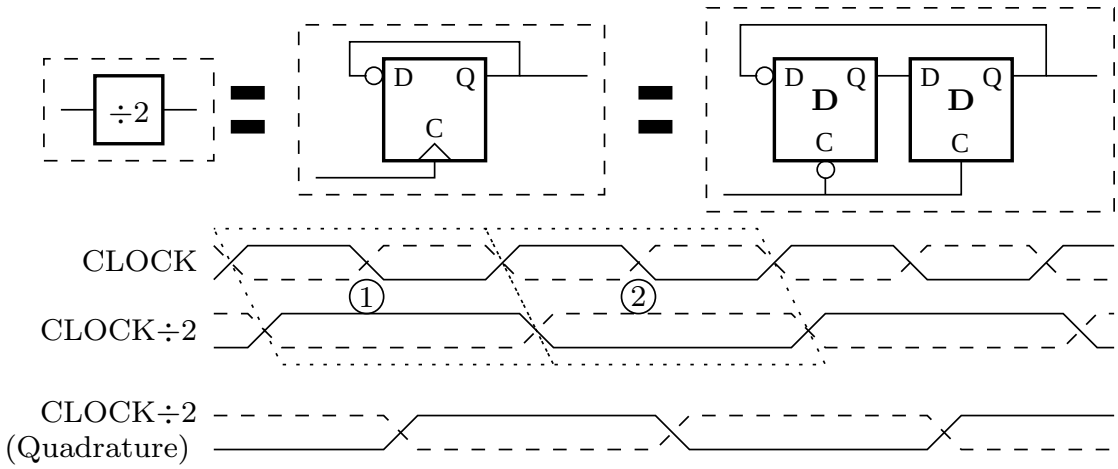

Figure 7.7: CML Divide By Two Circuit. This diagram illustrates the hidden state stored in the feedback.

The upper part of Figure 7.7 shows that a divide-by-2 circuit is effectively a Master-Slave latch with its output fed back into its input after an inversion. On the far right is an equivalent diagram showing the master-slave latch represented as two D-Latches. In addition to the output wire shown in the figure, the quadrature of the output signal is also available from the line connecting the left latch output to the right input.

Looking at the input waveform during the time the output is high**(1)**, it can be seen that it is identical to the input waveform when the output is low**(2)**. This is due to a necessary electronic state being retained in the circuit for a time longer than the input waveform period. When this circuit is powered up, it can initially be in either one of the two states, and this is the key to the problem. There are effectively four global clocks, each with specific phase offsets, and their divided signals must

retain the same offset in time in order for the whole multiplexer to operate properly. In the absence of initialization circuitry, the four dividers can be powered up in different states, leading to the divider outputs at times being 180 degrees out of phase with respect to the desired output.

This leaves only two chances out of sixteen that the circuit will power up and run correctly. Even combining the dividers for those phases which are complementary, we are still left with a 50% chance of the phase being incorrect, and that is only at the output of the first set of dividers. Since we go through two stages of division, the problem is further compounded and the chances of all the dividers producing the correct phases by chance becomes marginally slim. This same clock phase problem would also occur when using the four serial shift register architecture described in the previous section if separate divider circuits are used to generate each load pulse.

There are several methods which can be employed to correct this situation. One method might be to use a global reset which brings the divided outputs into correct alignment. However, the introduction of reset circuitry would be prohibitive in terms of speed loss. Added delays inside the highest speed divide by two circuits would immediately impact the top speed of the circuit.

Secondly, circuitry to add inversions as necessary at the output of each divider could be added. This might consist of exclusive or gates which would allow the signals to pass either complemented, or uncomplemented. Unfortunately extensive control circuitry would be required to detect out of phase conditions. An external control might be realized for this test circuit, however detecting the condition externally would be difficult using the equipment we presently have. Regardless, the presence of the exclusive or gate itself would introduce an additional gate delay into critical timing areas.

Lastly, it might be possible to use a single state machine which might generate all the divided signals while ensuring correct phase. That type of solution, if workable, would also be applicable to the shift register architecture. Unfortunately, a simple unified circuit cannot be achieved with a short enough feedback path to operate at the desired rates. In addition, different wire lengths going to different multiplexers from a centralized state machine circuit would lead to earlier timing

violations. In order to test the edge steering architecture it was desirable to have as fast a multiplexer as possible.

7.2.2 Solution

The solution to this problem came in the realization that the symmetry of the divide-by-two circuit itself can be exploited in realizing a unified clock generating circuit that was also distributed equally throughout the entire 16-to-4 multiplexer. As mentioned previously, one of the first stage dividers can be shared between two of the individual 4-to-1 multiplexers because pairs of multiplexers make use of complementary clocks. Because the divide-by-two circuit is made up of two D-latches, one can be located inside each of the paired 4-to-1 circuits. The inversions are handled via wiring, so there is no source of asymmetry.

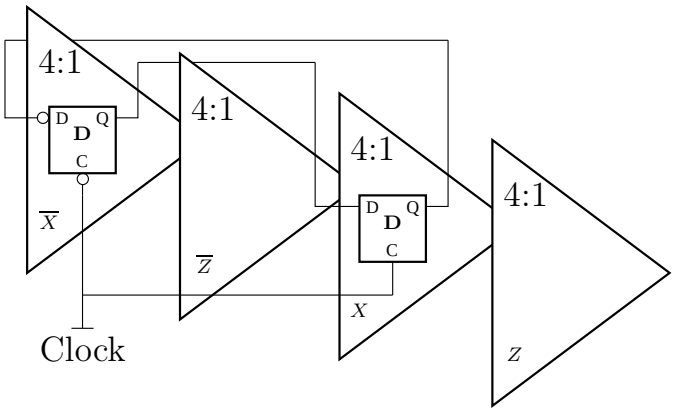

Figure 7.8: Distributed ($\div 2$) clock architecture chosen for the 16-to-4 multiplexer.

Shown in Figure 7.8 is a first level divider circuit. The input clock shown can be considered to be both the X and $\overline{X}$ phases. The output of the upper D-Latch is equivalent to $\overline{X} \div 2$, while the output of the lower D-Latch is equivalent to $X \div 2$, making the required signals available where they are needed.

The other pair of 4-to-1 multiplexers require their own phases and so cannot directly use the outputs of these dividers. If we used the same circuit design with a different phase as inputs we would still have the hidden state problem as mentioned previously. To ensure this does not occur, the design uses the MS-Latch halves

present in the second pair to **sample** the originally generated half-speed clock, generating a new half-speed clock aligned on the alternate phase.

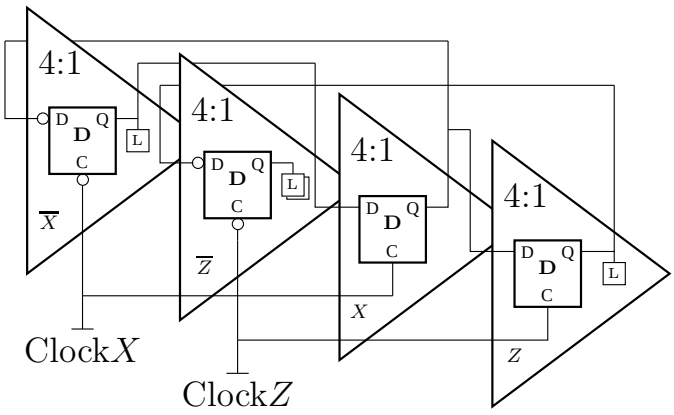

Figure 7.9: Distributed ($\div 2$) clock architecture plus sampling.

Figure 7.9 shows the method. The "Primary" master/slave pair in $\{X, \overline{X}\}$ generates a half rate clock using a full rate input phase X, while the secondary master/slave pair in $\{Z, \overline{Z}\}$ generates a delayed version aligned with the other full rate clock input, Z. The key is that there is only one actual divide-by-2 for this particular divider level. Each of the 4-to-1 blocks is virtually identical. Wiring is how the difference in functionality is achieved, and dummy loads, (shown as L-blocks), are added to balance the circuit in "holes" in the layout specifically left available for that purpose.

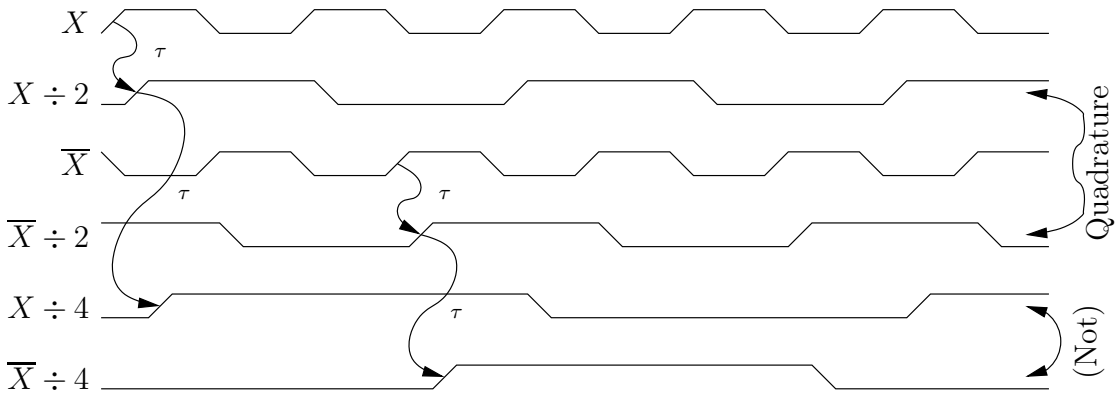

Figure 7.10: Clock Division Details. $X \div 2$ is in quadrature with $\overline{X} \div 2$, but $X \div 4$ is not in quadrature with $\overline{X} \div 4$

With the addition of the second stage divide-by-two requirement, it turns out

that a full pair of master/slave latches are required inside each 4-to-1 multiplexer to generate the divided-by-four clock required. This is because the quadrature of a half-rate clock generated from a specific phase will not necessarily be aligned with the quadrature of that phase. Figure 7.10 makes this clear. However, $\overline{X} \div 4$ cannot be generated from the quadrature generation of $X \div 4$.

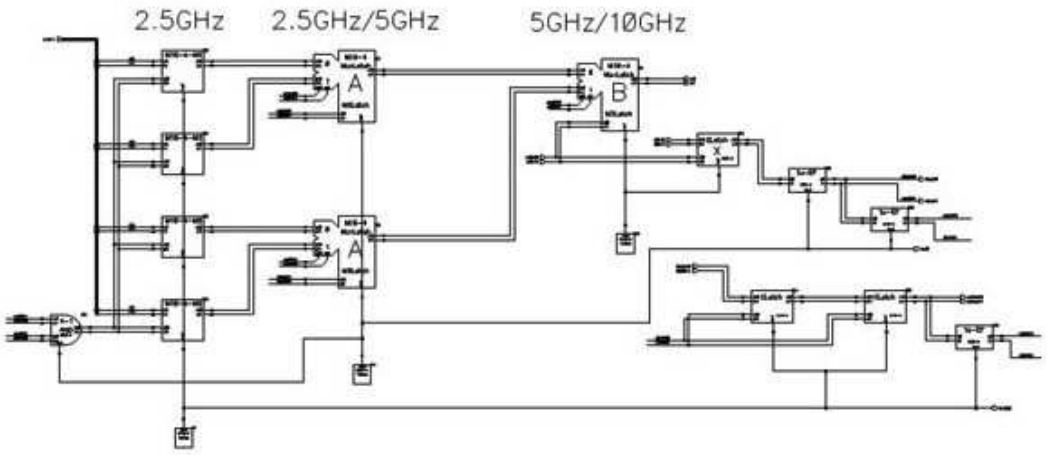

Figure 7.11: Schematic of the 4-to-1 multiplexer used inside the 16-to-4 multiplexer.

Figure 7.11 shows the actual schematic of one of the final four-to-one multiplexer blocks. Most of the components in the schematic above are which are not present in 7.6 are emitter-followers used to shift between voltage levels. To the right of the leftmost Mux/D-Latch block is a half divider, followed by two emitter-follower blocks. Below them is a complete master-slave pair to generate the divided-by-four signal for this multiplexer. These latches are clocked by the locally generated half rate clock. As with the divide-by-two latches, these cannot operate independently or they can be out of phase. To ensure correct operation, one of these pairs, (in this case the ones in the Z 4-to-1 multiplexer), acts as the divider and the other pairs sample this one. Again, this is all handled in the wiring, so each block is laid out identically. As with the divide-by-2 stage, dummy loads are added in the final layout to balance the clock loading.

The AND gate at the lower left of 7.11 shortens the effective load pulse length so that their union across all the 4-to-1 multiplexers leaves a larger update window for the LFSR to advance within. In operation, the AND takes in the complements of

the divided-by-two and divided-by-four signals and this results in a 25% duty cycle load signal with the high part appearing during the last quarter of the divide-by-four period. This means that the rising edge of the load signal actually occurs at the falling edge of the divide-by-4 clock inside each 4-to-1. Without this, the aggregate loads would be high for 200ps $+$ $3\,x\,25$ps of the 400ps period, leaving only a 75ps window during which it would be safe for the parallel input to be updated. With the load pulse shortening technique, only 100ps $+$ $3\,x\,25$ps is consumed, leaving 125ps free for loading.

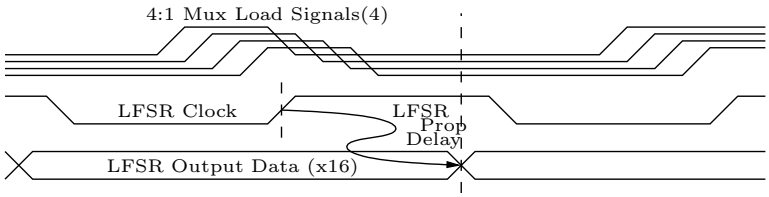

Figure 7.12: Generation of the clocking signal for the LFSR pseudo-random number generator.

The advance signal for the LFSR is derived directly from one of the quarter-rate clocks, and the rising edge coincides with a falling edge of the corresponding load signal. As the propagation time for the LFSR was at least 100ps in simulations, the clock from the $\overline{X}$ multiplexer was chosen as the trigger. This places the changes centrally within the update window at the targeted clock rate of 10GHz, but in retrospect a later phase should have been chosen to allow slower clocking to occur. This is shown in Figure 7.12.

The schematic for the entire 16-to-4 multiplexer is shown in 7.13. There is an extremely complex interconnection between the identical 4-to-1 blocks. The dummy loads are the components on the right. One of the dummy loads also serves to buffer the signal used to trigger the advance of the LFSR, which supplies the inputs to the 16-to-4 multiplexer.

The divided-by-four clock used to trigger the advance of the LFSR was chosen to cause the output to change inside an allowed "update window". Each 4-to-1 multiplexer making up the 16-to-4 loads on its own clock.

The full 16-to-4 multiplexer was laid out as is shown in 7.14. The four sub-multiplexers are each identical copies and the wiring in the center channel is re-

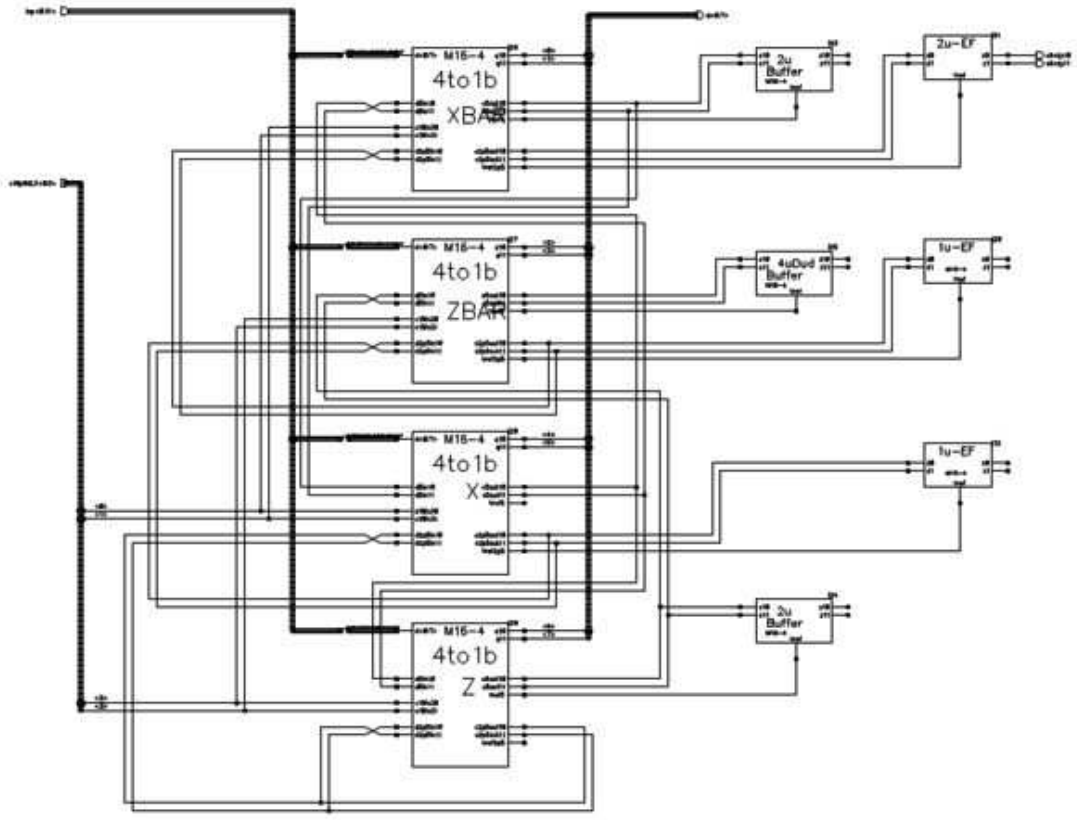

Figure 7.13: Complete schematic of the 16-to-4 multiplexer.

sponsible for making each divider half work correctly. Small open areas to the right of the central wiring channel were left to place the dummy balancing loads. The inputs for each of the 4-to-1 multiplexers is on the right side and matches up with a corresponding output on the LFSR described in the previous chapter. A control output to trigger the LFSR advance is at the lower left. Inputs from the clock phases enter from the right, and the four round-robin outputs are on the right as well.

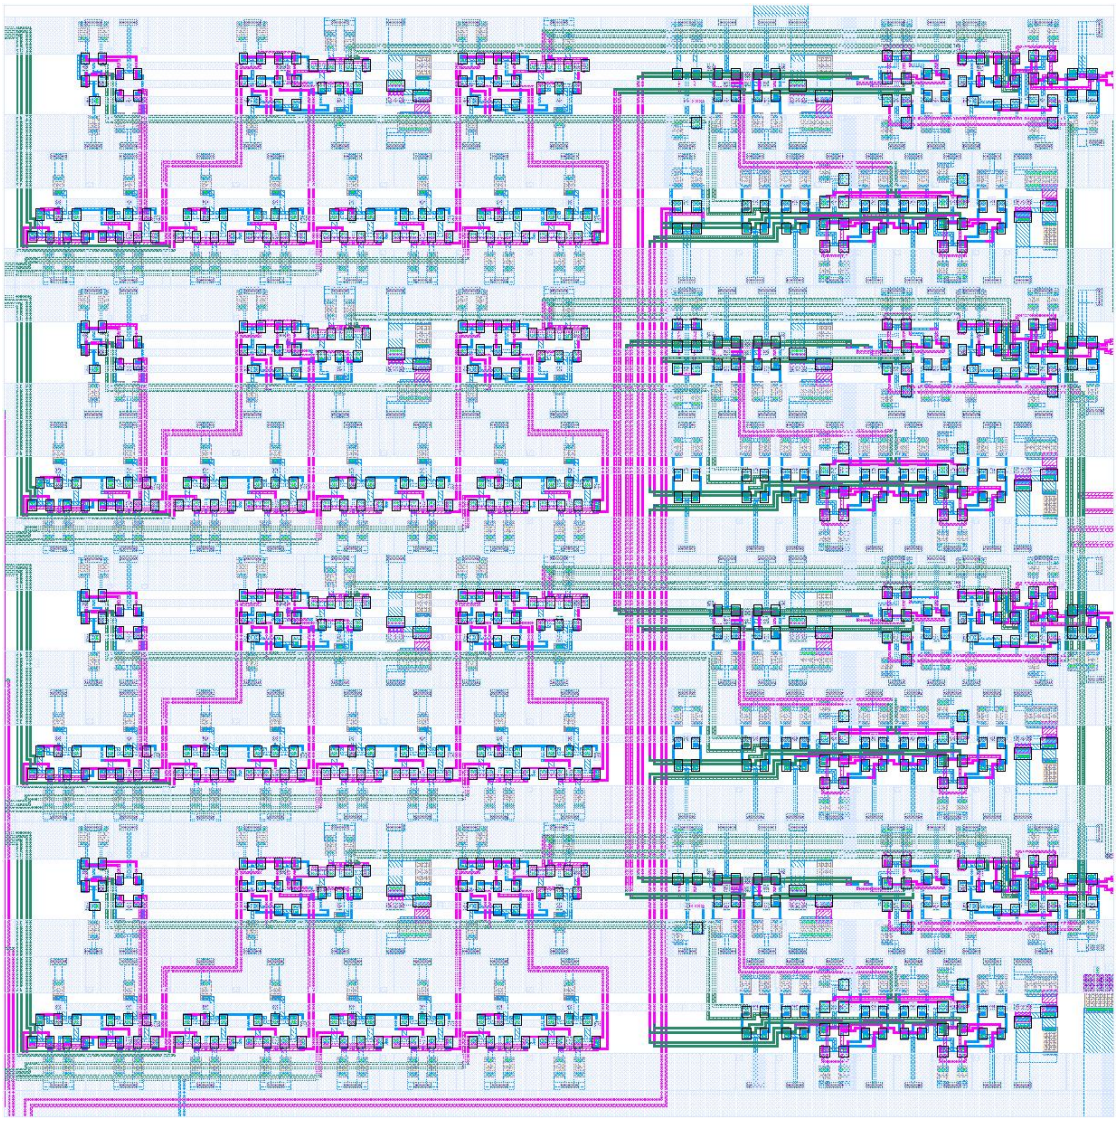

Figure 7.14: Complete Layout of the 16-to-4 multiplexer.

CHAPTER 8
Transmitter

In previous chapters various components necessary for a serial transmitter design were described. The 31-bit LFSR provides 16 bit pseudo-random parallel data, which is fed to the multiphase 16-to-4 multiplexer which presents that data in a round-robin fashion to the final 4-to-1 Edge Steering Multiplexer. Everything would driven by a four phase clock.

To test these components, a complete serial transmitter chip was designed. Originally, the design was to have been fabricated in the 5DM process. The 5DM processing option was later discovered to be unavailable, and so the design was converted to 5AM. In order to submit the design on time, this transmitter was re-designed for basic functionality. Most notably, the simpler and slower 5AM current-starving VCO was used which limits the expected operating range to below 30Gbps.

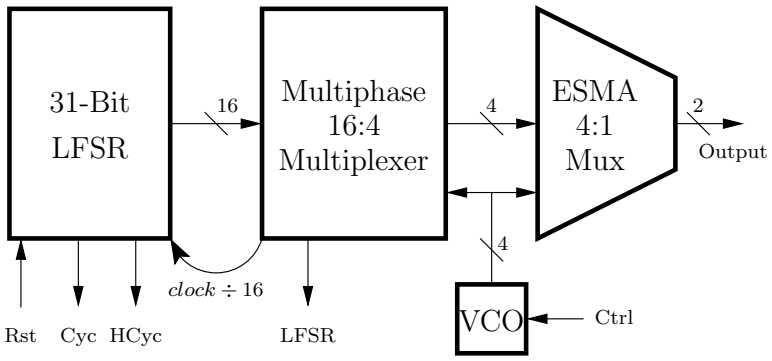

Figure 8.1: Block Diagram of the SERDES III Serial Transmitter.

The serial transmitter block architecture is shown in Figure 8.1. The three major components and the VCO which were described in previous chapters were combined to form a single chip. In addition to the blocks shown, additional circuitry was added to provide control signals to the VCO, the LFSR reset, and to get signals off chip for testing purposes.

8.1 ESMA Implementation Details

The previous chapter on the ESMA dealt with the design at the architecture level. In this section, the details of the design and layout for this chip will be presented.

Figure 8.2 shows the Edge Steering Multiplexer Schematic at the top level. Apart from the blocks along the bottom, which represent voltage references, each block represents a 2-to-1 multiplexer. The leftmost, or "top tier" multiplexer blocks are all identical in size, as are the three middle ones.

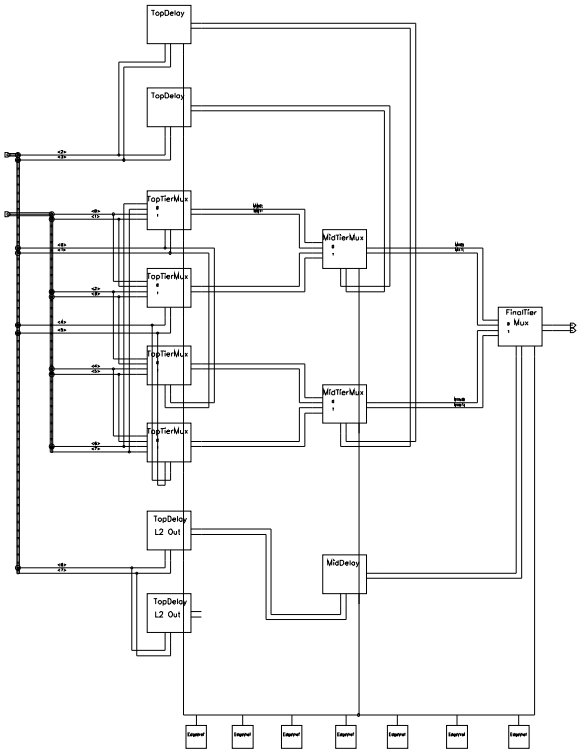

Figure 8.2: Edge Steering Multiplexer Schematic.

The central four multiplexers of the top tier correspond to those in the earlier presented block diagram of Figure 4.7. Those above and below the four are used as delay elements which delay the clock signals from reaching the next tier earlier than expected. The delay multiplexers are used in parallel pairs so that the clock loading on all phases will be identical. The bottom multiplexer of the middle tier is also a delay element, delaying the final Z-clock the proper amount before the final output multiplexer is reached.

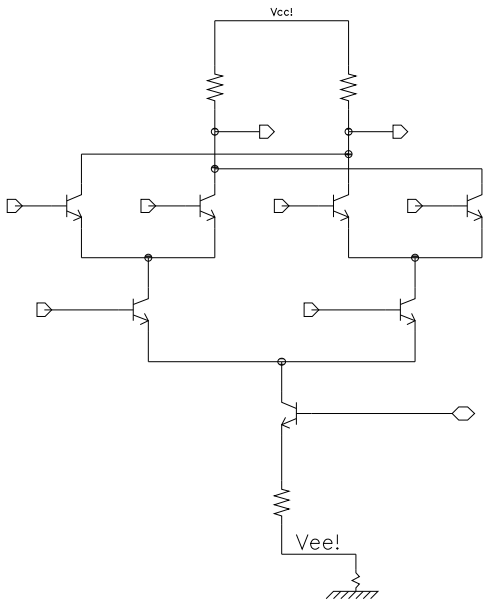

Figure 8.3: Top Tier 2-to-1 multiplexer.

The multiplexers of the top tier are of the standard CML type, (see Figure 8.3), and have emitter followers to produce outputs on level two for use by the next tier. The multiplexers use $3\mu m$ transistors, but the top level input transistors are $2\mu m$ transistors to reduce loading at the input. The emitter followers are also $3\mu m$, with extra diodes added as shown in Figure 8.4 to allow level three outputs as necessary.

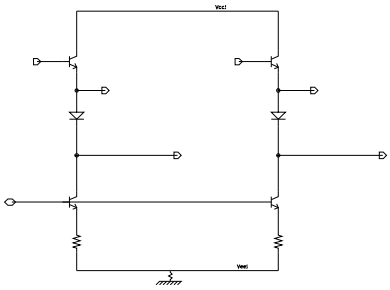

Figure 8.4: Top tier 2-to-1 emitter follower schematics.

The major differences between the blocks on the top tier are that the delay blocks have the data inputs to the multiplexers connected to locally generated voltage references. This way, the select input of a delay multiplexer generates edges in precisely the same way the edges are generated by the multiplexers which take data as input. Another difference is that since the mid-tier multiplexers need level

three inputs on their select lines, the level three outputs are used from the upper two delay multiplexers in the first tier. The lower two delay multiplexers on the first tier output on level two as they drive an additional delay before the final multiplexer. This mid tier delay should be the same as the delay encountered by an edge passing through from a data input, not one generated from a select line. Therefore, the mid-tier delay uses locally generated voltage references to hold fixed the select input to the multiplexer. The unused input to the multiplexer is tied to the common mode voltage to ensure the least imbalance. For the mid-tier, the multiplexers were composed of $6\mu m$ transistors and suitably biased, apart from the upper inputs, which used $4\mu m$. The emitter followers at this level were $6\mu m$ as well.

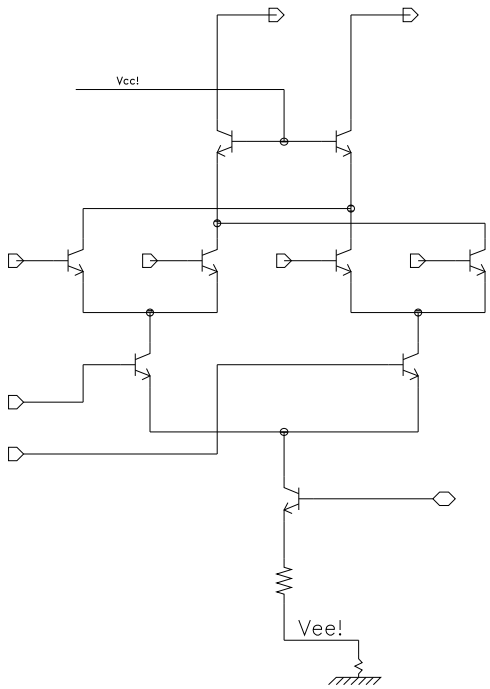

Figure 8.5: ESMA Final multiplexer schematic.

The last 2-to-1 multiplexer in the design drives the bond pads for sending the signal off-chip. It uses a cascode unity-current gain circuit at the top to supply current to the bondpads with minimal loading of the multiplexer tree. A schematic is shown in Figure 8.5. Each transistor in this multiplexer is $8\mu m$, biased at $0.6\text{mA}/\mu m$, giving output currents of 4.8mA. Driving the 50Ω impedance of the coaxial cable, this generates output swings of 240mV on each line, or 480mV if

measured differentially.

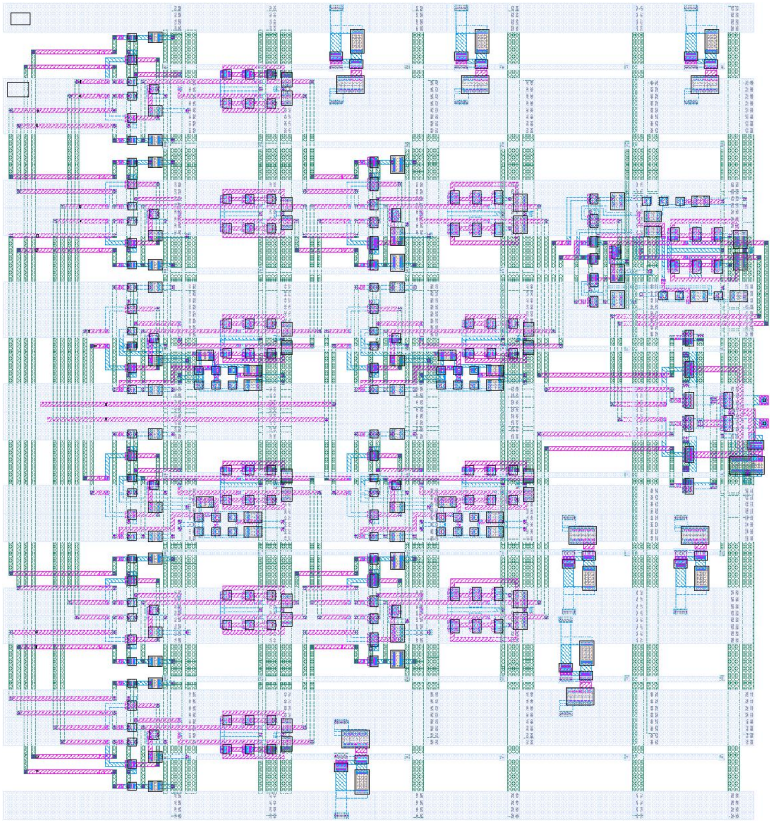

Figure 8.6: ESMA Layout for SERDES III Transmitter chip.

The layout for the final 4-to-1 Edge Steering Multiplexer is shown in Figure 8.6. It conforms closely in positioning with the schematic shown in Figure 8.2, apart from the voltage references across the bottom. The seven voltage references for the current mirrors are scattered throughout the layout and work in paralleled. Each uses $6\mu m$ transistors and sinks 3.6mA of current. The layout conforms to the power rail structure used in the other components so that the outputs on the right of the multiphase 16-to-4 multiplexer are aligned with the inputs on this circuit's right. In total, this circuit sinks 113mA and contains 196 HBTs. Using the 3.6V nominal supply voltage, this corresponds to approximately 0.4W for the final 4-to-1 ESMA.

8.2 SERDES III Transmitter Chip

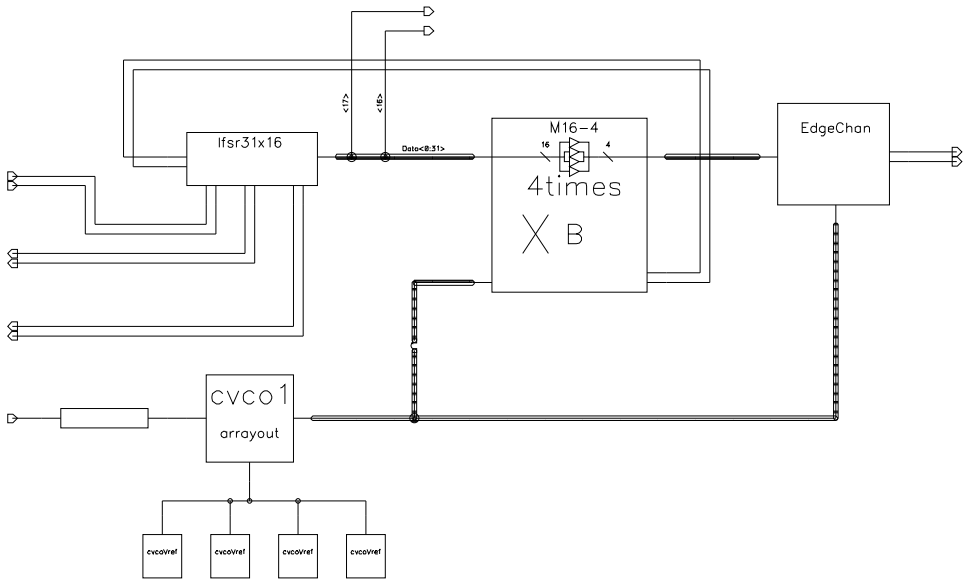

Figure 8.7: SERDES III transmitter schematic diagram.

All of the major components for the SERDES III transmitter chip were combined to form a complete integrated circuit. As seen in the schematic above, the layout design closely matches the block diagram. The connections on the left and the single connection on top are routed to the output pad circuits, (see Appendix A on Technology). The LFSR reset input is on the top left, followed by the LFSR Half-Cycle and Cycle signals below it respectively.

As discussed in the previous chapter on the 31-bit LFSR, the reset signal is used to ensure the state machine can be moved from the forbidden "all ones" state. The Cycle and Half-Cycle outputs are intended to indicate each time the LFSR completes its $2^{31} - 1$ output pattern. The Cycle output is a clocked value of the output of the high fan in state detector. It should produce a regular pulse the width of a clock period each time the LFSR "rolls over". The Half-Cycle output is equivalent to the Cycle signal fed through a divider, so that it is a square wave signal that changes state with each roll over. The LFSR takes a very long time to roll over. A 5GHz clock would trigger advances in the LFSR every 800ps, yielding a 800ps pulse every 1.718 seconds on the Cycle line. The Half-Cycle, or "HCycle" line would toggle every 1.718 seconds.

The fourth connection on the left of Figure 8.7 is a control signal for the VCO. This signal passes through a voltage divider and filter before reaching the control input of the 5AM current starving VCO. The connection at the top is bit 15 of the data passed to the 16-to-4 multiplexer. On the right is the output of the edge steering multiplexer. It does not make use of pad drivers but instead directly drives the bond pads.

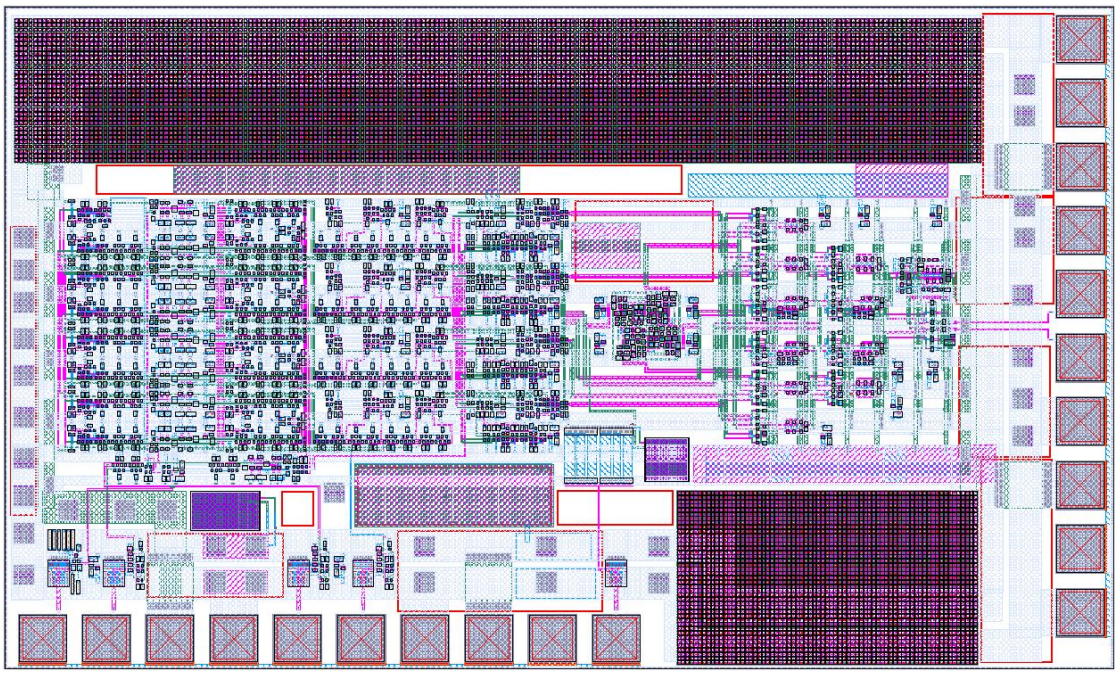

Figure 8.8: SERDES III transmitter layout.

As can be seen in Figure 8.8, the various components fit quite well together sequentially from left to right in the layout. These components are in order, the 31-bit LFSR, the multiphase 16-to-4 multiplexer, the VCO, and the edge steering 4-to-1 multiplexer. The output of the transmitter needed to be aligned with the central bondpads of one of our 10-pin probes, as these have the highest bandwidth. The row of pads along the bottom provide control signals and the lower speed test outputs for the chip. There are pad drivers and receivers attached to pad connections as well as ESD diodes for static protection. The total size of the chip was $2610\mu m\ x\ 1550\mu m$. This chip was manufactured through the MOSIS fabrication service. It was submitted in 2004 on June 28th, and we received the fabricated chip on the following November 1st.

8.3 Testing

The SERDES III transmitter chip has undergone testing using our probe station. The manufactured 5AM chip is shown in Figure 8.9. The chip size is $2.61mm\ x\ 1.56mm$ and it has 1825 HBTs.

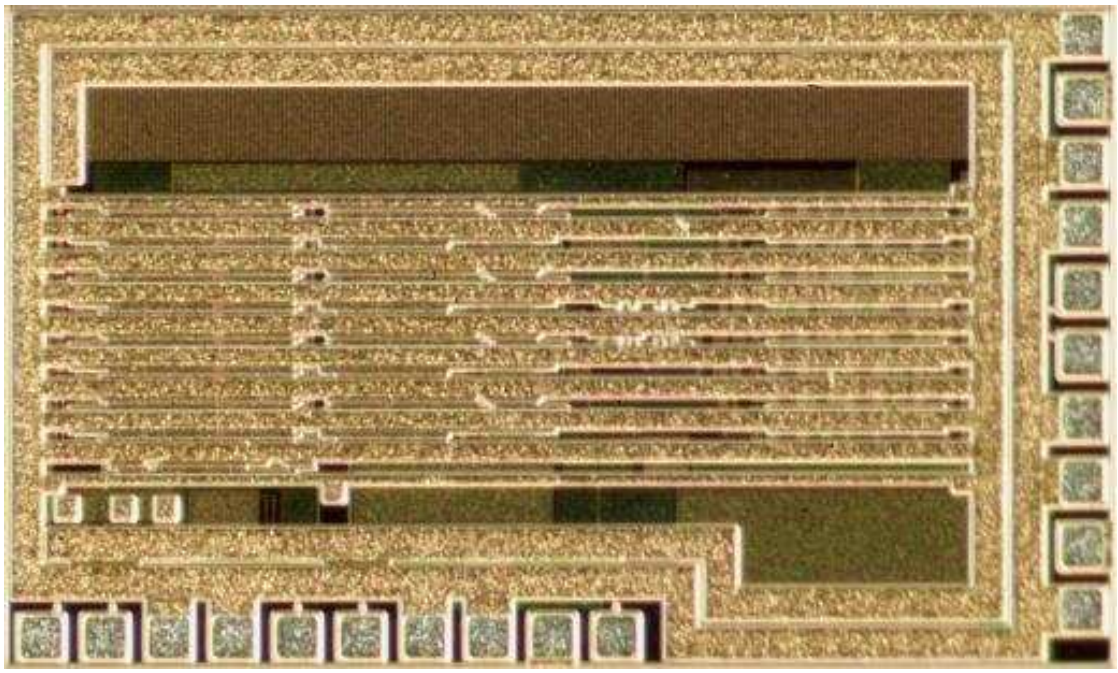

Figure 8.9: Microphotograph of the SERDES III transmitter.

Preliminary testing of this design included checking the supply current, and this was found to be the expected value of 760mA at the nominal supply voltage of V_{EE}=-3.6V. This yields a power consumption of around 2.7W for this chip. As power was applied, the signals were detected at the test signal outputs as well as at the transmitter output. The control voltage was varied to adjust the frequency to its maximum value, which was expected to be on the order of 7GHz due to the current-starving 5AM VCO design which was included in this chip.

The LFSR was observed to be operating correctly as the eye diagram in Figure 8.10 shows. This indicates that the dividers which make up the multiphase 16-to-4 multiplexer are providing a one sixteenth data rate to the LFSR. This corresponds to a 27Gbps output as shown in the output eye diagram in Figure 8.11

This eye is shown single-ended, and should essentially appear twice as "open"

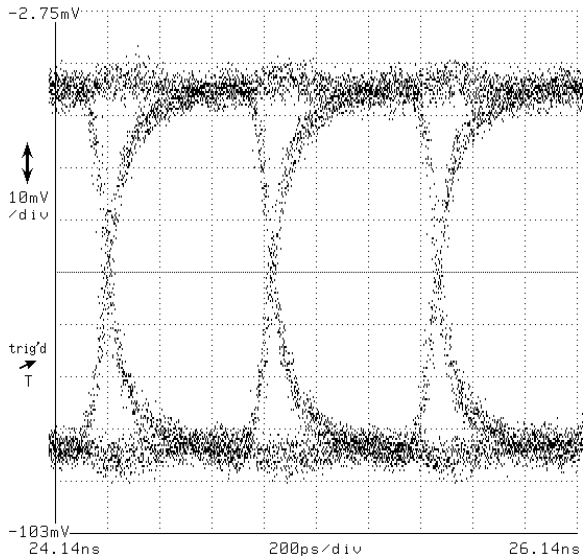

Figure 8.10: LFSR 1.6Gbps eye output.

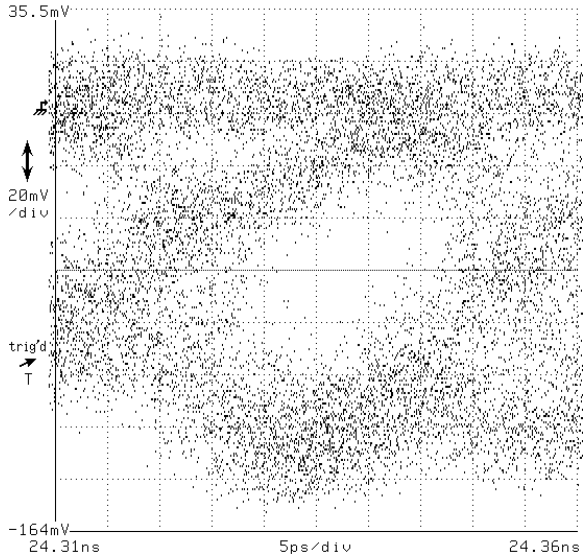

Figure 8.11: Transmitter 27Gbps output eye diagram.

when measured differentially. In addition, a small mistake in the connection between the 16-to-4 multiplexer and the VCO resulted in bits being updated during the incorrect phase as they pass into the edge steering multiplexer. This should be corrected in the next fabrication run, and the newer VCO designs, (see Chapter 3), should be capable of reaching much higher data rates.

To test the Cycle output, it was used as a trigger while sampling the LFSR

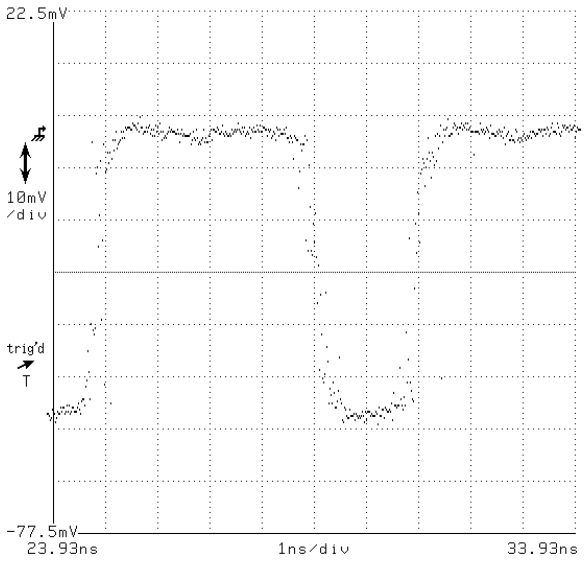

Figure 8.12: LFSR captured output waveform.

output. The scope was left in "infinite persist" mode. As can be seen in Figure 8.12 the LFSR steady waveform was produced which indicates that the Cycle and HCycle outputs were correctly detecting the "roll-over" of the LFSR.

In Figure 8.13, a series of open eyes is shown at approximately 15Gbps. The uniformity in the upper part of the eye is due to the stability of the upper ground reference. The differential output should show a smoother lower level. Even single-ended, this series of eyes does not show serious duty-cycle distortion, which is the primary design characteristic that the edge steering multiplexer architecture was intended to address. The alternate-eye size variations are explained in Chapter 9.

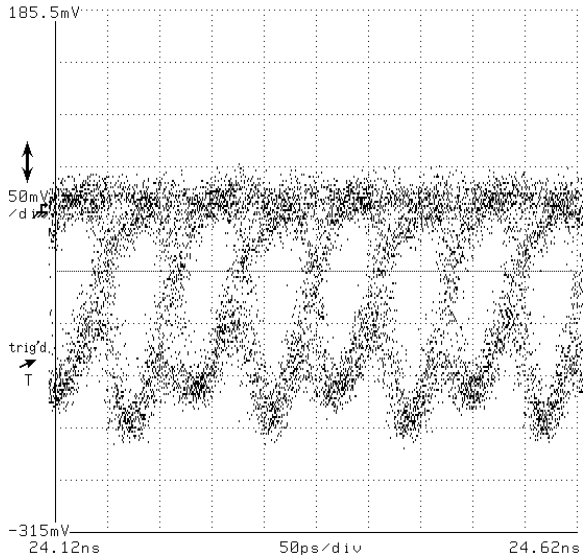

Figure 8.13: Captured output consisting of multiple eyes.

8.4 SERDES III Transmitter Revision

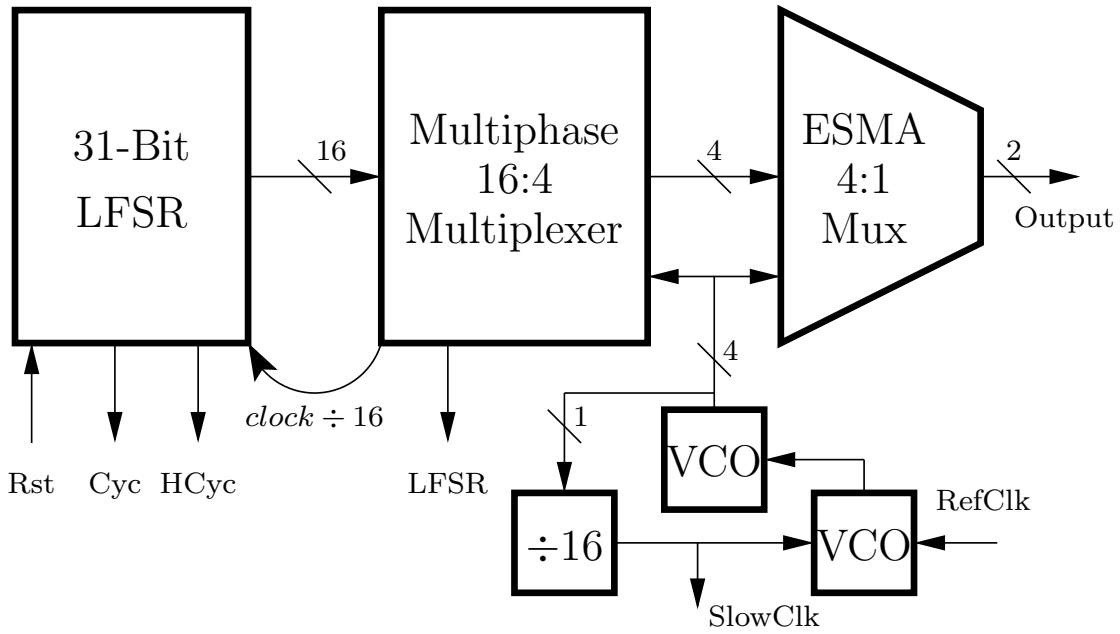

Figure 8.14: Revised Transmitter Block Diagram. This design features a new VCO, PLL circuitry, and other enhancements.

A additional pair of newer transmitter chips have been designed to further test the capabilities of the new architecture. The most significant changes are the inclusion of faster VCO designs, (see Chapter 3), which are embedded in a PLL. One chip is expected to operate in the 35Gbps range, while the other could potentially reach 40Gbps. The new chips have better load balancing, and use separate voltage references for each of the final 2-to-1 multiplexers in the ESMA. Other enhancements include more test inputs and outputs. The block diagram of the revised transmitter is shown in Figure 8.14.

8.4.1 PLL

I included the PLL circuitry originally intended for the last design in the latest revision. This circuit uses a divide-by-16 to act as a clock multiplication unit(CMU). A 625MHz input clock is sent through a capacitively-coupled pad receiver to the 3-state phase detector described in Appendix C.

An active loop filter was designed, which is also described in in the Appendix C. It consists of a simple two stage amplifier with RC feedback components to adjust

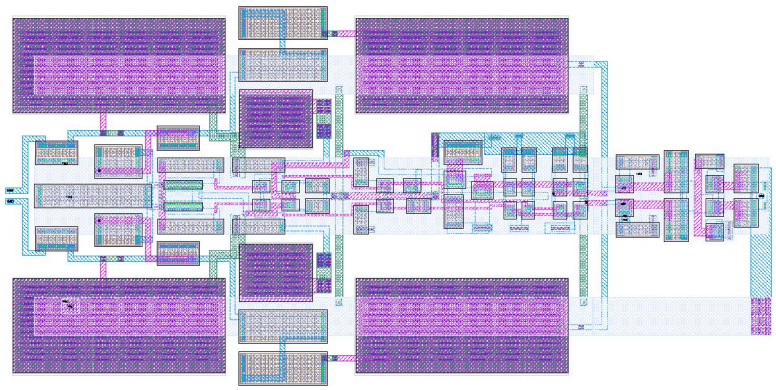

Figure 8.15: Loop filter circuit layout.

the pole locations of the PLL. The large capacitor areas and design rules relating to floating nodes made this circuit block layout particularly difficult.

8.4.2 Layout

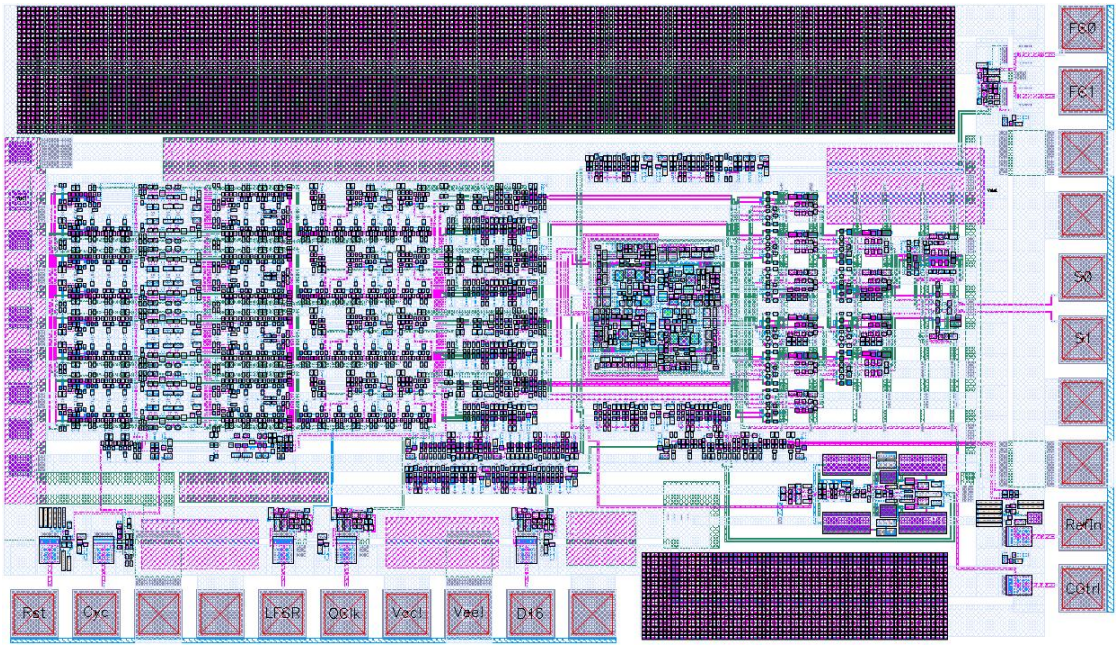

Figure 8.16: Revised Transmitter Layout

The finished layout for the revised Transmitter is shown in Figure 8.16. This can be compared with the layout in Figure 8.8. As can be seen, the new VCO is larger and actually required widening the entire chip by nearly $100\mu m$. Above the VCO is a clock load matching circuit, while directly below is the phase detector and

divider for the PLL. The loop filter is located in the lower right corner. This chip is 2702μm by 1550μm and contains approximately 2300 HBTs.

8.4.3 Simulations

Simulations of the revised chip, (Figure 8.17), show improved performance over that of the previous design, however, the higher data rates still suffer eye closure due to data dependent jitter effects, and the generally too low bandwidth of the output multiplexers.

As can be seen, the closure of the eye gets progressively worse as the data rates increase, despite the improvements to the clock loading balance and the isolation of the voltage references. Chapter 9, (on the performance limitations of multiplexers), further illustrates why it is difficult to alleviate this problem and why eye closures at rates approaching 40Gbps in this technology are to be expected.

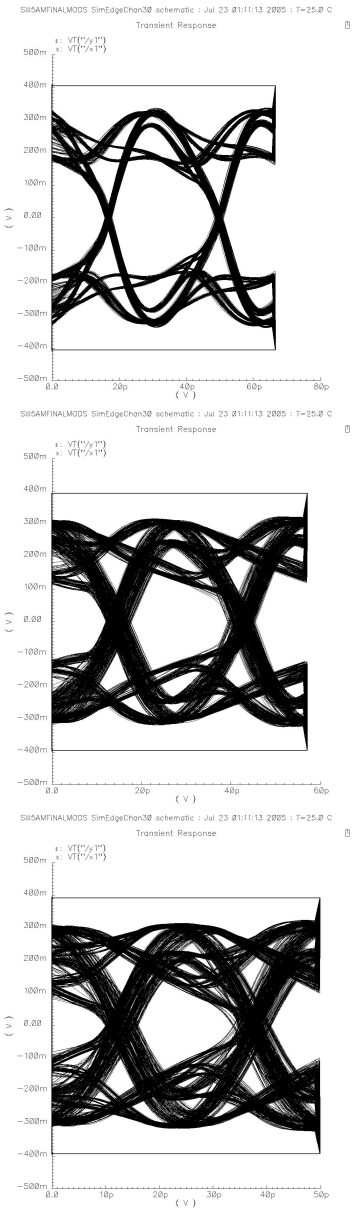

Figure 8.17: Revised Transmitter Simulations. This figure shows the simulated progression from 30Gbps, to 35Gbps, up to 40Gbps.

CHAPTER 9
Performance Limitations

This chapter covers some of the effects which fundamentally limit the further increase in operating frequency in this technology. In a traditional transmitter design, the output is retimed by a latch operating at the full data rate, or by a multiplexer operating at half the target data rate which thereby introduces all edges in the output stream from a single clock source. The architectures used in this work use quarter rate clocks with multiple phases to attempt to produce good output when the edges are generated in multiple locations, and then combined. The circuit which must operate at the highest rate is a 2-to-1 multiplexer, and this chapter describes some of the factors which ultimately limit performance when these 2-to-1 multiplexers are used in serializer circuits. Also, this work will be contrasted with that of others [43], [49], [45] who produce much cleaner eye diagrams and explain the apparent discrepancy.

9.1 Multiplexer Overview

The fully differential 2-to-1 multiplexer usually consists of 6 transistors, (not including any used in a current source or in input level shifting circuits). These are organized in differential pairs in a tree. The circuit can be seen in Figure 9.1. Constant current, I_{BIAS} flows through the tree, provided by the source at the root which is usually implemented with a current mirror or a resistor. The differential select inputs selectively enable either the A inputs or the B inputs to allow the current to be diverted to either of the top two resistors, causing the A or B signals to propagate to the Z outputs.

The primary limitations of this circuit are inherent in the devices. As described in the Technology Appendix A one of the primary limiting characteristics of the SiGe 5AM HBT device is the peak f_T current. The peak f_T current is the bias current which maximizes a transistor's unity gain cutoff frequency, the frequency when its current gain is reduced to unity due to the parasitic effects of the junction

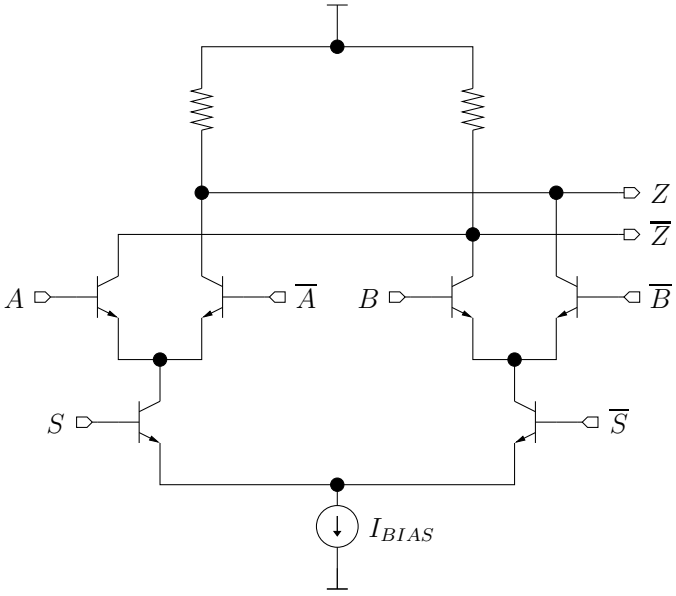

Figure 9.1: Standard fully differential 2-to-1 CML multiplexer used for many designs. The select inputs S at the bottom of the tree steer current to the left or right, activating alternate input pairs A and B to pass through to the outputs Z

capacitances. The equation for f_T is given in Appendix A, equation A.1. Plots of f_T verses bias current I_C give the peak f_T current as seen in Figure A.5, also in Appendix A. They show that approximately $0.6 \ldots 0.7 mA/\mu m$ of transistor emitter width gives the highest f_T of nearly 50GHz. Since the emitter length is fixed at $0.5\mu m$, this gives a current density of around $1.2 \ldots 1.4 mA/\mu m^2$.

In our designs, the magnitude of the I_{BIAS} current in a tree is usually chosen to be the peak f_T current of a *single* transistor, as opposed to twice that current. This means that when the differential input is zero, (at the midpoint of the transfer characteristic), each transistor is biased at **half** the peak f_T current. Because of this, the switching is perhaps not as fast as it ultimately could be. The reason we must work with this reduced bias current is that the maximum DC current allowed for these transistors is about equal to the peak f_T current. The transistors would be over their maximum value when they were switched on, at twice the peak f_T current. Note that emitter followers can be biased at the full peak f_T current as no steering is involved.

9.2 Switching Noise

Multiplexers **require** large signal non-linear operation. One signal must be blocked, and the other allowed to pass. Current is steered through one transistor of the lower differential pair between transitions, enabling only one of the upper pairs. The process of switching current to the other input buffer is a large signal operation, and the different effective loads and states of the transistors of the lower pair gives rise to asymmetric behavior in the transition.

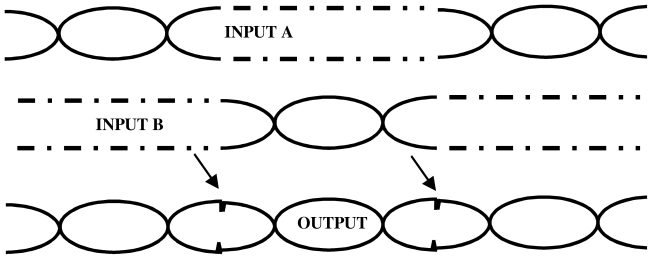

Figure 9.2: Switching noise in 2-to-1 multiplexers. The "glitch" is present even when the signals are the same.

The "on" transistor has a large C_{BE} diffusion capacitance, and because of this a downward step input change at the base is effectively echoed in the common-emitter node of the lower pair. At the same time, the "off" transistor gets a positive step input at its base, which due to the temporarily depressed shared emitter node, is actually larger than the steady-state V_{BE}-on value. Turn-on and turn-off times are asymmetrical, leading to a temporary deviation from constant current in the tree. Base currents can be very large depending on the rise/fall times, and as a result there is a significant "glitch" in the output waveform even when switching between two input signals which have the same value. The effect is more pronounced on the output which is transitioning from high-to-low, as it is turned on using an overly large V_{BE} signal, and the turn-on time is generally shorter than the turn-off time.

Figure 9.3 shows the output of the ESMA running at 10Gbps with sharper than required clock input transitions. The distortion caused by the switching noise lasts nearly as long as a 20Gbps bit time. When the edges are output closer together, this "glitch" intrudes into the next bit, causing an alternate-eye distortion like that seen in Figure 8.13. Note that the second figure is measured in a single-ended sense,

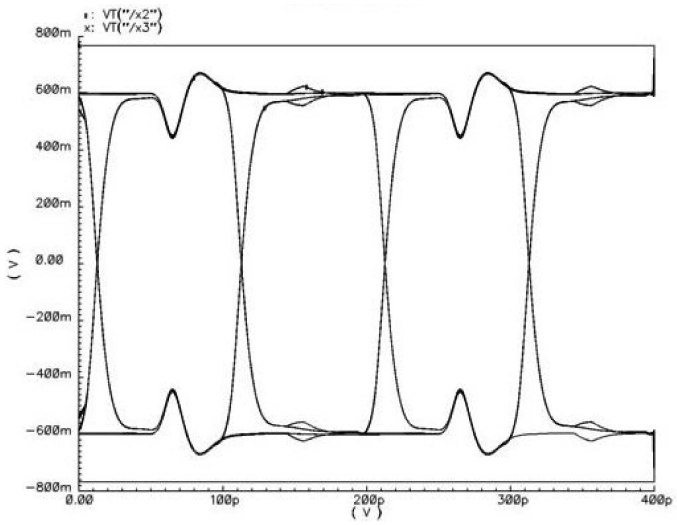

Figure 9.3: Simulated 10Gbps switching noise in the ESMA 4-to-1 multiplexer.

and that when measured differentially would show the distortions on both the upper and lower parts of the eye.

9.3 Weak Switching Distortion

As stated in the previous section, the ESMA architecture uses a final 2-to-1 multiplexer which switches between two serial input streams at a time when they have the same value. Before a switch, a bit is allowed to pass from one multiplexer input, along with a half-bit. Next, the other input is selected, and the half bit is completed, as well as the full bit from the second input. In essence, there are two "classes" of bits, those which are passed, and those which are shared. The eye diagrams of Figure 8.17 consist of superpositions of both kinds of bit eyes. The nearly horizontal lines passing through what would normally be the interior of the eye are really only generated by the output signal line with the lower voltage level. The differential measurement is what makes the effect symmetrical.

The output stream suffers from an additional alternate-eye distortion, (see Figure 9.4), which also shows the "weak low" output signal in alternate eyes. This "weak low" occurs only under the following conditions; The input selected before a select transition has a bit of value X to pass through the multiplexer, and the

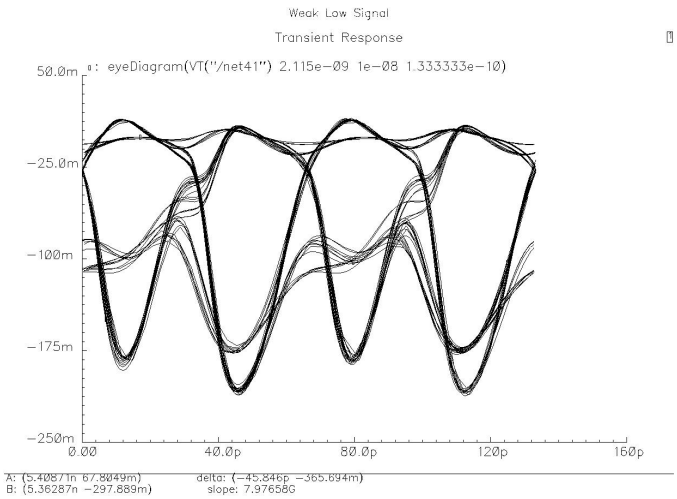

Figure 9.4: Single Ended Multi-Eye Plot. This shows the alternate-eye distortions due to switching noise, as well as the "weak low" distortions which result in a horizontal line cutting through the eyes at data rates approaching 40Gbps.

bit which is split between the two inputs is of the same value. In addition, the value of the bit which the multiplexer will pass after the select transition is $\overline{X}$. This $\ldots 110 \ldots$ or $\ldots 001 \ldots$ sequence with the center bit coinciding with a multiplexer SELECT change, (the shared bit), results in a weak low output signal for this middle bit.

The reason for this is that the select transition is not fully completed before the second input's bit starts to change in value from the shared bit value, to the complementary value. If there were symmetry in the rise and fall characteristics of the outputs of CML gates, the "weak low" would be accompanied by a "weak high". However, because the low output is driven by the collector at an "on" transistor, it changes more rapidly than the high output which is connected to 'off" transistors, which have a passive pull-up.

Alternately, one can view this as a lack of full current into the newly selected input buffer pair. The "off" transistor is not affected, so the high output remains high on the higher output line. However, the "on" transistor does not have full current until the select transition is complete. If the input to that "on" transistor also begins to fall, the effectively higher g_m of the "on" transistor will reflect that in the collector current, and the low output will begin to rise sooner than the high

output would begin to fall.

The end result is that because the select signal cannot "hard switch" inputs quickly enough, part of a bit change which should happen in the output later causes an earlier effect.

9.4 Multiplexer Bandwidth

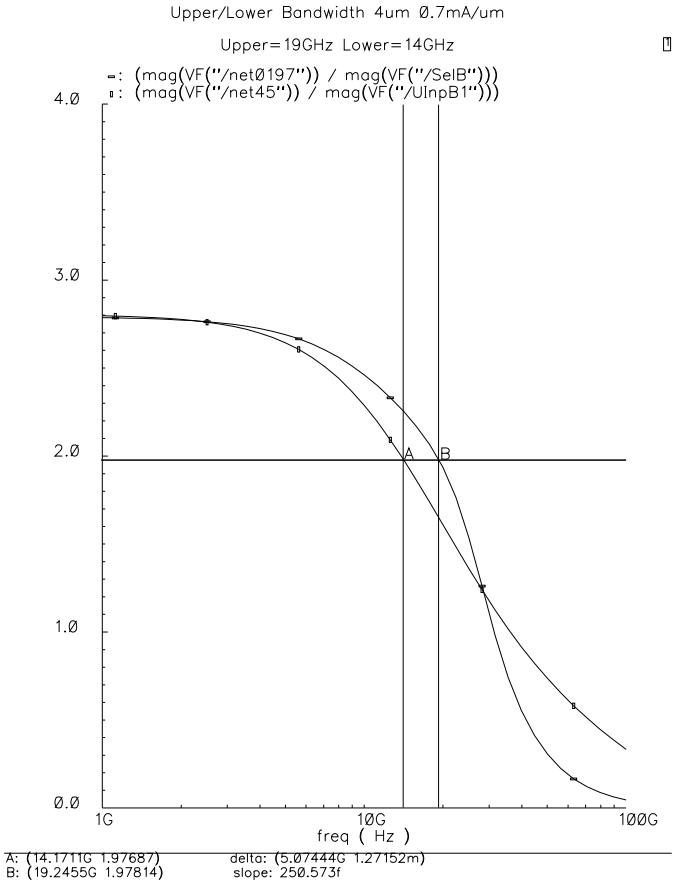

Figure 9.5: Small Signal Multiplexer Bandwidth.

Typical small signal bandwidth analysis can be misleading under these circumstances, (see Figure 9.5. This shows that the small signal bandwidth of both the upper and lower inputs is not strikingly different. however, as will be shown later the inputs behave differently in the large signal sense. For example, while the Data(Upper) Input to Output small signal bandwidth gives useful information regarding signal transmission while the Select input is fixed, the same cannot be said

for the Select(Lower) Input to Output small signal bandwidth while the Data Inputs are fixed. Both modes of signal propagation are important when the multiplexer is used in a serializer circuit, and both types of transitions need to happen inside a single bit time. This becomes the most critical consideration when the bit time has been reduced to the smallest possible values as it has in our circuits.

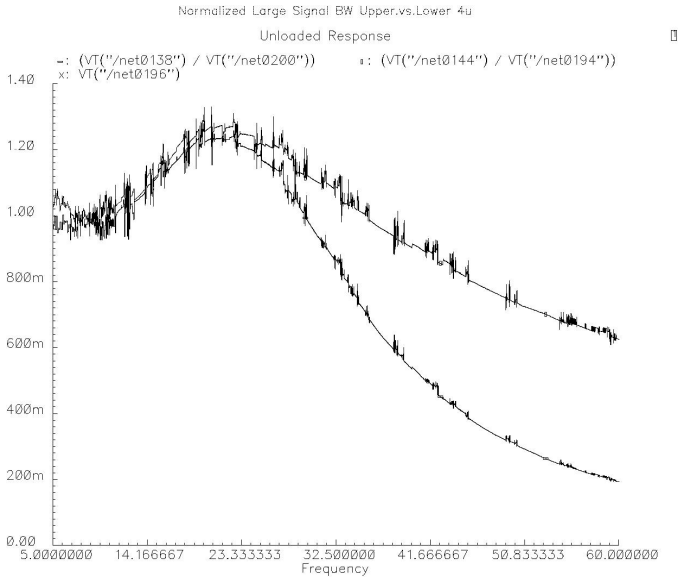

Figure 9.6: Normalized Large Signal Multiplexer Bandwidth.

Figure 9.6 shows the results of a large signal, (transient), simulation of a 2-to-1 multiplexer with the input frequency swept from 5GHz to 60GHz. To obtain a switching output for the select input, the two data inputs were fixed at complementary values. For the data-to-output measurement, the select signal was fixed, and the non-selected data input was set to a $V_{DIFF} = 0$ value. A simple aHDL module detected the amplitude at both inputs and outputs to provide an effective normalized output value for both select-to-output and data-to-output signals. The aHDL sampled the outputs to find the peaks, which effectively exhibits a degree of granularity which results in the "noisiness" seen in the plot. The upper trace in the figure indicates the effective bandwidth of the data-to-output signal is indeed larger and drops off less rapidly than that of the select-to-output signal, indicating it is indeed the bandwidth of the lower(select) input which is causing the earliest distortions.

9.5 Pass Through

As described in the Technology Appendix A, only a 120mV input difference at a differential pair is required to divert 99% of the tree current, which should prevent the disabled input from affecting the output. Unfortunately, at the high frequencies we are dealing with, this is simply not the case. Using very high frequencies, (large fractions of f_T), can cause unintended effects when attempting to use transistors as simple switching devices.

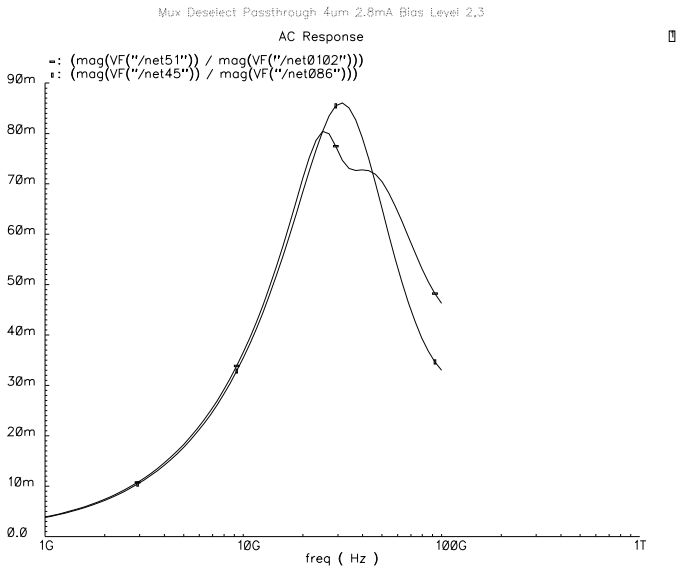

Figure 9.7: Multiplexer Signal Leakage. This plot shows the signal leakage from a disabled input pair to the output of a $4\mu m$ multiplexer. The difference is due to the different voltages present on the output lines which are driven by the enabled input.

In Figure 9.7, an AC plot of the normalized output magnitude of a $4\mu m$ multiplexer is shown when the input is **not** selected. As can be seen, at frequencies commonly present in 40Gbps signals, there is significant leakage through the multiplexer on the order of 10%. This occurs independently of the magnitude of the select signal, is worse in larger transistors, and it is due to the non-zero base-collector capacitances, (C_{bc}), of the uppermost transistors. At frequencies approaching f_T, the transistor junctions have reduced impedance. A C_{bc} of perhaps 5fF becomes only a 1.5KΩ impedance at 20GHz, and the multiplexer output signal can also affect the signal at the disabled input. Because of this, C_{bc} should be reduced as much as

possible, by using lower input voltages for the multiplexer. In the implementations of the ESMA in 5AM, the multiplexers use levels two and three rather than one and two, for the data and select inputs respectively. These are the lowest levels available. This reduced impedance of the junction capacitances also affects the bandwidth of the input which *is* selected, and reduces the output magnitude of the multiplexer at high frequencies as well. Note that because of the half peak f_T bias current, the effective f_T is lower than the optimal 50GHz.

Other factors limiting the simple multiplexer at high data rates include the collector-substrate capacitances, which are on the same order of magnitude as the base-collector capacitances. These C_{cs} capacitances also heavily degrade the performance of the current mirrors typically used in providing I_{BIAS} for the multiplexer, by reducing the effective output impedance of the source, rendering it little better, (or even worse), than a resistor at high frequencies.

9.6 Technology Comparison

The simple differential multiplexer can be placed in a feedback loop, (with optional emitter followers), to form the basic D-Latch used in CML circuits. This circuit forms the primary building block of the static clock divider. Because of this, the maximum performance of a multiplexer is related to that of the static divider in the same technology, barring loading considerations. In fact, with few exceptions, most "simple" dividers of this design type operate below $f_T \div 2$. Examples are [47], which uses 95GHz f_T devices to make a 50GHz divider multiplexer, [46] which uses 120GHz f_T devices to achieve a 35GHz clock divider, and [48] which attains a 97GHz static divider rate using 210GHz f_T devices.

In the cases where where static dividers exceed $f_T \div 2$, it is often accomplished through means not applicable to generic multiplexers. For example, in [49], current sources were added to the data input differential pairs so that they would retain some current flow when they were not selected. This allowed the de-selected input pair to more greatly affect the output, but not enough to upset the latch state. In [42] a static divider technique is described where the master and slave share two different current sources, with the "keeper" transistors in both latches making use of

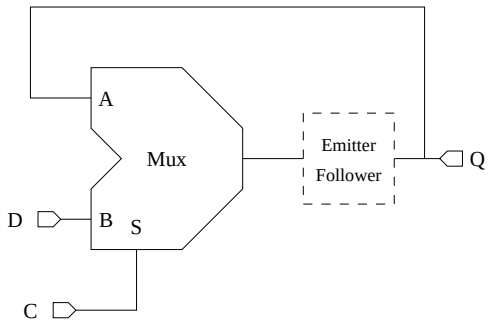

Figure 9.8: Multiplexer as part of a D-Latch. A D-Latch can be considered a multiplexer with feedback using optional emitter followers.

the smaller valued of the two. These techniques cannot be applied to a symmetrical 2-to-1 multiplexer.

There are other techniques which serve to increase the working frequency of a divider or a multiplexer. One such method is to use an even lower signal input voltage, thereby decreasing C_{bc} even further. Because of biasing concerns, the supply voltage must be increased in magnitude when lower input signal levels are used. This would be extremely wasteful of power if only a single supply were to be used and the circuit was part of a larger functional block, such as our monolithic transmitters. In some processes, such as 7HP, or 8HP, this cannot be done due to the low V_{CE} breakdown voltage. This technique is used in [37], [38] to achieve greater than $f_T \div 2$ performance. In addition, these circuits actually near 66% of f_T performance using microstrip lines as inductors to offset the capacitance of the RC delay at the top of the trees. In my opinion, this blurs the distinction between a static and dynamic divider as it is highly tuned for one frequency. One could also argue that this is similar to an injection-locked half-rate oscillator.

9.7 Bandwidth Improvements

There have been several papers which have claimed improvements in multiplexer bandwidth by the use of inductive peaking. In [43], peaking is used to improve the bandwidth of the simple multiplexer which is made up of 100GHz f_T devices. This technique at first glance seems to offer a huge benefit in terms of performance, but the actual utility may prove to be somewhat less.

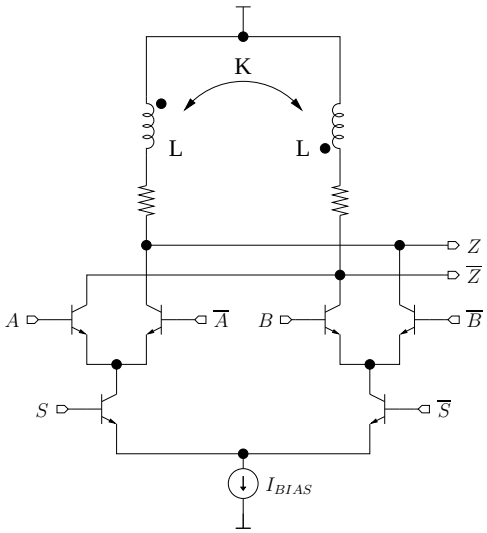

Figure 9.9: Generic inductive peaking circuit added to a standard 2-to-1 multiplexer.

Figure 9.9 is a schematic showing the addition of on-chip inductors to the standard multiplexer. The same circuit can serve to simulate both inductive peaking, $(K = 0)$, and cross coupled-transformer peaking, $(K \neq 0)$.

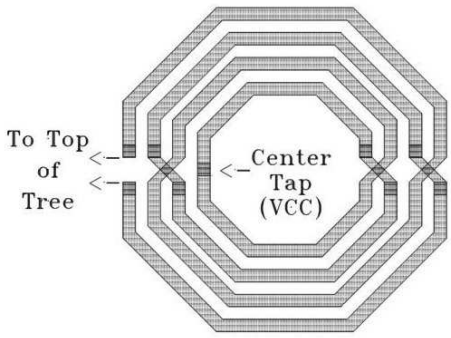

Figure 9.10: Transformer Layout

Transformer peaking makes use of processes with multiple thick metal layers to create symmetric interleaved inductors like that shown in Figure 9.10. The center tap can be connected to the upper supply rail, V_{CC}. The cross-coupling allows the decrease of current in one output branch to aid the increase in the other, giving rise to greater symmetry in output rise and fall times, as well as increasing the effective value of the peaking inductance.

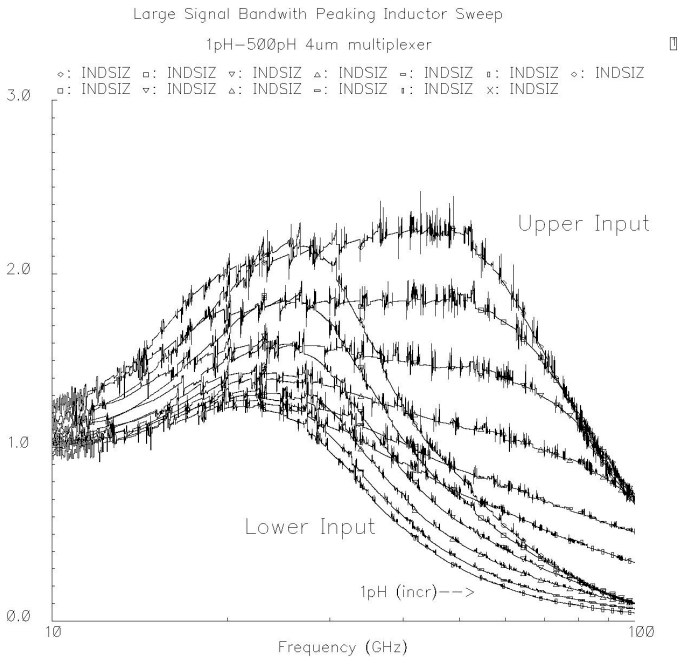

Figure 9.11: Large Signal Inductive Peaking Sweep. The inductors at the top of the circuit are swept from 1pH to 500pH.

As can be seen in Figure 9.11, inductive peaking greatly enhances the bandwidth of the Upper(data) input to the multiplexer. However, the Lower(select) input gains much less, and the response falls off much more rapidly. High speed operation requires that both inputs be capable of large signal switching in less than a bit time.

As shown in Figure 9.12, the use of cross-coupled transformers also does not radically improve the the response of the lower input. It mostly allows the use of smaller inductors while retaining the same peaking effect.

An additional means of increasing small-signal bandwidth in multiplexer circuits is through the use of negative feedback. Emitter degenerating resistors can linearize a differential pair as is described in Appendix A, (see Figure A.21). This reduces the gain, while the gain-bandwidth product (GBP) remains constant. As before, this technique works well on the upper input differential pairs, but does not have the desired effect on the lower select inputs. The reduced gain actually slows the current steering, although the magnitude of the steering is more *uniform* over a wider range. The addition of bypass capacitors across the R_E resistors can

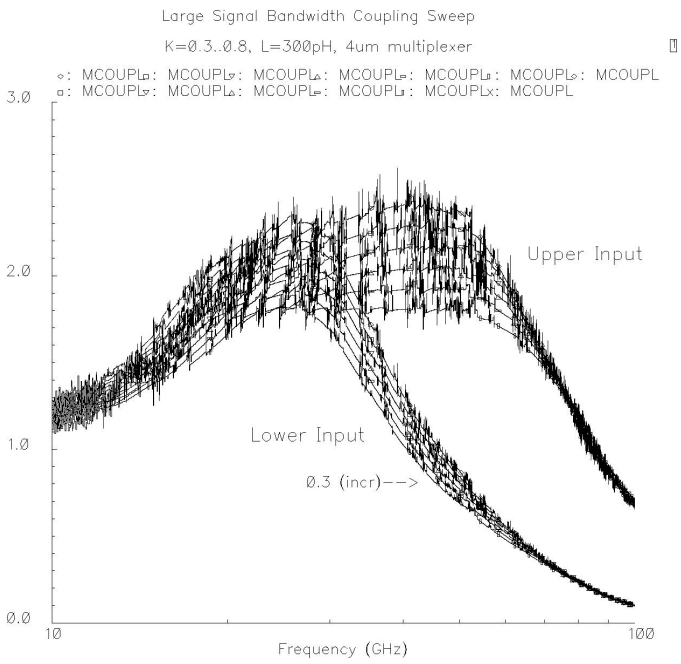

Figure 9.12: Large Signal Inductive Coupling Sweep. With the inductive value fixed at 300pH, the coupling coefficient K is swept from 0.3 to 0.8

increase the bandwidth even further, as at higher frequencies the resistors will be bypassed. Unfortunately, the addition of these components does not increase the large signal switching speed beyond the that of the original circuit without them. The additional parasitic effects of the components do more harm than good when the switching rate is required to be the maximum possible. However, when the multiplexer is used at lower data rates, these techniques can definitely improve the output signal quality. For example, linearization would greatly reduce the switching noise present in Figure 9.3.

9.8 Eye Comparisons

In comparing output eye diagrams between my monolithic serializer and multiplexers in the literature, allowances must be made for several factors. The multiplexers cited in the literature mainly consist of only that circuit, which leaves the testers with several degrees of freedom in applying test signals and supply voltages. For example, small phase adjustments between the data and clock can be made

externally. The signal levels and magnitudes can be adjusted independently of each other and of the power supply voltages.

The monolithic serializer chip contains its own clock and data sources which use the same process technology as the final multiplexers, and as such, they provide less than ideal inputs. Because the latches in the circuit are effectively the same as multiplexers, they have the same limitations with regard to switching noise as does the final output multiplexers. Often multiplexers in the literature are tested using commercial signal generation equipment which may use higher performing technologies to produce a superior input. For example, there are commercial 10Gbps...60Gbps test pattern generators which presumably make use of 200GHz or greater InP devices.

Lastly, the final output of the SERDES III is the result of a cascade of three separate levels of 2-to-1 multiplexer stages. If the output of the multiplexer in [43] were to be fed back into another identical multiplexer, for three stages, it is likely the output would not appear to be very much better than that of Figure 8.17, despite the fact that the devices in [43] have twice the peak f_T of the devices used in this work.

CHAPTER 10
Conclusion

As part of a group design effort, two complete serializer/deserializer chip designs were successfully fabricated and tested. This researcher contributed a great deal of the digital portions of those designs, including LFSRs, clock recovery circuits, state machines, multiplexers and demultiplexers. In those designs circuits operating at clock rates of up to 5GHz were created, which yielded serial data rates of approximately 20Gbps. Several problems with the prior designs were identified using simulations, and corrections and improvements were proposed.

For the next generation of chips a completely new 4-to-1 output multiplexer architecture was created, (the ESMA), which performs better than the types used in any of our previous chips. This multiplexer operates at data rates which were fast enough to require a completely new 16-to-4 multiphase multiplexer architecture in order to supply it with data. Also developed was an LFSR which serves as both a parallel shift register and a random number generator by converting a traditional design into a multiple phase one. The novel LFSR architecture which was used in the final transmitter design was generated using a program to transform regular state machines into an "advance-by-N" form. This researcher designed and fabricated a 10GHz 5DM-based ring oscillator test chip, which performed as designed. The final 5AM transmitter chip was fabricated by MOSIS under their educational program. At that time, the 5DM process was not available and so a lower-performing, but already available 5AM current starving VCO was used instead. The chip was received and its operation was verified up to 27Gbps, at which time the lower performance of the current steering VCO precluded testing individual components to failure. Regardless, the architecture has been shown to be successful and is an improvement on previous designs.

In order to better characterize the performance of the newer designs, an extra pair of chips were designed and submitted to MOSIS with improved frequency range, better load balancing, and more outputs for test signals. At the higher rates, output

quality is expected to further degrade, however each of the subsystems can be better individually evaluated. These chips will not be ready for several more months, and the other members of the research group will be involved in their testing. They should hopefully yield the highest data rate serial transmitter system ever produced, relative to the process f_T.

10.1 Future Work

There are several different methods which may still prove fruitful in increasing the performance of SiGe HBT Serializer circuits. There is no hard line which can be drawn indicating that nothing more can be achieved. Notably, the use of peaking inductors or cross-coupled transformers may slightly improve 2-to-1 multiplexer bandwidths to allow slightly better output eye quality for a given data rate. Distributed designs may be developed which allow common collector loading effects to be lessened, perhaps leading to a distributed multiplexer architecture.

Power consumption in this design is higher than in necessary, and can doubtless be reduced by converting some parts of the design to CMOS, yielding a truly BiCMOS design. External analog signals and extra supply voltages might serve to fine-tune the delay elements built into the ESMA in order to minimize the switching distortion.

LITERATURE CITED

[1] D. C. Ahlgren, G. Freeman, S. Subbanna, R. Groves, D. Greenberg, J. Malinowski, D. Nguyen-Ngoc, S. J. Jeng, K. Stein, K. Schonenberg, D. Kiesling, B. Martin, S. Wu, D. L. Harame, and B. Meyerson, "A SiGe HBT BiCMOS Technology for Mixed Signal RF Applications," BCTM Tech. Dig., pp 195-197 12.3, 1997. 8, 150

[2] R. C. Walker, K. Hsieh, T. A. Knotts, and C. Yen, "A 10 Gb/s Si-Bipolar TX/RX Chipset for Computer Data Transmission," IEEE International Solid-State Circuits Conference, pp. 302-303, 1998. 13, 16, 195

[3] J. Scheytt, G. Hanke, U. Langmann, "A 0.155-, 0.622-, and 2.488-Gb/s Automatic Bit-Rate Selecting Clock and Data Recovery IC for Bit-Rate Transparent SDH Systems," IEEE Journal of Solid-State Circuits, vol. 34, no. 12, pp. 1935-1943, December 1999. 13

[4] D. Friedman, M. Meghelli, B. Parker, H. Ainspan, and M. Soyuer, "Sub-picosecond SiGe BiCMOS Transmit and Receive PLLs for 12.5 Gbaud Serial Data Communication," Symposium on VLSI Circuits, pp. 132-135, 2000. 13

[5] Z. Wang, M. Berroth, J. Seibel, P.Hofmann, A.Hulsmann, K. Kohler, B. Raynor, J. Schneider, "19GHz Monolithic Integrated Clock Recovery Tuning PLL and 0.3um Gate-Length Quantum-Well HEMTs," 1994 IEEE Int. Solid-State Circuits Conf. (ISSCC), pp. 118-119, February 1994. 13

[6] A. Felder, M. Möller, J. Popp, J Böck, H.M. Rein, "46 Gb/s DEMUX, 50Gb/s MUX, and 30 GHz Static Frequency Divider in Silicon Bipolar Technology," IEEE Journal of Solid-State Circuits, vol. 31, no. 4, pp. 481-486, April 1996. 13

[7] H. Knapp, T. Meister, M. Wurzer, D. Zoschg, K. Aufinger, L. Treitinger, "A 79GHz Dynamic Frequency Divider in SiGe Bipolar Technology," IEEE Int. Solid-State Circuits Conf. (ISSCC), Feb. 2000, Dig. Tech. Pap., pp. 208-209.

[8] M. Wurzer, J. Bock, W. Zirwas, H. Knapp, F. Schumann, A. Felder, L. Treitinger, "40 Gb/s Integrated Clock and Data Recovery Circuit in a Silicon Bipolar Technology," IEEE BCTM, 1998, pp. 136-139.

[9] T. Masuda, K. Ohhata, E. Ohue, K. Oda, M. Tanabe, H. Shimamoto, T. Onai, K. Washio, "40Gb/s Analog IC Chipset for Optical Receiver using SiGe

HBTs," IEEE Int. Solid-State Circuits Conf(ISSCC), Feb. 1998, Dig. Tech. Pap., pp. 314-315.

[10] G. Freeman, M. Meghelli, Y. Kwark, S. Zier, A. Rylyakov, M. Sorna, T. Tanji, O. Schreiber, K.Walter, J. Rieh, B. Jagannathan, A. Joseph, S. Subbanna, "40-Gb/s Circuits Built From a 120-GHz fT SiGe Technology," IEEE Journal of Solid-State Circuits, vol. 37, no. 9, pp. 1106-1114, September 2002.

[11] G. Georgiou, Y. Beyens, Y. Chen, A. Gnauck, C. Gropper, P. Paschke, R. Pullela, M. Reinhold, C. Dorschky, J. Mattia, T. Mohrenfels, C. Schulien, "Clock and Data Recovery IC for 40-Gb/s Fiber-Optic Receiver," IEEE Journal of Solid-State Circuits, vol. 37, no. 9, pp. 1120-1125, September 2002.

[12] M. Reinhold, C. Dorschky, E. Rose, R. Pullela, P. Mayer, F. Kunz, Y. Baeyens, T. Link, J. Mattia, "A Fully Integrated 40-Gb/s Clock and Data Recovery IC With 1:4 DEMUX in SiGe Technology," IEEE Journal of Solid-State Circuits, vol. 36, no. 12, pp. 1937-1945, December 2001.

[13] H. Kroemer, "Heterostructure Bipolar Transistors and Integrated Circuits," Proceedings of the IEEE, vol. 2, no. 1, pp. 13-25, January 1982. 13 13 14 14 14 14 13

[14] C. McAndrew, J. Seitchik, D. Bowers, M. Dunn, M. Foisy, I. Getreu, M. McSwain, S. Moinian, J. Parker, D. Roulston, M. Schroter, P. Wijnen, L. Wagner, "VBIC95, The Vertical Bipolar Inter-Company Model," IEEE Journal of Solid-State Circuits, vol. 31, no. 10, pp. 1476-1483, October 1996.

[15] Jan M. Rabaey, Digital Integrated Circuits a Design Perspective., Englewood Cliffs, NJ:, Prentice Hall, 1996.

[16] R. L. Treadway, "DC Analysis of Current Mode Logic", IEEE Circuits and Devices Magazine, Volume: 5 Issue: 2, March 1989.

[17] David A. Hodges, Horace G. Jackson, Analysis and Design of Integrated Circuits., New York, NY: McGraw-Hill, Inc, 1988.

[18] Dan H. Wolaver, Phase-Locked Loop Circuit Design, Englewood Cliffs, NJ: Prentice Hall, 1991.

[19] T. H. Lee, and A. Hajimiri, "Oscillator Phase Noise: A Tutorial," IEEE Journal of Solid-State Circuits, vol. 35, no. 3, pp. 326-335, March 2000.

[20] S. A. Steidl, "A 32-Word by 32-Bit Three-Port Bipolar Register File Implemented Using a SiGe HBT BiCMOS Technology," Candidacy document, Rensselaer Polytechnic Institute, Department of Electrical Engineering, May 1999.

[21] S. Finocchiaro, G. Palmisano, R. Salerno, C. Sclafani, "Design of Bipolar RF Ring Oscillators," IEEE International Conference on Electronics, Circuits and Systems, vol. 1, pp. 5-8, September 1999. 52, 151, 152 158 158, 162 162 189 193, 200, 201 194 55, 195

[22] A. W. Buchwald, and K. W. Martin, "High-speed voltage-controlled oscillator with quadrature outputs," Electronics Letters, vol. 27, no. 4, pp. 309-310, February 1991.

[23] H. Matsuoka, and T. Tsukahara, "A 5-GHz Frequency-Doubling Quadrature Modulator with a Ring-Type Local Oscillator," IEEE Journal of Solid-State Circuits, vol. 34, pp. 1345-1348, September 1999.

[24] P. Kinget, R. Melville, D. Long, V.Gopinathan, "An Injection-Locking Scheme for Precision Quadrature Generation," IEEE Journal of Solid-State Circuits, vol. 37, no. 7, pp 845-851, July 2002.

[25] S. Lee, B. Kim, and K. Lee, "A Novel High-Speed Ring Oscillator for Multiphase Clock Generation Using Negative Skewed Delay Scheme," IEEE Journal of Solid-State Circuits, vol. 32, no. 2, pp. 1451-1454, February 1997.

[26] S. K. Enam and A. A. Abidi, "A 300-MHz Voltage-Controlled Ring Oscillator," IEEE Journal of Solid-State Circuits, vol. 25, no. 1, pp. 312-315, February 1990.

[27] L. Sun, T. Kwasniewski, and K. Iniewski, "A Quadrature Output Controlled Ring Oscillator Based on Three-Stage sub-feedback Loops," IEEE Internation Symposium on Circuits and Systems, vol. 2, pp 176-179, June 1999.

[28] J. A. McNeil, "Jitter in Ring Oscillators," IEEE Journal of Solid-State Circuits, vol. 32, pp. 870-879, June 1997.

[29] G. Niu, Z.Jin, J. Cressler, R. Rapeta, A. Joseph, D. Harame, "Transistor Noise in SiGe RF Technology," IEEE Journal of Solid-State Circuits, vol. 36, no. 9, pp. 1424- 1427, September 2001.

[30] T. W. Krawczyk, "Circuits for the Design of a Serial Communication System Utilizing SiGe HBT BiCMOS Technology," PhD Thesis, Rensselaer Polytechnic Institute, Department of Electrical Engineering, Nov 2000. 25, 195 195 195 195 195 195 198, 200 200 36, 56

[31] Rensselaer Polytechnic Institute. (2003). Symmetric multiplexer. US Patent, us 6,531,910. 76

[32] T. W. Krawczyk, P. F. Curran, M.W. Ernest, S.A. Steidl, S. R. Carlough, J. F. McDonald, R. P. Kraft, "SiGe HBT serial transmitter architecture for high speed variable bit rate intercomputer networking," IEE Proc-Circuits Devices Syst, vol. 151, no. 4, Aug 2004. 11

[33] R. Singh, Y. V. Tretiakov, J. B. Johnson, S. L. Sweeny, R. L. Barry, M. Kumar, M. Erturk, J. Katzenstein, C. E. Dickey, and D. L. Harame, "Parasitic Modeling and Noise Mitigation in Advanced RF/Mixed-Signal Silicon Germanium Processes," IEEE Transactions on Electron Devices, vol. 50, no. 3, March 2003. 157

[34] A. Barna, "Rise time optimization of high-speed digital fanout circuits," IEEE Journal of Solid-State Circuits, vol. 4, no. 3, pp. 159-161, June 1969. 169

[35] A. Barna, "On the transient response of emitter followers," IEEE Journal of Solid-State Circuits, vol. 8, no. 3, pp. 233-235, June 1973. 169

[36] A. Barna, "Rise Time of Emitter-Coupled Logic Circuits Including the Effects of Collector-to-Base Capacitances," IEEE Journal of Solid-State Circuits, vol. 8, no. 4, pp. 284-286, August 1973. 169

[37] M. Wurzer, T.F. Meister, H. Schäfer, H. Knapp, J. Böck, R. Stengl, K. Aufinger, M. Franosch, M. Rest, M. Möller, H. -M Rein, A. Felder, "42 GHz Static Frequency Divider in a Si/SiGe Bipolar Technology," IEEE International Solid-State Circuits Conference, 1997 Digest pp. 122-123. 138

[38] M. Wurzer, T.F. Meister, H. Knapp, K. Aufinger, R. Schreiter, S. Boguth, L. Treitinger, "53 GHz static frequency divider in a Si/SiGe Bipolar Technology." IEEE International Solid-State Circuits Conference, February 2000 Digest pp. 206-207. 138

[39] C-H Wu, C-C Tang, and S-I Liu, "Analysis of On-Chip Spiral Inductors Using the Distributed Capacitance Model," IEEE Journal of Solid-State Circuits, vol. 38, no. 6, pp. 1040-1044, June 2003. 157

[40] E. Hegazi, and A. A. Abidi, "Varactor Characteristics, Oscillator Tuning Curves, and AM-FM Conversion," IEEE Journal of Solid-State Circuits, vol. 38, no. 6, pp. 1033-1039, June 2003. 69

[41] J. D. Cressler, G. Niu, Silion-Germanium Heterojunction Bipolar Transistors, Norwood, MA: Artech House, Inc., 2003. 13, 150, 153

[42] B. Razavi, Design of Integrated Circuits for Optical Communications, McGraw-Hill, 2003. 104, 137

[43] D. Kehrer, H.D. Wohlmuth, H. Knapp, M. Wurzer, and A.L. Scholtz, "40-GB/s 2:1 Multiplexer and 1:2 Demultiplexer in 120-nm Standard CMOS," IEEE Journal of Solid-State Circuits, vol. 38, no. 11, pp. 1830-1837, November 2003. 14, 129, 138, 142

[44] J. Lee, B. Razavi, "A 40-GHz Frequency Divider in 0.18-μm CMOS Technology," IEEE Journal of Solid-State Circuits, vol. 39, no. 4, pp. 594-601, April 2004. 14

[45] M. Meghelli, "A 132-Gb/s 4:1 Multiplexer in 0.13-μm SiGe-Bipolar Technology," IEEE Journal of Solid-State Circuits, vol. 39, no. 12, pp. 2403-2407, December 2004. 14, 129

[46] M. Meghelli,A.V.Rylyakov, Lei Shan, "50-Gb/s SiGe BiCMOS 4:1 Multiplexer and 1:4 Demultiplexer for Serial Communication Systems," IEEE Journal of Solid-State Circuits, vol. 37, no. 12, pp. 1790-1794, December 2002. 137

[47] K. Washio, E. Ohue, K. Oda, R. Hayami, M. Tanabe, H. Shimamoto, "A 50-GHz static frequency divider and 40-Gb/s MUX/DEMUX using self-aligned selective-epitaxial-growth SiGe HBTs with 8-ps ECL," IEEE Transactions on Electron Devices, vol. 48, no. 7, pp. 1482-1487, July 2001. 137

[48] A. Rylyakov, T. Zwick, "96-GHz Static Frequency Divider in SiGe Bipolar Technology," IEEE Journal of Solid-State Circuits, vol. 39, no. 10, pp. 1712-1715, October 2004. 137

[49] T. E. Collins, V. Manan, and S. I. Long, "Design Analysis and Circuit Enhancements for High-speed Bipolar Flip-Flops," IEEE Journal of Solid-State Circuits, vol. 40, no. 5, pp. 1166-1174, May 2005. 129, 137

[50] H. -M. Rein, M. Möller, "Design Considerations for Very-High-Speed Si-Bipolar IC's Operating up to 50 Gb/s," IEEE Journal of Solid-State Circuits, vol. 31, no. 8, pp. 1076-1090, August 1996. 169

[51] D. A. Huffman, "A Method for The Construction of Minimum Redundancy Codes," Proceedings of IRE, 40(9):pp. 1098-1101, September 1952. 8

[52] A. X. Widmer, P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," IBM Journal of Research and Development, vol. 27, no. 5, pp 440-451, 1983. 8

APPENDIX A
Technology

A.1 Initial Choices

Our decisions in choosing the process and logic for our project were heavily influenced by both ambition and practicality. Having been given early access to IBM's advanced SiGe HBT process made the process choice clear. The only other options would involve expensive low yield compound semiconductor processes. Industry seems to be embracing SiGe because of its obvious benefits. Work by other members of our research group in the FRISC project influenced the decision to go with fully differential CML logic with optional emitter followers.

A.2 IBM SiGe 5HP BiCMOS Process

IBM's SiGe 5HP BiCMOS process allows the use of $0.5\mu m\ x\ 1.0\mu m$ emitter HBTs and $0.35\mu m\ L_{eff}$ CMOS on the same die [1]. It is a full featured process with numerous types of resistors, and includes components like varactors, capacitors, inductors, and diodes. Our logic designs primarily use the npn HBT, and the polysilicon resistor.

A.2.1 The SiGe HBT

IBM's SiGe HBT differs from a normal BJT in that the base contains germanium. The base is epitaxially grown with a non-uniform doping profile. The introduction of germanium into the transistor base results in a reduction of the base bandgap, increasing emitter injection efficiency [41] and reducing base transit time. The improvement in injection efficiency would normally result in improved β, but this is traded off for reduced base resistance by increasing the base doping. This increases the maximum operating frequency of the transistor, (see Equation A.2.)

In Figure A.1 above, the general layout of the npn HBT is shown. The layout cell in the design kit is parameterized by the emitter length (L), while the width (W) is fixed at $0.5\mu m$, as indicated in the top view. Only discrete values of length

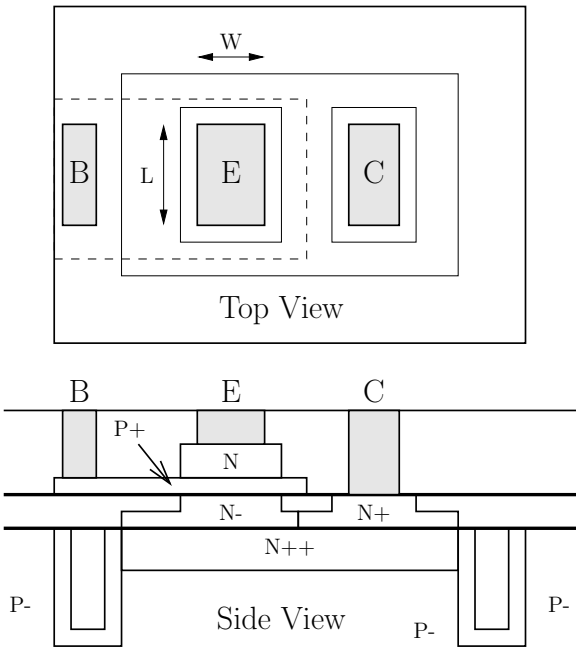

Figure A.1: Layout of the 5HP SiGe HBT showing top and side views.

are permitted by the kit, with $1\mu m$ being the smallest, and $20\mu m$ being the largest, with $2.5\mu m$ being the only non-integral value allowed. The emitter and collector areas, (as visible from the top), are surrounded by a deep trench isolation moat that serves to define the N++ subcollector, the collector itself, and the subcollector connection. The P+ SiGe base is epitaxially grown. There are two types of HBTs available, npn, and npnhb. The npnhb is a high breakdown version that has a different collector doping in order to allow higher collector-emitter voltages, up to 4.5V.

The IBM design kits contain subcircuit models for the HBT which contain extrinsic resistances, a VBIC(vertical bipolar intercompany model) [14], and an optional current source to degrade performance appropriately when a high breakdown HBT is instantiated. The design kit is optimized for Cadence's Spectre circuit simulator that includes a VBIC model with parameters for self heating and impact ionization. The VBIC model schematic is shown in Figure A.2.

We have had access to the design kit through several iterations of model refinement, so over time the device model parameters have varied by several percent. Presented here are some device characteristic curves produced by simulation using

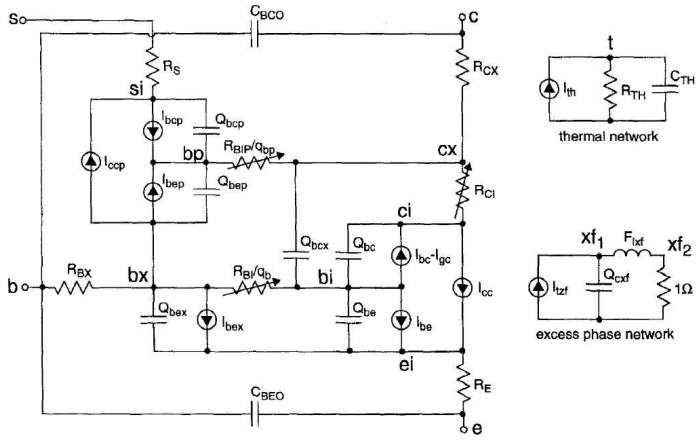

Figure A.2: Spectre VBIC Model [14].

the design kit.

Figure A.3 is a forward active Gummel plot of a $2.5\mu m$ npn showing β to be approximately 100 at 70C with V_{CE}=0V. The β varies only slightly with transistor size, and the high breakdown device with a designation of "npnhb" shows a β of approximately 90. As temperatures go up, so do the absolute currents, but the ratio stays approximately the same as expected.

Forward active curves for a typical npn HBT are shown in Figure A.4. The early voltage, ($V_A > 120$V), gives nearly flat curves from 0.5V to 2.8V of V_{CE}. Note that the early voltage parameter, **vef**, used in the VBIC model is actually half the value normally specified as the Early voltage.

An important device parameter for very high-speed digital circuit operation is f_T, defined as the frequency at which the transistor exhibits unity current gain with the output shorted. It is related to the forward transit time at high frequencies, or equivalently by the relation in Equation A.1.

$$f_T = \frac{g_m}{2\pi(C_{be} + C_{bc})} = \frac{1}{2\pi\tau_f} \tag{A.1}$$

This value is a fairly strong function of bias current through g_m, and in order to make ensure maximum performance of our circuits, we want the bias current to coincide with the peak f_T. Understanding f_T is crucial to designing high speed circuits. Essentially at higher frequencies, the voltage dependent junction capacitances

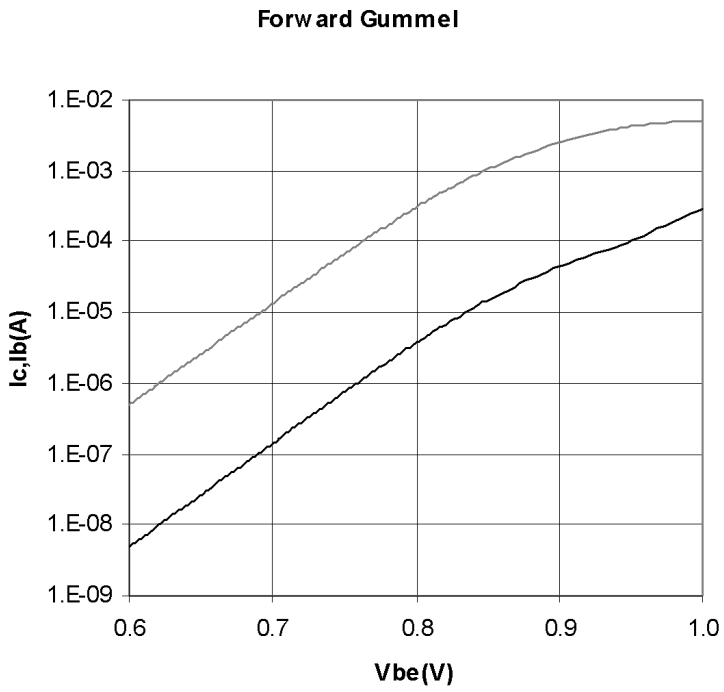

Figure A.3: Gummel plot of 2.5μm npn HBT at 70C.

partially short out the gain of the g_m controlled source, reducing the gain to unity at f_T. Because the output of the transistor is shorted when f_T is being characterized, there is no voltage gain and the Miller effect does not change the effective value of C_{bc}. It should be noted that f_T may not represent a real physical cutoff, but rather it is interpreted from the earlier 20db/dec roll-off [41]. In reality, other effects such as substrate capacitance C_{CS} may limit performance even earlier.

Operating digital circuits at or near f_T is problematic because these capacitances are large enough to pass a significant portion of a signal even when the transistor is in cut-off, if the signal has high frequency content. For example, a multiplexer will still transfer part of the non-selected input to the output, (see Chapter 9).

As can be seen in Figure A.5, larger transistors have a larger f_T, but also sink a larger amount of current. It is also important in that larger transistors will also induce a larger load on the circuits driving them due to the larger base currents, as well as having larger parasitic capacitance effects. From the above chart an

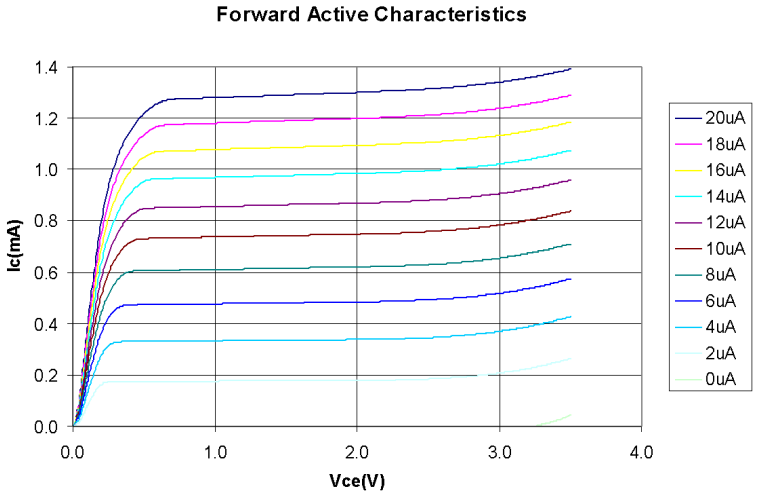

Figure A.4: Forward active characteristics of the 2.5μm npn HBT at 25C.

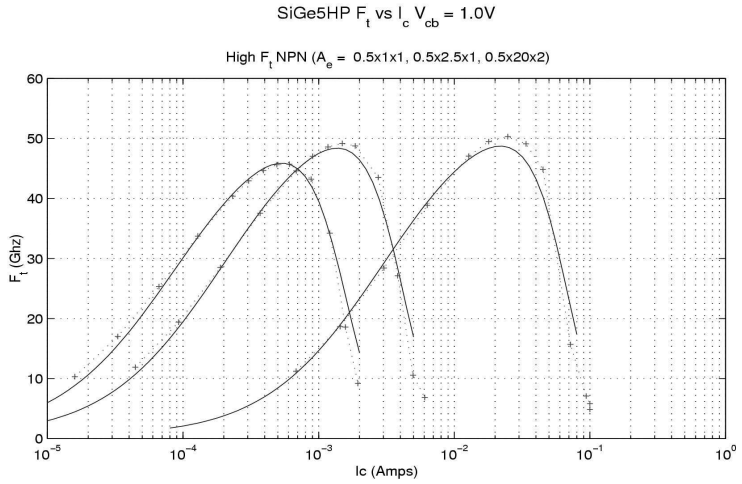

Figure A.5: Plots from IBM showing simulated f_T as a function of collector current bias for three different sized HBTs at 25C.

approximate maximum f_T value occurs around I_C=0.7mA/μm for all transistors. Referring back to Figure A.3, we can see that in our desired I_C operating range, the forward V_{BE} drop is between 0.85 and 0.9 volts. This is what sets the separation between our voltage levels, as is discussed later.

Another important parameter is f_{max}, the maximum oscillation frequency. This is defined as the unity power gain frequency. It is related to f_T through the relation:

$$f_{MAX}{}^2 = \frac{f_T}{8\pi R_B C_{BC}} \tag{A.2}$$

In the above equation, R_B is the base resistance and C_{BC} is the base-collector junction capacitance. This parameter determines maximum bandwidth, which is important for the analog sections of our designs. According to the documentation, the peak f_T for the minimum sized device is 47GHz, while f_{max} is specified as 65GHz at 75C.

A.3 Resistors

IBM's technology offers a variety of resistors with different sheet resistivities which range from 8.1Ω/□ to 1750Ω/□. Some of the other characteristics considered in choosing a type of resistor are the temperature coefficient(s), the parasitic capacitances, and the tolerances. The calculation of the actual tolerance for a particular resistor is quite complicated as the resistor instance is dependent on multiple mask and material tolerances, and the total resistance is dependent on a complicated geometry consisting of the primary resistive material along with contacts, vias, and other initial wiring. The polysilicon over deep trench resistor, (pbdtres), was chosen for general use because of its relative temperature insensitivity, low parasitic capacitance, and good (relative) tolerance. It features 220Ω/□, an almost zero temperature coefficient, and a low capacitance per unit area of 0.0667fF/μm^2.

In resistor layout, the length over width ratio determines the actual resistance. Ideally the resistor layout could be made as small as desired, but process variations will cause a relatively larger effect on a smaller physical resistor layout than on a larger one. When a larger resistivity material is used, this effect is increased. A spreadsheet was used to calculate tolerances of pbdtres resistors for particular resistance ranges used in our designs. A minimum layout dimension of $4\mu m$ for both the width and length of the polysilicon was found to be adequate to achieve tolerances on the order of 10%, with little incremental improvement beyond that. In the actual layout tool, one can specify a desired total resistance and the polysilicon width. To ensure a proper resistor, the width parameter should be increased until the calculated length is $4\mu m$ or greater. The current passing through the resistor must

also be considered when choosing the geometry as the maximum current density is specified as 0.6mA/μm. As this is nearly equal to the maximum f_T current parameter ratio for the HBT, the resistor width for circuits scales linearly with the transistor width for those circuits with transistors 4μm or greater.

A.4 Interconnect

IBM fabricates 5HP with several metalization options. Up to five metal layers may be used and there is an option, (5AM), available for creating a thicker last metal and insulating layer which is useful for creating higher quality inductors and low resistance wires. The 5AM option replaces the last metal layer(LM) and the one below that(M2,M3, or M4) with both the MT and AM(topmost) layers. A recently introduced option is 5DM, which is quite similar to 5AM except that the layer below AM is copper. The kit documentation for each option gives complete design criteria for creating reliable layouts. For a wire with a particular width and at a specific level the maximum safe DC and RMS currents can be calculated using the equations therein. Vias are also similarly covered. Most of the wiring in our designs falls into two broad categories, high current, and signal. The high current lines occur primarily inside gates and usually either carry a constant current or switch on and off with currents of at least 0.6mA. In some cases the currents go as high as tens of milliamps.

Table A.1: Metal Interconnect Characteristics

Layer	Metal Thicknesses	Oxide Below	Min. Width	Max(mA) at MinWidth	Max mA/μm	$\Omega/\square$	TCR %/C
M1	0.63 ± 0.06	1.90 ± 0.2	0.8	0.39	0.49	.076	.37
M2,M3,M4	0.85 ± 0.08	1.2 ± 0.31	0.9	0.95	1.1	0.45	.39
LM(5HP)	2.07 ± 0.2	1.2 ± 0.31	2.4	8.17	3.4	0.15	.42
MT	0.83 ± 0.08	1.2 ± 0.31	0.9	0.95	1.1	.045	.35
E1(5DM)	3.0 ± 0.5	4.0 ± 0.5	1.5	26.1	17.4	.0063	0.33
AM(5AM)	4.0 ± 0.4	3.0 ± 0.5	4.0	23.4	5.9	.00725	.38
MA(5DM)	4.0 ± 0.4	4.0 ± 0.4	4.0	24.4	6.1	.00725	.38

In the Table A.1, TCR is %/C, (percentage change per degree C.) The values above are at 25C. Vias are generally 0.9μm square, and can handle 1.2mA each.

For large current situations, a rectangular "via bar" is allowed. The selection of the 5AM/5DM options greatly ease the layout of power rails as they can take up one third the size if placed on AM as opposed to LM. The E1 layer of 5DM can be used as an orthogonal power rail distribution network, and it is also very useful in making inductive elements.

A.4.0.1 Capacitors

There are two types of regular capacitors, metal-insulator-metal(MIM), and decoupling caps which are substrate junction based. The MIM has fewer parasitics and is used whenever a capacitor is needed other than for power supply decoupling. The basic MIM capacitance per unit area is 1.35fF/μm^2, with a maximum value as high as 7pF. The insulator is an extra nitride layer which has a frequency dependence which causes the capacitance to decrease drastically above 10GHz. In contrast, the dcap has 1.5fF/μm^2, but is much more non-linear, and the capacitance increases beyond 10GHz.

In addition to the two fixed-value capacitors, a varactor diode with 1.37fF/μm^2 at 0V can be used. This varactor has a minimum available size of 6μm^2. Some of the 5HP versions also offer a MOS-based varactor with 3μm^2 which has a high roll-off closer to 2GHz.

A.4.0.2 Inductors

Inductors are available in the 0.6nH to 15.8nH range with Q-factors as high as 10, or between 2.8nH and 83nH with Q-factors as high as 19 if the 5AM or 5DM metalization option is chosen. I have experienced some difficulty using the IBM-kit supplied inductors in designs primarily due to the fact that they are modeled for small signal circuits rather than for digital ones. The inductance, Q, and other factors are functions of frequency [33].

Simple discrete π-models containing only a few elements are created by the IBM kits during netlist generation which are optimized for a specific operating frequency. More complex models include many more components [39] such as inter-turn capacitances. The digital signals employed in my designs contain large frequency ranges, and the use of the simple inductor models has returned somewhat suspect

results in transient mode simulations. An attempt was made to use $Sonnet^{TM}$ modeling software to create an n-port s-parameter table based model from the geometry of one of the parameterized inductor instances. Inexperience with Sonnet's 2.5D system made modeling the deep-trench mesh below the inductors an exercise in "guesstimation." In simulations, the newly generated model's performance diverged from the one supplied from IBM's tools, leading to a lack of confidence in the results.

A.5 Other Components

The technology offers a large number of additional devices. Many of these were not used in my designs. They include a gated lateral PNP(GLPNP) transistor, Schottky barrier diodes(SBD) with a 310mV drop and breakdown voltage of 5.5V, and a P-i-N diode with 15V breakdown.

A.6 Logic

Our digital circuits are implemented using differential current mode logic, (CML). This is essentially the same as differential ECL, but with the optional omission of emitter followers when feasible [15]. Analysis of this logic family draws upon [16], but is adapted for differential circuits as opposed to single-ended as is presented there. In a differential logic system, a pair of wires is used to indicate a boolean condition with the voltage difference between the pair indicating the state, (as opposed to the absolute voltage on a single wire.) The difference voltage, (V_{DIFF}), and common mode voltage, (V_{CM}), for a pair of wires (a_0,a_1) is shown below. Ideally, a differential logic system should be immune to variations in V_{CM} and the outputs should be solely a function of V_{DIFF}.

$$V_{DIFF} = a_0 - a_1 \tag{A.3}$$

$$V_{CM} = (a_0 + a_1)/2 \tag{A.4}$$

The individual voltage signals can also be expressed in terms of the difference and common mode voltages as:

$$a_0 = V_{CM} + V_{DIFF}/2 \tag{A.5}$$

$$a_1 = V_{CM} - V_{DIFF}/2 \tag{A.6}$$

Since all inputs and outputs are differential, there is no need for inverters since wires of a differential pair can simply be crossed to produce a complementary signal. The circuits mainly consist of current steering logic trees that use differential transistor pairs to divert current through one or another "branch" as it travels from V_{CC} to V_{EE}. Resistors at the "top" of the tree are used to generate voltage outputs that change when the current through a particular resistor's branch changes. These outputs swing through a voltage V_S, and drive other gates.

A.7 CML Gates

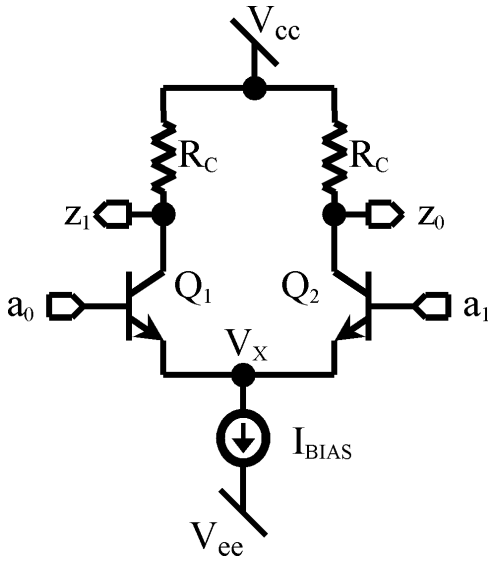

Figure A.6: Simple CML buffer/inverter.

In Figure A.6 you see a simple buffer/inverter element. The current source at the bottom ensures that a fixed amount of current I_{BIAS} flows through the tree at all times. In general, circuits that have different sized switching transistors will use a different value for I_{BIAS}, and the resistors at the top are chosen to provide a single design-wide uniform swing. The V_{CC} supply voltage is traditionally set at 0V, (which is the easiest to provide a stable value for), due to the direct effect this

supply has on the signal outputs. V_{EE} is set to some negative value dependent on considerations that will be covered later.

A desirable property of current steering logic is that the constant current per gate prevents or minimizes the problems associated with switching noise on the power supply lines. This leads to smaller allowable logic swings, and ultimately to a smaller power supply voltage. The two complementary differential inputs, a_0 and a_1, cause the current to flow in one of the two branches through the collector resistors R_C, affecting the corresponding outputs z_0 and z_1.

The transistors in these CML trees are meant to remain exclusively in the active mode, so simplifications can be made regarding modeling their collector currents. In the standard Ebers-Moll model, with α approximately equal to one and the reverse saturation current I_{CS} negligibly small, we can express the collector current as:

$$I_C = I_S(e^{V_{BE}/\Phi_T} - 1) \tag{A.7}$$

Note that the I_S above is actually I_{ES}. If we were to solve this for V_{BE}, it would be obvious that the "1" term is insignificant since collector current I_C is much greater than I_S. Accordingly we can say:

$$I_C = I_S e^{V_{BE}/\Phi_T} \tag{A.8}$$

Using this, the ratio of the collector currents in the two branches are exponentially related to the difference voltage as shown in the equation below. The common mode voltage V_{CM} isn't a factor, at least to a reasonable approximation.

$$\frac{I_{C1}}{I_{C2}} = \frac{I_S e^{(a_0 - V_x)/\Phi_T}}{I_S e^{(a_1 - V_x)/\Phi_T}} = e^{V_{DIFF}/\Phi_T} \tag{A.9}$$

Figure A.7 shows the percentage diversion as a function of the input difference voltage. A nearly complete swing is desirable due to the consideration that we wish to preserve and enhance the digital characteristics of signals. If we arbitrarily define "nearly complete" to mean 99%, the resulting current ratio in the equation above can be solved to give a minimum V_{DIFF} of only approximately 120mV, which is only

about 5 thermal voltages at 25C.

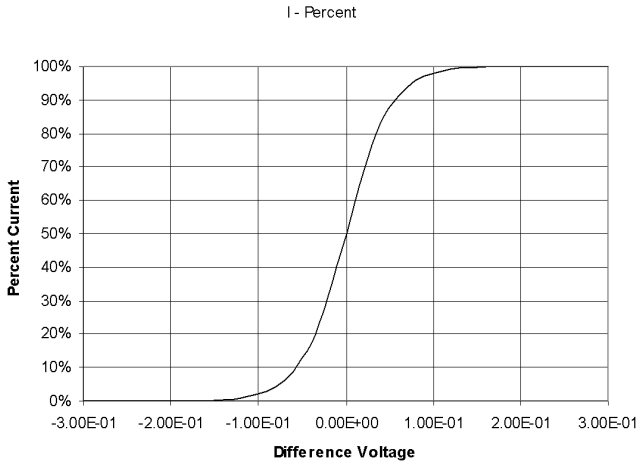

Figure A.7: Plot of percentage of current down a branch as a function of the input difference voltage.

There is a dependence on temperature such that increasing temperature reduces the percentage of current diverted. At room temperature, (25C), a 60mV difference voltage is required to divert 90% of the current. At 155C, (the highest qualified temperature for our SPICE models), nearly 100mV is required for the same effect.

$$I_{C1} = I_{BIAS}\frac{e^{V_{DIFF}/\Phi_T}}{1 + e^{V_{DIFF}/\Phi_T}} \tag{A.10}$$

The current ratio can be easily recast to obtain the single sided current as a function of the input difference as shown above.

Obviously when the full current is switched down one particular branch, the corresponding z_x output would switch between V_{CC} and $V_{CC} - R_C * I_{BIAS}$, which we refer to as V_S. The choice of desired output swing V_S determines the R_C resistor size. Figure A.8 is a normalized voltage transfer characteristic(VTC). The gain in the middle region is about 18*V_S.

If the input voltage swing were not large enough, small changes in the input differential magnitude would cause larger changes in the output, amplifying noise. To prevent this, minimum high and maximum low input values(V_{IL},V_{IH}) are calculated to determine when the gain is unity. ($dV_{out}/dV_{in} = 1$)

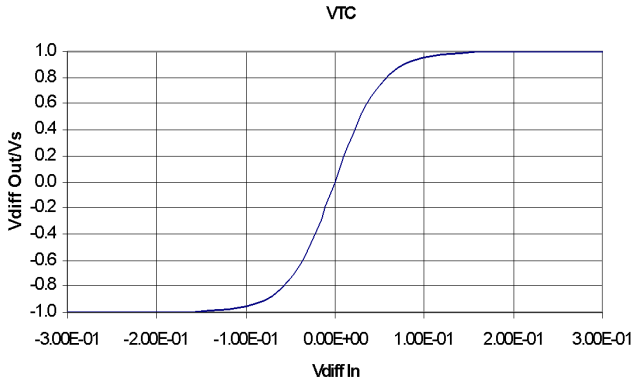

Figure A.8: Normalized voltage transfer characteristic relative to chosen V_S.

$$V_{slope=1} = \Phi_T \ln \frac{-K \pm \sqrt{K^2 - 4}}{2} \tag{A.11}$$

where

$$K = 2 - \frac{2R_C I_{BIAS}}{\Phi_T} \tag{A.12}$$

Using the above equations, VIL and VIH can be determined. The differences between V_{IL} and $V_{OL}(-V_S)$, and between V_{IH} and $V_{OH}(0V)$ are the static DC noise margins NML and NMH respectively [17]. They represent the magnitude of noise that could be present on a nominal output signal between a pair of gates that would still be valid at the input of the second gate. For a differential circuit, this would represent a change in the magnitude of the differential signal. It should be noted that the NM=0 boundary sets an absolute lower boundary of V_S at $2^*\Phi_T$. This is half of that calculated in [16] for single ended CML as expected. The low and high margins are equal due to the symmetrical nature of our circuits. A large noise margin is of course desirable, and this noise margin value grows with increasing V_S. However, increasing V_S beyond a certain point would eventually result in saturation of the transistors at the top of the tree, resulting in a slower circuit. We can see that the transistors would enter saturation when $V_S = V_{BE}(on) - V_{CE}(sat)$. Providing a noise margin for that condition would require $V_S = V_{BE}(on) - V_{CE}(sat) - NM(sat)$.

Solving for the noise margins simultaneously and plotting the results we obtain the plot in Figure A.9.

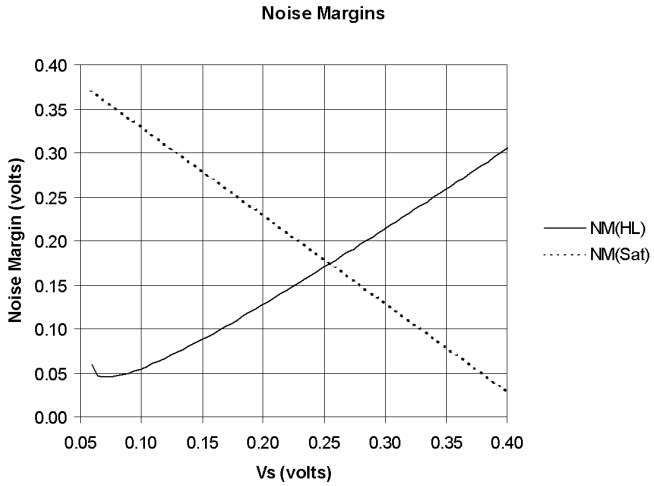

Figure A.9: Optimal noise margin plot for determining V_S. The plot was derived using a temperature of 70C, $V_{CE}(sat) = 0.45V$, and $V_{BE}(on) = 0.88V$.

The intersection of the curves in Figure A.9 indicates the optimal choice for V_S using the specified transistor parameters and temperature. The selection of 250mV for V_S seemed a good engineering choice as the optimal value varied between 200mV and 300mV as the temperature, $V_{BE}(on)$, and $V_{CE}(sat)$ were varied within reasonable ranges.

The choice of I_{BIAS} is determined by the f_T of the two identical transistors used in the circuit. Given the f_T curve in Figure A.5, we choose to have approximately $0.7\text{mA}/\mu m$ based on the transistor size. For the smallest size transistors, that means we would require $R_C = 416\Omega$ resistors. Trees with $4\mu m$ emitter transistors would require an R_C of 104Ω.

In Figure A.10, a 2-input, two level gate is shown. When the a_0 input is high relative to a_1, the current is steered by the b_x inputs and the z_x outputs change accordingly. When the a_1 input is high relative to a_0, the z_1 output is pulled low regardless of the values at the b_x inputs while the z_0 goes high. This single circuit can be used as a logical AND, NAND, OR, or NOR, merely by labeling the pins accordingly. Two level trees can also be constructed to form any two input

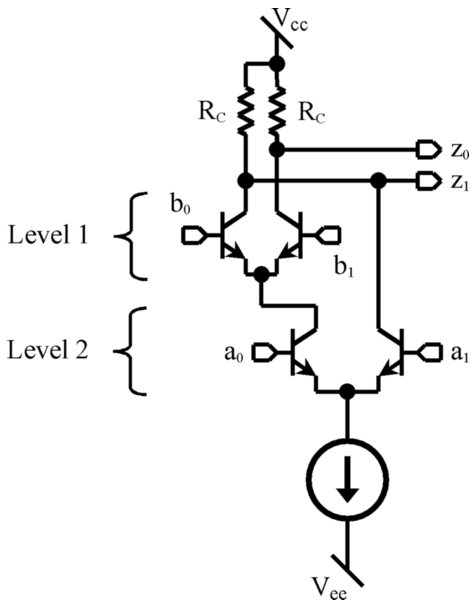

Figure A.10: Simple CML gate with two inputs.

combinatorial logic function and can even incorporate feedback to form latches as in Figure A.11.

When C_0 is high and C_1 is low current is steered through the left level-one transistor pair. Inputs D_0 and D_1 are then active and cause a corresponding output at the Q_0 and Q_1 terminals. When C_1 is high and C_0 is low, the current is steered through the right upper level pair of transistors and these feedback transistors hold the previous state of the D_X inputs. Because there is no current in their branch the D_X inputs are no longer active. This stacking of differential pairs is called "series gating", and it is another reason to require nearly complete switching of current. The percentage of current switched is repeated at each level, and if only 90% of the current were switched at each level in a three level circuit, less than 75% would be available at the top to generate the output swings. With 99% switching we have 97% of I_{BIAS} available at the top.

In general, an N-level tree can have up to 2(N-1) inputs and can be used to create any N- input combinatorial logic function. We use up to three levels in our designs. One side effect of adding multiple levels is that the propagation delay via lower level inputs is generally longer than that for upper ones due to the loading of the transistors above them. Another unfortunate effect of adding levels is that

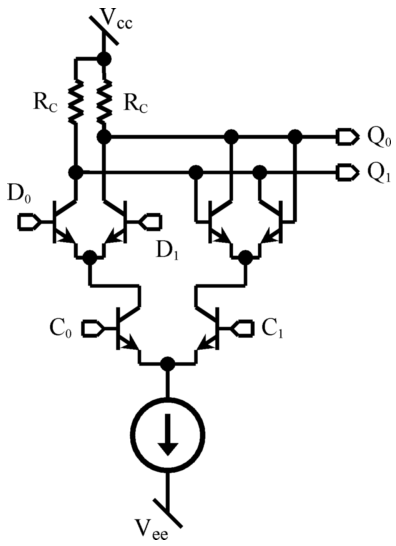

Figure A.11: CML D-Latch circuit diagram.

the signals to a set of inputs effectively sets the collector voltage for the pair of input transistors below them. If the lower level inputs had the same common mode voltage as the upper ones, their transistors would saturate with V_{CE} ≈0V, and the circuit switching performance would be impaired.

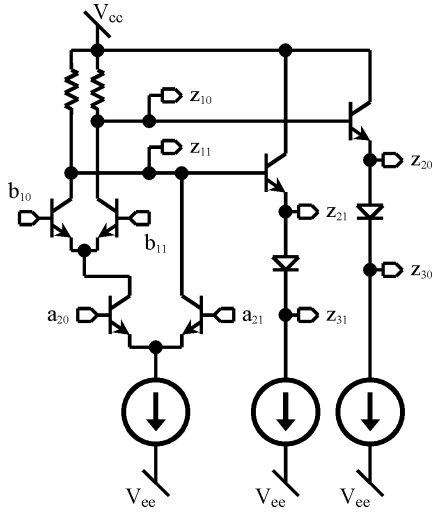

Figure A.12: CML two level gate, with emitter followers providing different output levels.

To address this issue, when lower level common mode voltages are needed, emitter followers like those in Figure A.12 are added. Our level naming convention is also illustrated, with inputs and outputs having subscripts indicating the level

(1-3) and whether it is the main signal line(0) or its complement(1). Each junction in the emitter follower drops the output signal by $V_{BE}(on)$, preparing it for driving gate inputs which require that level. The emitter followers also improve the single-sided rise times over the passive resistor pull-ups. In order to adequately drive different loads, the transistors in the emitter followers can be of a different size than those in the main tree. The bias current for these followers would also be adjusted accordingly.

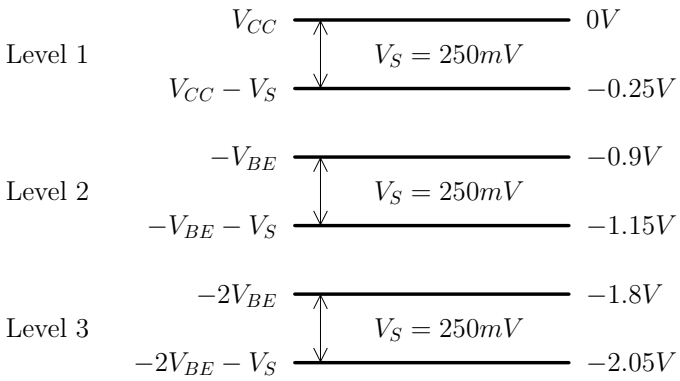

Figure A.13: Nominal voltage levels for CML inputs.

The chosen voltage levels for our three level circuits is illustrated in Figure A.13. The large separation between levels relative to the voltage swing ensures common mode noise will not induce saturation for the transistors on the level below. It should be noted that the inputs to gates with less than three input levels need not necessarily use the specific levels designated, but rather as long as the levels are separated and in the proper order the circuit will work correctly. The levels differ only by a common mode value, and CML is generally indifferent to that. Because of this, a single level buffer can accept a Level 1, 2, or 3 input signal. A gate with two input levels can accept the level pairs: 1 & 2, 2 & 3, or even 1 & 3. In circuits with 4 levels there would need to be additional restrictions due to the possibility of exceeding the collector/emitter breakdown voltages. A Level 4 signal into a simple single level buffer would cause V_{CE} on the input transistor to be 3.34V, which is above the minimum rated breakdown voltage.

It should be noted that fan-out for these circuits is very large, in the sense of static driving capability. The emitter followers can be increased in size to handle

almost any desirable number of loads. With large numbers of loads though, the rise, fall, and propagation delay characteristics suffer. Since the primary impetus for using CML is for performance, we limit loading for mostly speed considerations, and size the transistors accordingly.

A.8 Current Sources

One final piece for our logic family is required, and that is the current sources positioned at the bottom of each of the trees. We decided to use an active current pull-down as opposed to a passive resistor for performance reasons. Simulations show a definite performance increase, especially under high loading conditions.

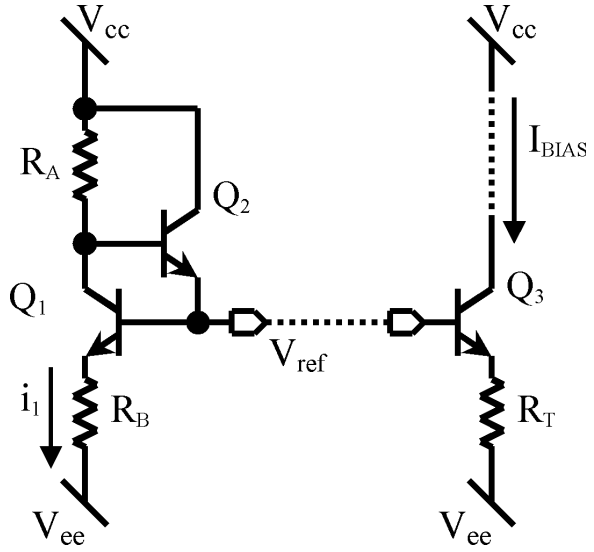

Figure A.14: Current mirror schematic diagram for providing I_{BIAS} current for CML trees and emitter followers.

The constant current source at the bottom of the tree is provided using a Widlar current mirror circuit. In Figure A.14, the circuitry on the left provides a nearly fixed voltage reference, V_{ref}, by using forward biased junctions and resistive drops. On the right, the I_{BIAS} current line represents a gate or emitter follower circuit as described previously. In some implementations of a current mirror, Q_2 is absent and a base-collector short is placed across Q_1, effectively turning it into a diode. The presence of Q_2 allows greater current drive capability, with nearly all of the V_{ref} drive current arriving via its collector. The V_{ref} voltage is $i_1 * R_B + V_{BE}(on)$

above the V_{EE} rail, and is derived via the following equation:

$$V_{ref} = V_{EE} + V_{BE}(on) + R_B * \frac{V_{CC} - V_{EE} - 2 * V_{BE}(on)}{R_A + R_B} \tag{A.13}$$

Since the reference isn't switching, the current i_1 is sometimes set to a value below that of the current maximizing f_T in order to save power. The tail resistor R_T allows us to adjust the current for individual trees serviced by the same V_{ref}. This is useful as trees using larger transistors require a larger bias current in order to maximize their performance. For a particular tree,

$$I_{BIAS} = \frac{V_{ref} - V_{BE}(on) - V_{EE}}{R_T} \tag{A.14}$$

The derivative of the above with respect to V_{ref} shows we need to have a larger R_T if we want to reduce the variations in the bias due to variations in the reference. One of the V_{ref} generators can supply a reference voltage to many gates because of the small base currents involved. If the loads is given in terms of the total microns of emitter length, the load microns can have a 50:1 ratio over the reference microns if the peak f_T current is used in the reference, neglecting things like voltage drops due to wiring.

The collector voltage at Q_3 can be fixed as low as -2.7V by a level three signal above it. Avoiding saturation in Q_3 effectively sets a maximum allowable value of V_{ref} as ≈2.25V if $V_{CE}(sat)$ is considered to be 0.45V and $V_{BE}(on) = 0.9V$. The emitter of Q_3 is then one V_{BE} less than V_{ref}, or ≈3.2V. Choosing the difference between this value and V_{EE}, (deciding on V_{EE}), determines the size of the individual R_T resistors. This choice is somewhat arbitrary, but the size of the tail resistors should be on the order of that of the R_C resistors in the current trees, as the same tolerance concerns apply here as there. The trees with the smallest transistors will require the largest resistors, and it was decided that the voltage developed across the tail resistors for these should be roughly the same as the voltage swing. Thus, V_{EE} was decided to be equal to ≈3.2V-V_S, or 3.45V. As V_{EE} is lowered, the chip power requirements grow proportionally as the chip is designed to use a constant current. In the actual design of the SERDES II chip, the V_{EE} rail was designed to

be -3.4V, with the V_{ref} voltage designed to be around -2.2V. The $V_{CE}(sat)$ value of 0.45V is conservative, and the current source transistors do not in fact become saturated. Additionally, the highest speed circuits tend to use less than three levels. In retrospect, and in future designs, I plan to design for a -3.6V V_{EE}, to allow a common mode noise margin for the Level 3 inputs.

Another design concern is the maximum V_{CE} which can appear across the current source transistor at the bottom of the tree, which is $\approx V_{BE}(on) - V_{EE} - V_S = 2.45V$ for V_{EE}=-3.6, or 2.25V for V_{EE}=-3.4V. The first SERDES design was planned around a 4.5 V_{EE}, which would have resulted in a maximum V_{CE} of 3.35V, which would have exceeded the breakdown voltage of the normal HBT(npn). For that reason, the SERDES and SERDES II designs used a high breakdown HBT(npnhb) as the current source, as this transistor has a minimum breakdown of 4.5V. The f_T of this transistor is drastically reduced however, and it is thus a less responsive current source. Also, the bias current for the maximum f_T for the npnhb is only 1/6th that of the npn. To achieve optimal performance the npnhb would need to be six times larger which would increase parasitic effects and layout requirements unreasonably. For these reasons, future designs will use a regular npn as the current source.

A.9 Performance

The performance characteristics of the CML logic family using this technology are extreme. Unloaded gate delays on the order of 10-20ps are certainly achievable, leading to clocking rates as high as 10GHz. Due to the excellent device performance, our circuits can be limited by the parasitic effects of interconnect almost as much as by device characteristics. Using the devices at the upper limits of their performance [50], [34] capability brings to light many effects which are not seen at slower rise time rates. Ringing of outputs can occur if certain rise time thresholds are exceeded [35], [36].

Below are plots showing simulated simple buffer performance both with and without the RC parasitic effects from actual layout interconnection, and how performance is affected by the addition of emitter followers. The simulations were all

done at 75C, and used the built-in parasitic extraction engine in Cadence which creates a RC ladder network for each interconnection.

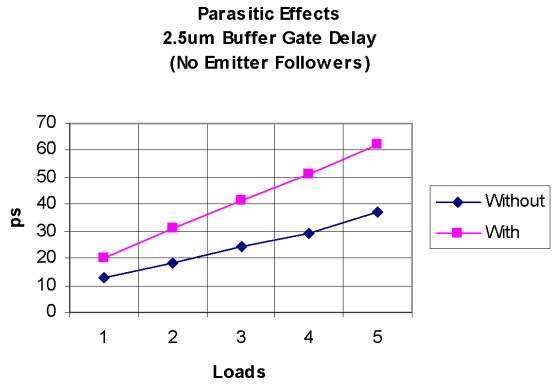

Figure A.15: Plot showing large effect of interconnect on buffer performance for differing numbers of loads.

As can be seen in Figure A.15, the interconnect can degrade performance by as much as 33%. In practice, the amount of degradation is closer to 10%-15% as the simulation layouts used in these plots were not as tightly laid out as are the circuits which we have fabricated. The interconnect effects are largest for the smaller transistors ($1\mu m$) of course, but increasing the transistor size tends to yield marginal improvements on tight layouts when the transistor sizes reach around $6\mu m$.

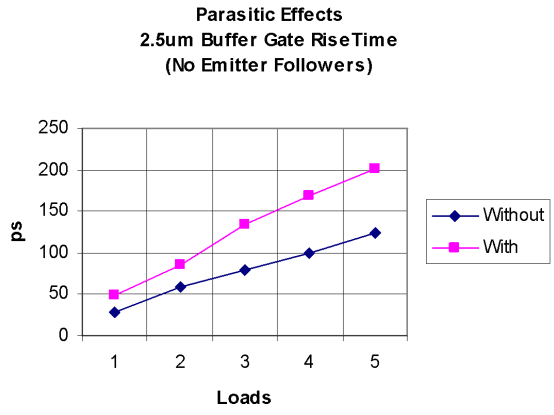

Figure A.16: Interconnect effects on buffer rise/fall times for various loads.

In Figure A.16, it can be seen that the interconnect has an even larger effect

on the rise and fall times of gates. Luckily, the introduction of emitter followers greatly enhances the driving capability of the circuits. This is shown in the next figure.

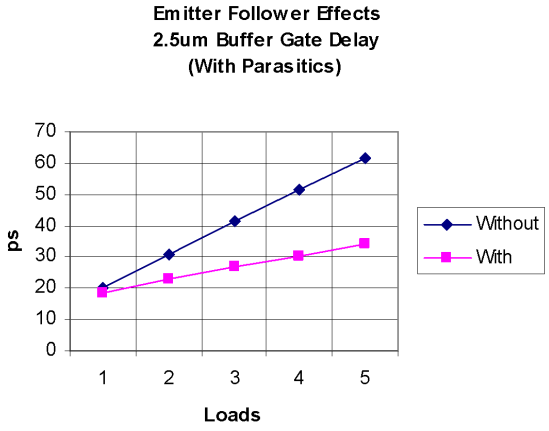

Figure A.17: Gate delay reduction in 2.5μm buffers with the use of 2.5μm emitter followers.

As can be seen in Figure A.17, the introduction of emitter followers greatly reduces the effective delay in CML gates, for nearly all conditions. In fact, the only time it has been observed to actually increase the effective gate delay has been when only a single load is used, and the interconnect parasitics were small. Under those circumstances the added delay of going through the emitter follower transistors was more than that made up for by the extra driving capability.

With emitter followers added, the rise and fall times are greatly improved as well. The only real disadvantage to the use of emitter followers is that the signals must always output on level 2 or lower, and the emitter follower circuits often more than triple the power consumption of a given gate. However, since our focus is on raw performance, we use them quite liberally in our designs. There is an increase in layout size as well, but as we are not creating large circuits it is not of primary concern at this time. As mentioned in the logic section, the emitter follower transistors can be of a different size than the transistors that make up the logical functioning part of the gate, and should be biased accordingly. Increasing the size of the emitter followers improves the driving capability of the gate, but with diminishing returns. In addition, larger emitter followers have a greater load

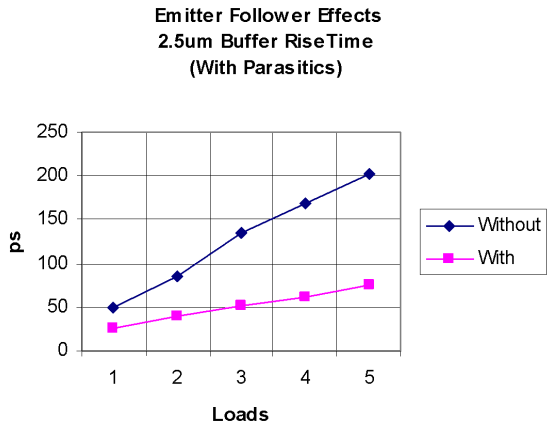

Figure A.18: Effect on rise time of adding 2.5μm emitter followers to 2.5μm buffers.

burden on the logic transistors, and beyond a certain point this will actually degrade performance.

A.10 High Frequency Response

Another concern in the design of our circuits is the effective bandwidth of buffers, gates, and multiplexers. Because we are operating close to the effective limits of the technology, we need to consider that our circuits might not return to a steady state condition in a single bit time. This is the direct result of bandwidth limitations, and the effects are expressed as data-dependent jitter.

Note that the normal small-signal bandwidth calculation methods will not result in an accurate representation as our circuits do not operate in a small signal fashion. The plot in Figure A.19 was generated by sweeping the input frequency of a sinusoid in the time domain. The amplitude and phase were extracted using an aHDL, (analog Hardware Definition Language), program. The simulation was performed on the data inputs of a 2-to-1 MUX, as opposed to the select input, which generally has a smaller bandwidth.

Phase plots are important for determining sources of pulse distortion. As can be seen in Figure A.20 the phase response of the buffers grows more non-linear at the higher frequency extremes. The L_1 output signals show the greatest distortion,

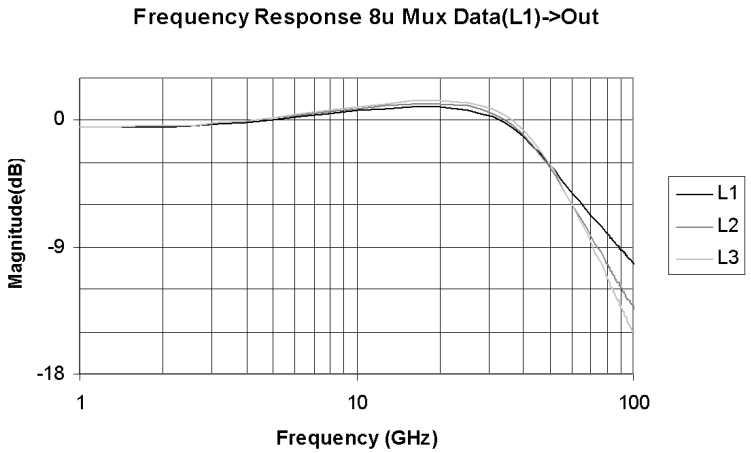

Figure A.19: Large signal bandwidth of $4\mu m$ buffer with three output levels.

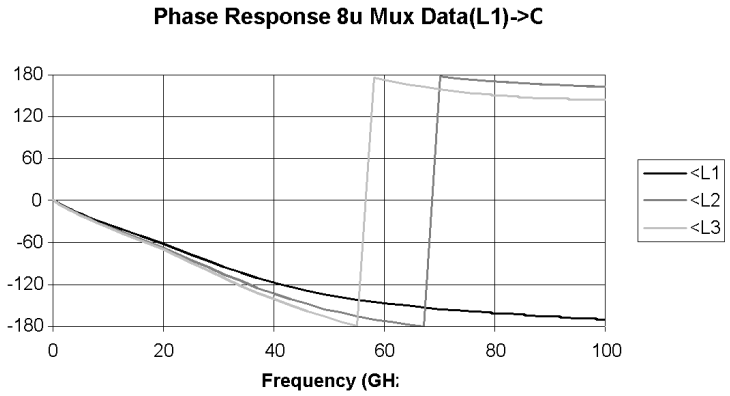

Figure A.20: CML buffer phase response plot using transient large signal swing simulation.

due to their inferior current sourcing capability.

A.10.1 Linearized Differential Buffers

In some circumstances it is necessary to either reduce or eliminate the gain from a common differential pair, for example, when attempting to increase bandwidth. This can be accomplished by the addition of emitter resistors, R_E, as shown in the figure below.

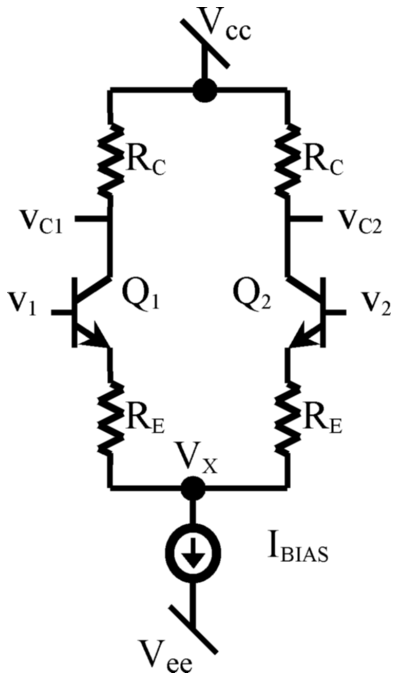

Figure A.21: Linearized buffer circuit with emitter degenerating resistors.

The differential transfer characteristic can be derived from the loop of voltage drops encompassing the base-emitter junctions, the emitter resistors, and the input difference voltage $v_{ind} = v_1 - v_2$. Assuming the following:

$$I_E = I_{ES} e^{V_{BE}/\phi_T} \tag{A.15}$$

V_{BE} can be cast in terms of the emitter current, and the loop equation can be solved for V_{ind} to give:

$$v_{ind} = R_E i_{outd} + \phi_T \ln\left[\frac{I_{BIAS} + i_{outd}}{I_{BIAS} - i_{outd}}\right] \tag{A.16}$$

$$v_{outd} = -\alpha R_C i_{outd} \tag{A.17}$$

$$v_{outd} = v_{C1} - v_{C2} \tag{A.18}$$

$$i_{outd} = i_{E1} - i_{E2} \tag{A.19}$$

$$v_{ind} = v_{BE1} - v_{BE2} \tag{A.20}$$

In the equations above, i_{outd} is the difference current $i_1 - i_2$. Multiplying this difference current by the collector-emitter current ratio alpha, and by the value of the collector resistors R_C, the input value v_{outd} can be plotted in terms of the output value v_{ind}. With the axis exchanged, the familiar transfer characteristic can be displayed.

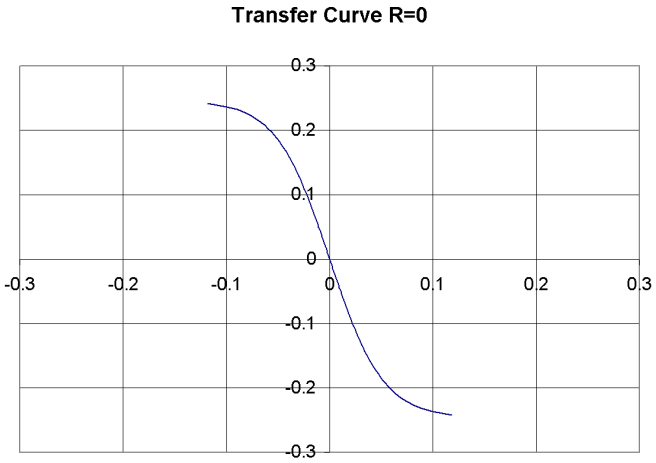

Figure A.22: Inverter transfer characteristic without emitter resistors.

The figure A.22 above represents the transfer characteristic of a non-linearized one micron buffer biased at 0.6mA using T=25C. The collector resistors are 416Ω and β=86 yielding α=0.99. This base plot with no emitter resistors can be compared with the figures below. Note that the derivative of the previous equation can be taken and it yields the reciprocal current gain, which we can recast into the buffer voltage gain.

$$\frac{dv_{ind}}{di_{outd}} = R_E + 2\phi_T \frac{I_{BIAS}}{{I_{BIAS}}^2 - {i_{outd}}^2} \tag{A.21}$$

$$i_{outd} = -\frac{v_{outd}}{\alpha R_C} \approx -\frac{v_{outd}}{R_C} \tag{A.22}$$

$$\frac{dv_{outd}}{dv_{ind}} = \left[\frac{dv_{ind}}{di_{outd}}\right]^{-1}\left[\frac{di_{outd}}{dv_{outd}}\right]^{-1} = \left[\frac{1}{R_E + 2\phi_T \frac{I_{BIAS}}{{I_{BIAS}}^2 - (v_{outd}/(\alpha R_C))^2}}\right][-\alpha R_C] \quad \text{(A.23)}$$

At the midpoint of the transfer characteristic, this reduces to:

$$\frac{dv_{outd}}{dv_{ind}}|_{v_{outd}=0} = \frac{-R_C}{R_E + \frac{2\phi_T}{I_{BIAS}}} = \frac{-R_C}{R_E + 2/g_m} \quad \text{(A.24)}$$

This gives a gain of 4.8 for the 1 micron non-linearized buffer described above. Unity gain can be achieved with a 373Ω R_E.

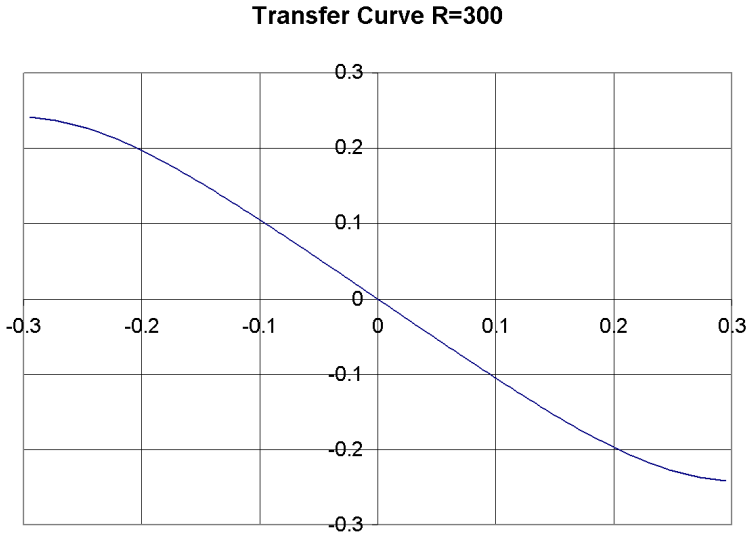

Figure A.23: Linearized buffer transfer characteristic.

The same buffer with 300Ω emitter resistors is shown in Figure A.23 above. Note the near linearity around the center. However, at the extremes, voltage gain can be less than unity. It is also of concern that the addition of emitter resistors will effectively reduce the voltage at the base of the tree V_x, just above the current reference. This voltage will also swing through a much larger range. If a current mirror is used to provide the bias, care must be taken to ensure that the transistor is not driven into saturation.

A.11 Pad Drivers

Our chips have for the most part been probed on wafer, using 50Ω probes, SMA cables and test equipment. Due to the limitations of our testing set up, we are usually limited to examining at most one high speed (>1GHz) differential signal at one time. We make use of DC input control signals to change on-chip configurations to present different output characteristics at the high speed pad. We also have a number of medium-speed outputs and inputs which we can use for sending in reference signals for PLLs, or to obtain a trigger signal for generating an "eye" diagram of a high speed output.

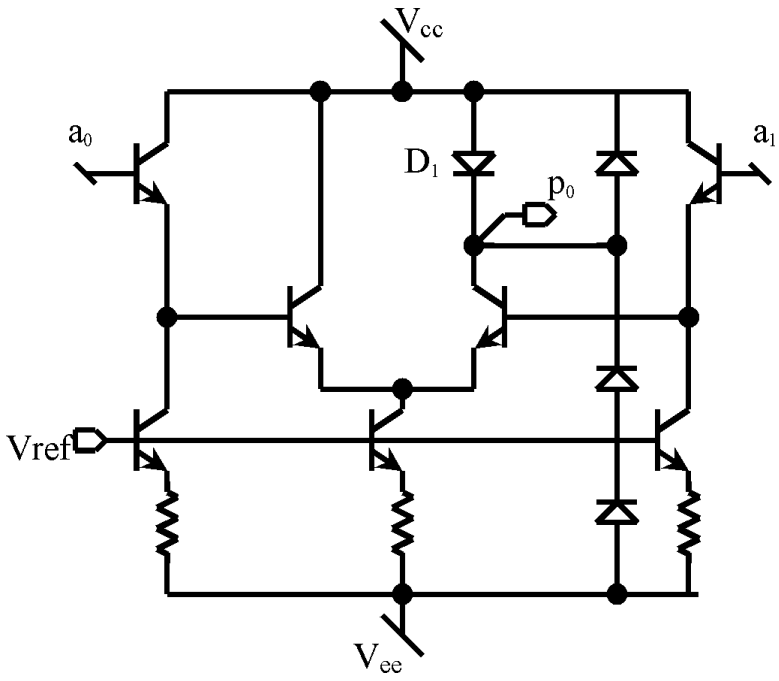

Figure A.24: Single-ended medium frequency pad driver for signals under 1 GHz.

Above in Figure A.24 is a medium speed single-ended pad driver. Not shown is a $1\mu m$ buffer input stage with outputs to a_0 and a_1. A large voltage swing is developed across diode D1, providing a signal sent to a bond pad via p_0. The other diodes are present for electrostatic protection, but regular diodes are used instead of the provided large ESD components in order to reduce the amount of parasitic capacitance introduced. The central core has $6\mu m$ transistors, while the emitter followers, (actually, leaders as shown), are $2\mu m$. With $6*I_{BIAS}$ current across the

diode, we can obtain nearly 800mV DC swings. The pad itself is a large LM plate with a DT mesh or NS plane beneath, and it is modeled as a capacitance with a parasitic reverse biased diode to the substrate.

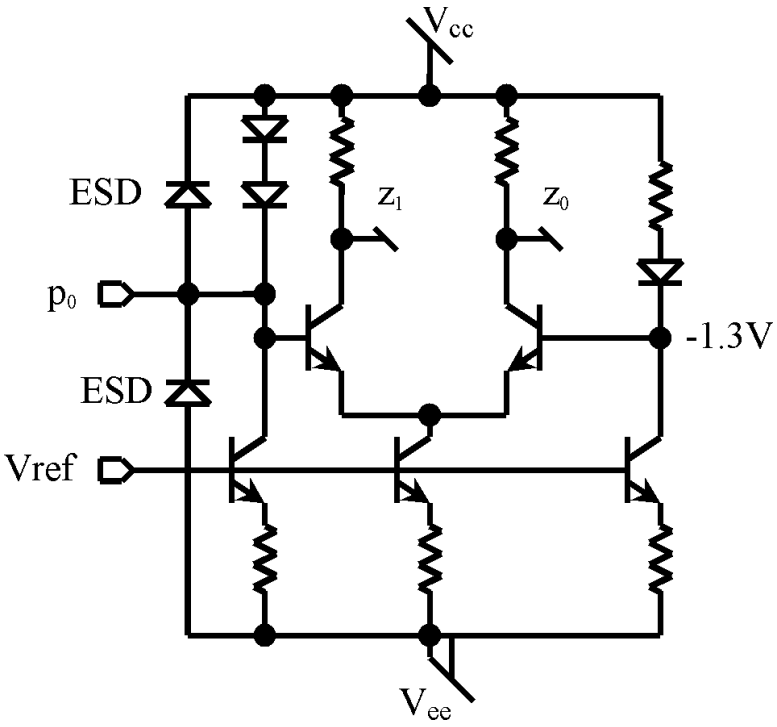

Figure A.25: Single-ended DC signal pad receiver with ESD protection. This pad driver is designed to go to a default state if left floating.

Figure A.25 shows a single ended DC control signal pad receiver. It uses the design kit suggested ESD diodes for static protection. The central buffer is uses $1\mu m$, and the z_x outputs go to $1\mu m$ emitter followers (not shown). These emitter followers are small because they are intended to drive a single load at level two, and the input is not intended to change during normal circuit operation. The other diodes and sources are $2\mu m$ sized and provide biasing so that the current in the central buffer switches when the input is varied around ≈1.3V. When not connected, p_0 will be pulled to ≈1.7V, and z_1 will be brought low. The medium speed input circuits are similar, but have a 50Ω matching resistor added at the pad.

The circuit above in Figure A.26 is identical to the previous one, except that it incorporates 100Ω emitter degenerating resistors R_E, which linearize the buffer

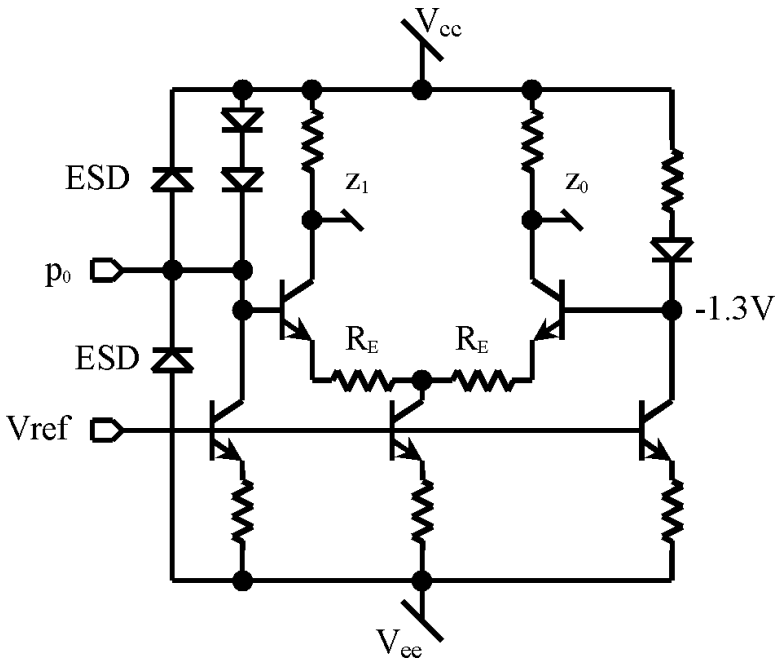

Figure A.26: Linearized analog DC signal pad receiver. Allows linear input voltages to be passed differentially to internal circuits.

response. This is used to directly apply a bias signal when desired, such as when a VCO is controlled by an external analog voltage.

Figure A.27 shows a differential high speed pad receiver with $R_T = 50\Omega$ input resistors. It is a simple $6\mu m$ buffer with resistors sized accordingly, but with 50Ω resistors added for input matching.

In order to drive our signals off chip, we need to consider transmission line characteristics. Our test stand uses probes and cables with 50Ω characteristic impedances (Z_0). The test equipment is assumed to be matched, which means the effective load our driving circuits will encounter is equal to Z_0. Ideally, the source resistance(Z_S) of our output circuits should be 50Ω as well, in order to attenuate any signals reflected back into the output pads. As a first approximation for this, the collector pull up resistors should be 50Ω. This has the effect of creating a 25Ω (Two 50Ω in parallel), effective load for the collectors to drive. Developing a voltage of 400mV would mean supplying as much as 16mA. Because of this, we have opted in some cases to use high impedance output circuits and trust the end termination to hopefully absorb the signals. Future work would involve improving

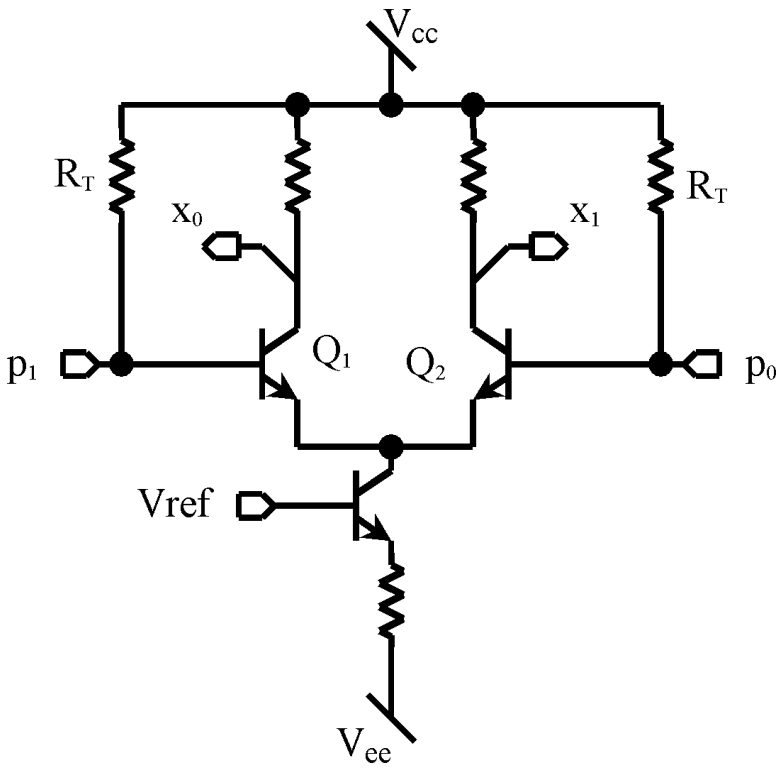

Figure A.27: High-speed differential pad receiver circuit with resistive termination on chip.

on this. At 40Gbps, our fundamental frequency is 20GHz. The signal is made up of frequency components many times higher. At 5 times the fundamental, 100GHz, the wavelength is ≈1.5mm.

A.12 Power Rails

The design of power rails for logic cells arranged in long rows is straightforward if the power usage is fairly even over the length of the rail. This is true for the majority of our CML logic circuits. We can derive some equations useful in design. Consider Figure A.28, which shows the relationship between the voltage along the V_{CC} power rail as a function of distance.

Assuming the circuit is powered from the left, V(0)=V_{CC}. The parameter λ is a measure of how much current is drawn by circuitry along the rail. This can be arrived at by looking at the amount of current per average cell, divided by that cell's length. If you examine a long row of cells, calculate the total current and divide it

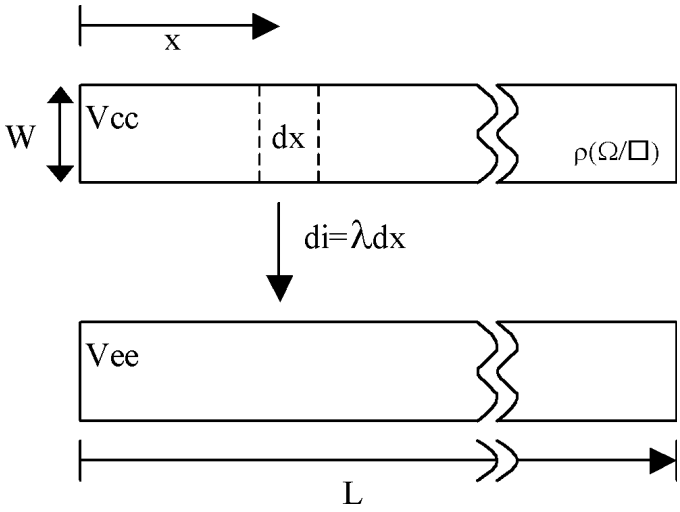

Figure A.28: Power rail dimensions and characteristics.

by the length of the row. The parameter ρ is the sheet resistivity of the metal layer the power rails appear on.

$$\%droop = \frac{V_{CC} - V(L)}{V_{CC} - V_{EE}} \, 100\% \tag{A.25}$$

Droop is the percentage of the original voltage lost to resistance in the power rails. The above value should be set to some reasonable limit, such as 1%. It should be noted that this value is for one-sided supply droop only, the reduction in effective supply voltage is actually twice this. We find the droop by examining the differential element '*dx*'. The current through that element and the material characteristics cause the change in voltage. The current varies along the rail as follows:

$$i(x) = I_{TOTAL} - \lambda x = L\lambda - x\lambda = \lambda(L - x) \tag{A.26}$$

Next, the voltage drop across the differential element is given as:

$$dV = i(x)\frac{\rho}{W}dx = \frac{\lambda\rho}{W}(L - x)dx \tag{A.27}$$

The integration of differential voltage changes, with the addition of the boundary constant V_{CC}, gives the voltage at any given point:

$$V(x) = V_{CC} - \frac{\lambda\rho}{W}\int_0^x (L-\mu)d\mu = V_{CC} - \frac{x\lambda\rho}{W}\left(L - \frac{x}{2}\right) \tag{A.28}$$

We are interested in V(L), which when evaluated and used in the previous droop equation gives:

$$\%droop = \frac{L^2\lambda\rho}{2W(V_{CC} - V_{EE})}\ 100\% \tag{A.29}$$

This equation allows a number of different possibilities to be explored. It can be solved for the maximum length, the necessary width, etc. For illustrative purposes, a high value for λ might be 87.5A/m. The rail width(W) might be $20\mu m$. $V_{CC} - V_{EE}$ might be 3.6V. Lastly, the sheet resistivity of the LM layer is $0.015\Omega/\square$. Using a 1% power droop criterion, we can find the maximum allowable length of the rails to be $1047\mu m$. If the power were distributed on one of the intermediate metal layers, M2..M4, the maximum safe length value would drop to $605\mu m$. Connecting the power rails at both ends would double this length, while requiring a minimum 1% change in total supply voltage would again halve it.

APPENDIX B
LFSR

B.1 Linear Feedback Shift Registers

Linear Feedback Shift Registers(**LFSRs**) are a class of binary pseudo-random number generators(PRNGs). Shown in Figure B.1 and Figure B.2 are the two equivalent simulation diagram architectures which are used to realize them.

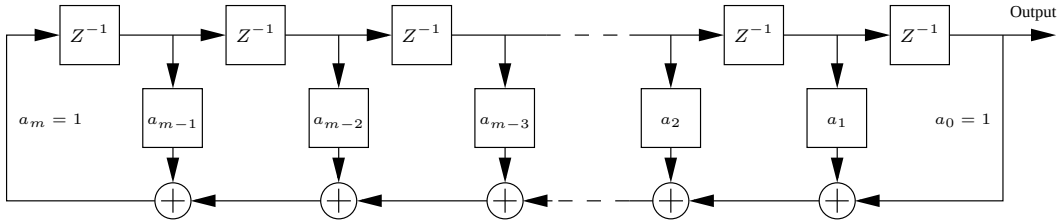

Figure B.1: Fibonacci LFSR

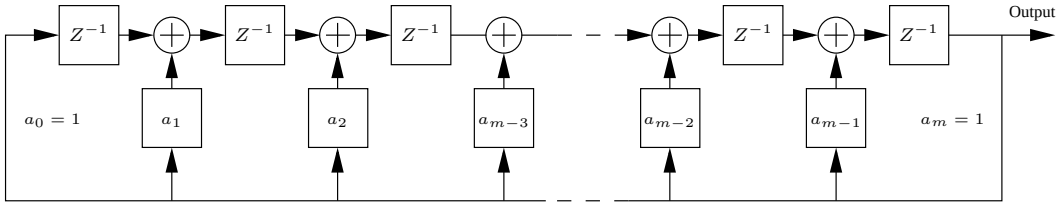

Figure B.2: Galois LFSR

The mathematics governing their behavior is based on boolean finite difference equations. The coefficients $a_0 ... a_m$ are equal to either zero or one. Actual circuit implementations use an exclusive-or, (XOR), operation in place of an adder because in this situation exclusive-or is equivalent to modulo-2 addition. Obviously, when the coefficient is zero, no gate is necessary. The delay blocks are implemented as shift registers. The two forms, Fibonacci and Galois, are mathematically identical. However, the Galois form has advantages over the Fibonacci form. If the feedback in the Fibonacci form is implemented with 2-input XOR gates as shown in Figure B.1, there will be several propagation delays in the path which will lower the maximum clocking rate. This can be partly compensated for using n-input gates, but in general

this becomes infeasible as n grows arbitrarily large. In the case of differential CML logic, the maximum fan-in is directly related to the power supply voltage.

Depending on the choice of coefficients, the registers produce a repeating sequence of bits up to $2^n - 1$ bits in length. With a correct initialization pattern, the Fibonocci and Galois forms can be made to produce the same sequence. One particular state pattern, (all zeros), must be disallowed in these types of PRNGs because it would remain in that state. Some type of reset or method of ensuring a non-all-zero condition on power-up must be designed into the system.

B.1.1 Characteristic Polynomial

The LFSR can be represented by a characteristic or *generator* polynomial $G(Z)$ as shown in Equation B.1. The coefficients in the polynomial are identical to those in the previous LFSR simulation diagrams. The order m of the polynomial is equal to the number of registers which are present.

$$G(Z) = a_m Z^m + a_{m-1} Z^{m-1} + a_{m-2} Z^{m-2} + ... + a_2 Z^2 + a_1 Z + a_0 \tag{B.1}$$

Some coefficients are more useful than others in that they generate long pseudo-random sequences with approximately the same number of ones and zeros. These maximal length, or *m-sequence*, coefficient groups are such that the characteristic polynomial is not factorable, and in addition the characteristic polynomial must divide $Z^N + 1$, where N in this case is the length of the maximal sequence, which was previously given as $2^m - 1$. If these criteria are met by the coefficients, the sequence will be of maximal length, have approximately the same number of ones and zeros, and the maximum number of consecutive bits of the same value is equal to $m - 1$.

Because of symmetry and the modulo-2 nature of addition and subtraction, each set of maximal length coefficients can be reversed in order to derive another set of coefficients which also have maximal length. In fact, the sequences generated by these pairs of sets of coefficients will be the same except that they are in reverse order with regard to each other. Another useful characteristic of the *m-sequences* is

that the same sequence appears at the output of each delay, but at a different offset in the sequence.

It is a common convention to list the indices of the non-zero coefficients when describing a particular LFSR, with the understanding that $a_0 = 1$ in all implementations. Despite the fact that a_m is also always one, it is nevertheless usually listed, perhaps to indicate the order of the LFSR.

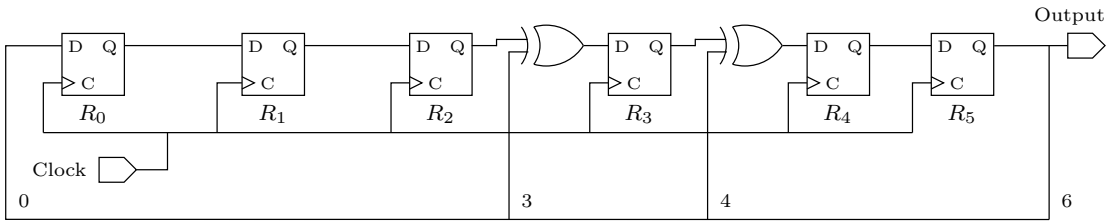

Figure B.3: Galois 6th order {6,4,3} LFSR

Using this convention, for a 7th order LFSR with $a_6 = 1$, $a_4 = 1$, $a_3 = 1$, $a_0 = 1$, and all other coefficients zero, we would indicate this as $\{6, 4, 3\}$. Such an LFSR might be realized using a circuit like that in Figure B.3. Note that the reset or initialization circuitry is not shown. As can be seen in this diagram, the registers are labeled starting with R_0 and continuing up to R_{m-1}. This is the convention I used in my LFSR designs. Note that because R_0, R_1, and R_2 are in series with no intermediate gates, the same sequence will appear at the output of each register, staggered by one delay. In general the same repeating sequence is available at the output of **each** register, but there can be large offsets between positions in the sequence for different registers.

It is worth noting that the number of non-zero coefficients can be different for *m-sequences* for a given m. Usually, having a smaller smaller number of non-zero coefficients will result in a simpler circuit.

Table B.1 provides coefficient specifications for maximal LFSRs for specified values of m. The boldface values are ones I have used in my designs. Note that in some cases four or even more XORs are required to achieve an *m-sequence* LFSR, while in others, as few as two are needed. This is the case for 12-bit and 31-bit LFSRs respectively, which I used in my designs.

Table B.1: Selected LFSR coefficients

n	coefficient sets
3	{3, 2}
4	**{4, 3}**
5	{5, 3}, {5, 4, 3, 2}
6	{6, 5}, {6, 5, 4, 1}
7	{7, 6}, {7, 4}, {7, 6, 5, 4}
8	{8, 7, 6, 1}, {8, 7, 6, 5, 4, 2}
...	...
12	**{12, 9, 8, 5}**, {12, 11, 10, 9, 8, 4}
...	...
16	{16, 15, 9, 6}, {16, 14, 13, 11}
...	...
31	**{31, 28}**, {31, 25}, {31, 30, 29, 20}

B.1.2 Sequence Characteristics

Pseudo-random data is useful for testing serial communication circuits because is is spectrally similar to real world data. Compressed data appears practically random, and various encoding schemes used in serial communication force some of the random-like qualities upon even non-random data. One of the most important characteristics is a zero DC component. (In actuality, this is a 0.5 DC requirement when considering "0" and "1" signals as opposed to "-0.5" and "+0.5".) This is achieved by ensuring there are as many ones as zeros. This allows capacitive coupling of the serial data without significant integrated changes in the coupling capacitor voltage.

There are 2^{m-1} ones and $2^{m-1}-1$ zeros in a 2^m-1 sequence. The maximal run lengths are m ones, and $m-1$ zeros. These each occur exactly once in the sequence. The length of the 31-bit LFSR sequence is 2147483647. There can be no sequence of ones or zeros more than 31 bits in length, and there are effectively as many ones as zeros. However, there can still be a significant DC average offset from the desired value considering short terms with respect to the whole sequence period. Using a C program, I investigated DC averages over what are considered long periods with respect to the thermal time constant of transistors. Despite the maximal run lengths being quite short, the DC average can deviate to the point where the effects must

be taken into account to ensure self heating is not a factor. Using various encodings would eliminate this problem, but such encodings also reduce the need for pseudo-random test data as the encodings incorporate stricter rules on edge frequencies and average value.

APPENDIX C
Clocking Components

To realize complete SERDES systems, we need more components than just logic gates, buffers, and pad drivers. For the receiver, a local oscillator must generate regular pulses so that the input data signal can be sampled at regular intervals. This clock must adapt to variations in the input in such a way that bit errors are minimized. The transmitter requires an accurate clock to generate a clean data signal for the receiver. Because highly accurate oscillators are difficult to create on-chip, the internal transmitter clock is usually derived in some way from an accurate external reference. The adaptation of the local VCO output to an external source is handled by a phase locked loop (PLL) circuit in both cases.

C.1 Phase Locked Loops

Phase locked loops form crucial elements of SERDES circuits. In simplest form they consist of three primary components; a phase detector (PD), loop filter (LF), and a voltage controlled oscillator (VCO). These elements operate together in a closed loop control system as shown in Figure C.1. The phase detector detects a difference in phase between the input and output signals and generates a phase error signal. This phase error is passed to the loop filter that filters out high frequency components, and has a high low frequency gain. The output of the filter becomes the VCO input signal, either increasing or decreasing the frequency until the signals are again in phase. Often a frequency divider is employed in the feedback loop so that the VCO generates an output that is some multiple of the input frequency. In control system terminology, this is a single-input, single-output closed loop system with the variables of interest being the phases of the signals. Since frequency is the derivative of the phase, holding the phase differences at a constant value causes the frequency differences to be zero as well, and the output frequency will track the input over some acceptable range of values. The model is usually linearized, and the VCO is modeled as an integrator. Standard control system analysis usually results

in a second order model.

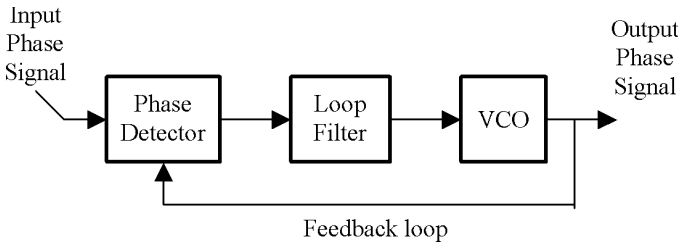

Figure C.1: Simple phase locked loop diagram.

C.2 Phase Detectors

The phase detector is the circuit that compares the input signal to the signal from the VCO and produces an output signal conveying information regarding the phase difference between the two. An "ideal" dynamic detector circuit would produce a linear output over the range $\pm\pi$, as it is usually not possible to detect more than one full period of phase difference since the signals are assumed to be periodic. (Actually, you can have such a detector, but it must incorporate internal states.) There is typically a gain K_d associated with the phase detector that adds to the total open loop gain. In addition, when the phase detector has no input signal applied, it is often useful to have the output of the phase detector be the same as it is when the signal is exactly in phase, which is called the free running voltage,V_{do}. This prevents the absence of the input from immediately causing the VCO frequency to change. In situations like clock extraction, this is very important. The parameter V_{do} also in part determines the PLL's static phase error, and it is desirable to have this parameter be close to zero [18].

One common type of phase detector is simply a multiplier, whereby one signal is used to modulate the other producing a signal with spectral components centered at twice the input frequency and at a frequency of zero. The higher frequency spectral components are either naturally attenuated or intentionally filtered out, ideally leaving a DC signal proportional to the sine of the phase error, if the inputs were originally sinusoidal. For low values, the sine function is approximately linear. The indication of phase error polarity is correct over the range $\pm\pi/2$. A Gilbert Multi-

plier is an example of this type of phase detector. It suffers from the disadvantage that the magnitude of the output is in actuality dependent on the magnitude of the inputs.

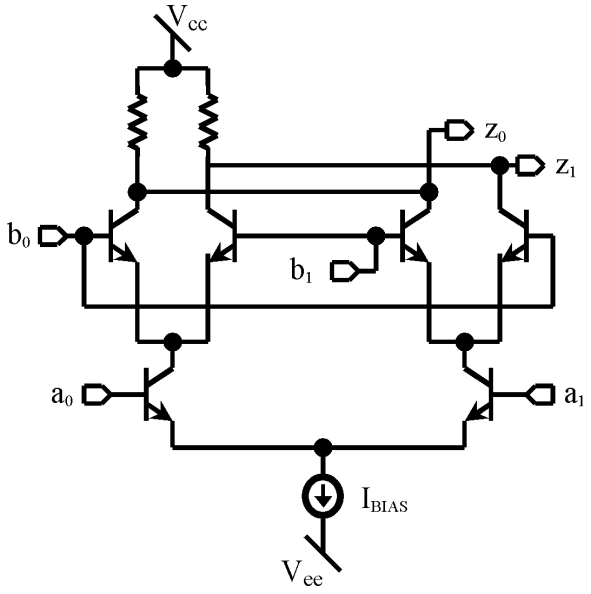

Figure C.2: Gilbert Multiplier/XOR phase detector.

A CML XOR gate has an identical circuit diagram when compared to the Gilbert Multiplier, (Figure C.2), but when used on digital signals its outputs swing from one extreme to the other, removing the dependence on input amplitudes. This is referred to as an XOR phase detector. With square input pulses, the XOR phase detector produces a digital stream of pulses that have a duty cycle proportional to the difference in phase. This pulse signal has an average value that can be used as the detected phase and it is linear over a $\pm\pi/2$ region.

Another common type is the 3-state phase detector, (Figure C.3). This is a digital circuit that incorporates latches, and thus can indicate a phase error greater than $\pm\pi$. It is simple and works well, which is the reason for its popularity. It incorporates two edge-triggered latches and gives output pulses with widths that are proportional to the time between the input pulse rising edges. The average of the output(s) is a linear signal valid over $\pm 2\pi$ radians. These outputs can be used as a single differential signal. This is a good phase detector for clock synthesis purposes, and it is what we use in our designs.

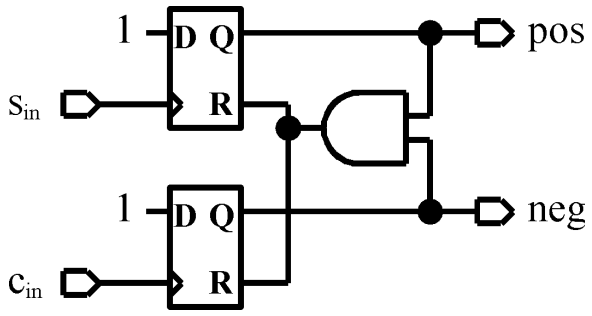

Figure C.3: Three state phase detector. The two outputs indicate edge position and can be combined differentially.

The above phase detectors are mostly suitable for clock synthesis, or for other uses when the PLL input is regular. For our receiver designs, we have to consider the case when edges in the bit stream are not present. When this occurs, the phase detector must supply the same signal as if the input edge arrived in perfect synchronization with the local clock edge. To achieve this, we over-sample the input stream by a factor of two, and compare the values of adjacent samples to determine if an edge is present. The sampling takes place in a round-robin fashion using eight sample latches. If the edge is present, a "fast" or "slow" pulse is generated. Since there are many data sample latches, there are many fast and slow lines. The fast and slow pulses are summed, and the resulting averaged signal is sent to the loop filter. This is in essence a "bang-bang" phase detector when a single edge is examined, but the averaging over several edges smoothes out the response.

C.3 Loop Filters

The loop filter is the primary design control point. By modifying this filter, the behavior of the PLL control system is adjusted. A resistive dividing attenuator can reduce the AC-gain, and hence the bandwidth, but the DC-gain is affected as well. This tends to reduce the signal swing available to the VCO. A passive loop filter using resistors and capacitors can affect the AC and DC gains separately. However, a PLL loop filter is frequently implemented as a proportional-plus-integral active circuit, using an op-amp. By designing the filter to have infinite gain at DC, the static phase error can be made independent of the oscillator center frequency, (V_{CO}), which is a difficult parameter to control.

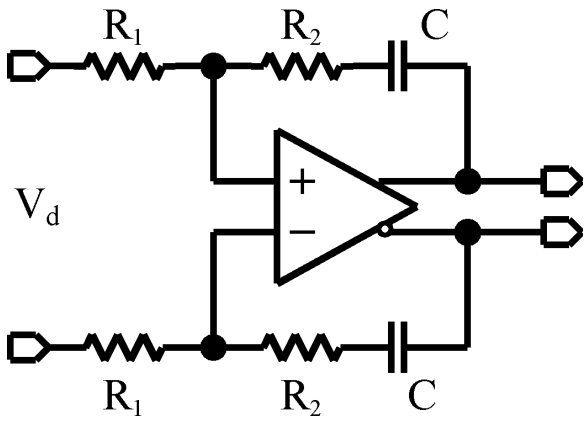

Figure C.4: Differential active loop filter.

Our circuits use differential implementations of phase detectors, loop filters, for the VCO as well. In Figure C.4 the basic circuit for our active differential filter is shown. This circuit has very high gain at low frequencies, and a flat gain K_h for high frequencies, derived from the following equations.

$$F(s) = K_h \frac{\omega_2 + s}{s} \tag{C.1}$$

$$K_h = \frac{R_2}{R_1} \tag{C.2}$$

$$\omega_2 = \frac{1}{CR_2} \tag{C.3}$$

The op-amp introduces an additional pole that eventually causes the output to decay, giving a low pass characteristic with a large pass-band. We actually introduce an earlier pole using a preceding RC low pass filtering stage so the op-amp pole doesn't come into effect. The op-amp is realized using an nfet differential buffer followed by a CML buffer with emitter followers.

The primary purpose of the loop filter is to adjust the PLL bandwidth. This changes how well the local VCO can track the incoming reference signal. PLLs can be tuned to screen out undesirable effects. For example, if the input reference clock has an accurate average frequency, but possesses more phase noise, (undesirable high frequency components), than the local oscillator, the loop bandwidth should be made small so that the output will have accurate frequency, with less noise. The loop bandwidth directly affects the step response, so the lock-in time of the PLL is also affected.

C.4 VCOs

Our designs make use of voltage controlled ring oscillators. Ring oscillators were chosen over other types because of the multiple phase requirements of our architectures. Often, when multiple phases are required, they can be generated from a single phase oscillator at a frequency several multiples higher. In our case we found we could not create single phase oscillators with a high enough frequency. This is unfortunate as ring oscillators tend to have very poor phase noise characteristics [19]. Because of this, we typically use a large bandwidth PLL to lock on to the external references and reduce the phase noise in our clocks.

The basic method of ring oscillator operation is that when an odd number of digital inverters are placed in a ring, the circuit becomes unstable with a period equal to 2N times the propagation delay in one of the inverters, (Figure C.5). The output of each of the N inverters carries the same signal, but at a phase offset of twice the propagation delay, thus supplying an N-phase clock. According to the Barkhausen criteria, for oscillation to occur there must be unity gain around the loop and a 180° total phase shift. This gain and phase shift is provided equally by the buffers. When using differential CML circuits, inversions are accomplished by wiring, (not shown).

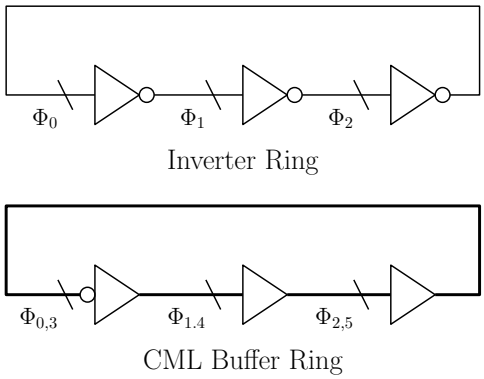

Figure C.5: CML Ring Oscillator .vs. Inverter Ring Oscillator.

This allows us to put an even or odd number of buffers in a ring, with a single wired inversion. The period of the CML ring is 2N times the propagation delay of one of the buffer elements, but because each of the buffer outputs can be used in an inverted or non-inverted sense, you can obtain 2N phases separated by a single propagation delay. The frequency of the ring can be modulated if desired by making the delays through the various elements adjustable using some sort of controlling signal. If the duty cycle of the phases isn't critical, a single delay element can be made adjustable, but it is more common to have all the delay elements identical.

In our original SERDES design, we used a ring VCO design conceived by Sam Steidl [20]. He came up with a method of speeding up or slowing down a ring VCO by modifying the I_{BIAS} current for the simple buffers,(see Figure A.6), by using a variable voltage reference for the current mirror. As can be seen in Figure A.5, the bias current affects the maximum switching speed. As the modified current mirror reference voltage reduced the bias current to slow the VCO, the VCO architecture was designated a "Current Starving VCO". This method had limitations, as the control exhibited a peaking response rather than a monotonic characteristic. In addition, the output voltage swing decreased with decreasing bias current.

In our search for better ring VCO designs, one of the discovered performance increasing methods was to use the interpolation of phases at buffer stages, shown in Figure C.6 on the right. Each buffer element would have two inputs, and the output signal consisted of an interpolation of the two. While one of the inputs was tied to the output of the previous element, (as in a conventional ring), the second

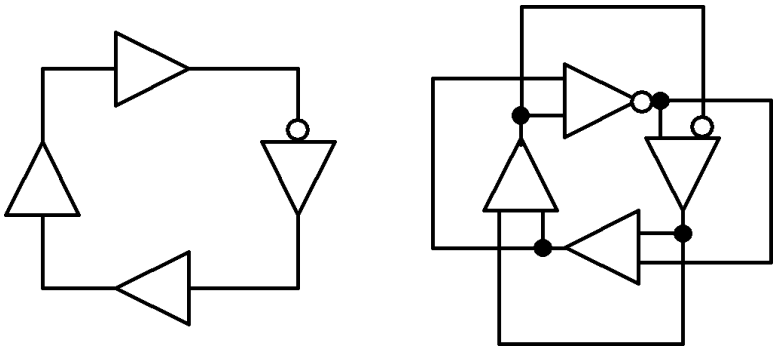

Figure C.6: Normal ring oscillator verses a "forward leap" architecture.

input is connected to the output of the element preceding the preceding element. In this way, the output of an element "leap- frogs" ahead of the element immediately ahead of it. This is the type of design used in the second SERDES prototype chip. The details of the actual phase interpolation can be seen in Figure 2.19.

The VCO has proven to be one of our most difficult design challenges. Research was done into methods of designing them for high speeds, approaching f_{MAX}. Simple bipolar ring oscillator design is discussed in [21]. A quadrature generating ring VCO at $f_{MAX}/4$ with a single phase at $f_{MAX}/2$ is described in [22], which uses a reverse-biased junction for control. The ring VCO described in [23] uses capacitors which bridge the input and output of the buffer for frequency control, and uses mixers to simultaneously increase the frequency and modulate incoming quadrature signals. A novel method of generating quadrature signals at the same frequency as a reference is given in [24]. The "leap frog" architecture mentioned in [2] achieves a higher frequency of operation over a traditional ring of buffers/inverters via the 50-50 interpolation of the phase entering each buffer with the previous phase signal. The frequency control is implemented via variable propagation delay in the buffers. The same interpolation scheme is also implemented in a fixed frequency CMOS ring oscillator in [25]. Earlier designs have also incorporated phase interpolation for speed-up in ring buffer oscillators. One of these is a VCO described in [26], where the phase interpolation ratio is variable, and is used to control frequency while allowing speed-up. However, that design was very asymmetric and unsuitable for uniform phase applications such as ours. Lastly, the VCO architecture in [27] is nearly identical to that in [2], except that the fed forward phase passes through a

buffer external to the ring rather than one inside it. That allows the external buffer to add delays without compromising the maximum frequency of the buffers inside the ring.

C.5 Jitter

No real world signal is completely periodic or perfectly timed. Characterization of these imperfections is called "jitter" in the time domain and "phase noise" when examined in the frequency domain. Examining individual edge placement events makes more sense in the time domain, and so "jitter" is used there. The difference in periods between two successive cycles in can be called "cycle-to-cycle" jitter. Measurement of the deviation between an ideal and a real signal period boundary is called "absolute" jitter.

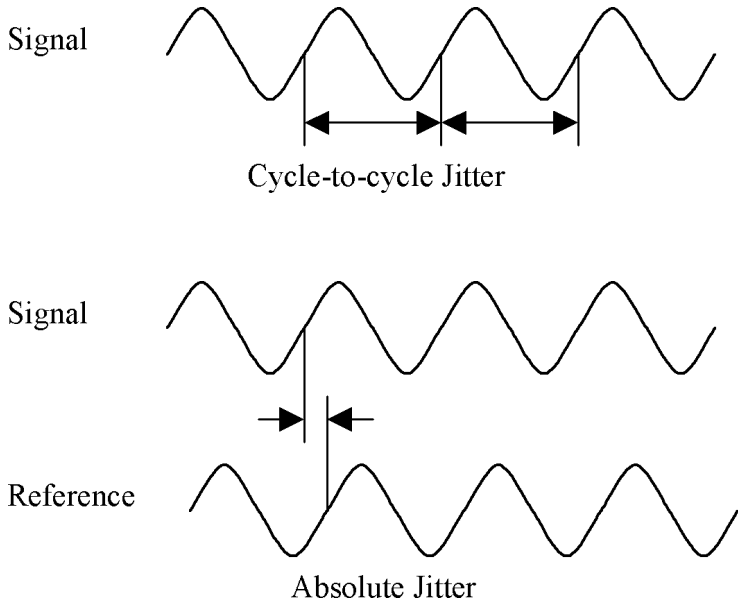

Figure C.7: Diagram showing jitter measurement.

Absolute jitter is the more commonly used term. Instantaneous jitter is the measurement of error for a particular event relative to an assumed "ideal" clock. Since jitter is a measure of time, it has the units of seconds. However, for purposes of comparison it is sometimes described relative to a period of the reference signal, or a bit, in which case it can be expressed as being a unit-less fraction of such a "unit interval" (UI).

Considering individual jitter events doesn't easily lead to characterizations of performance, therefore aggregate values such as average, RMS, and peak-to-peak jitter are usually calculated. Peak-to-peak jitter is only meaningful for jitter that is bounded.

There are two basic types of jitter usually considered. These are random jitter (RJ) and deterministic jitter (DJ). Random jitter is usually unbounded and gaussian in nature and cannot be predicted, while deterministic jitter is bounded and can be predicted with enough knowledge of the system characteristics and the values of the bit stream. For random jitter, the RMS value is equivalent to the standard deviation, and thus uncorrelated variances can be added.

$$\sigma_{TOTAL}^2 = \sigma_1^2 + \sigma_2^2 + \sigma_3^2 + \ldots \tag{C.4}$$

Deterministic jitter can have any distribution function, including uniform distribution or discrete values. These functions are not always known. Since deterministic jitter does have well defined bounds, peak-to-peak values are specified. When calculating aggregate deterministic jitter, the peak-to-peak values can be added to calculate a worst case bounds.

Deterministic jitter can be categorized into four general types. Pulse width distortion (PWD), is usually associated with different rise/fall time characteristics of circuits and leads to longer or shorter "high" times with respect to "low" times. Our use of differential logic effectively eliminates problems associated with PWD. Another name for PWD is Duty Cycle Distortion, or DCD. Sinusoidal jitter (SJ) is mostly of theoretical interest, as system performance is often tested with a sinusoidally time varying phase signal. Data dependent jitter (DDJ), is our biggest concern. This refers to the condition whereby the output of a bit is affected by the values of the prior bits. In essence, the circuit does not necessarily reach a steady state condition in a single bit time, and that state affects the next bit transmitted. DDJ is also referred to as Intersymbol Interference, (ISI). It is a direct result of bandwidth limitations in the circuit. It is possible for both low frequency and high frequency bandwidth limitations to cause this effect. For high frequency cutoffs, the signal might not reach a full high or low bit value in a single bit time, causing

the next bit if different to get a "head start" towards its final value. In a low frequency cutoff situation, a string of several bits of the same value will slowly decay towards the average value, and when a bit change does occur it has a "head start" towards the opposite value. Lastly, there are other deterministic jitter effects which are bounded, but uncorrelated with the bit stream. Jitter arising from power supply variations or cross- talk fall into this category.

Jitter can be measured directly from circuits in a variety of ways. A spectral analyzer can directly measure the standard deviation of an input signal given an appropriate trigger. For example, a clock recovery PLL output can be measured directly, or it can be measured relative to the transmitter clock. The PLL can also be tested with no input to check the spectral characteristics of the VCO. Due to the limitations of testing equipment, care must be taken when such measurements are made because often the equipment cannot sample waveforms in real time. Instead, periodic waveforms are built up over many cycles. Often the triggering event must occur dozens of cycles prior to the actual data acquisition, and this introduces an effective increase in the jitter measured proportional to the square root of the delay [28].

A single total jitter value can be calculated if a specification of allowable jitter magnitude, and the percentage of events that can violate that condition, are known. Usually this specification is based on the amount of jitter that would cause an incorrect bit value, and on an acceptable bit error rate (BER). These values can be used to find an equivalent peak-to-peak values for random jitter, which can then be added to the deterministic jitter.

Using knowledge of the jitter characteristics of a clock which is sampling a serial stream, the noise present in the signal, and the shape of the waveform, it is possible to directly estimate bit error rates. In Figure C.8, a pulse eye diagram is shown along with greatly exaggerated jitter and noise histograms. The noise histograms on the left are only representative of the noise at the sample instant T_S. If T_S is misplaced due to jitter, the shape of the noise distributions don't generally change, but if they move closer together. There is a finite probability that a bit error will occur even in the absence of jitter, as the noise distributions overlap. The

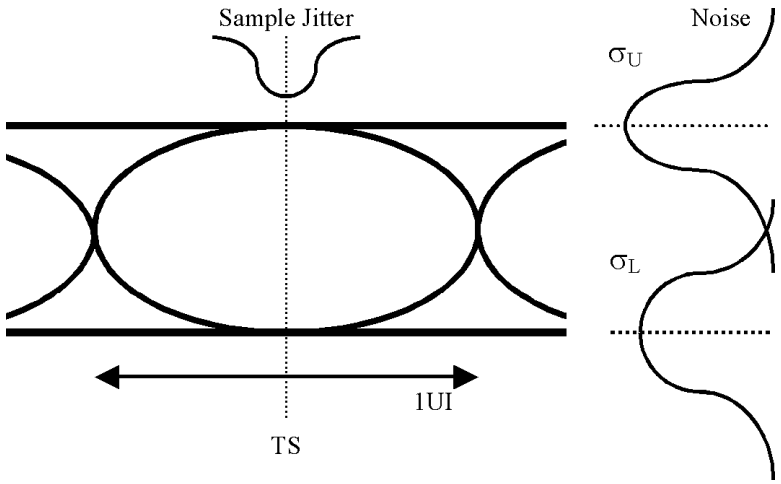

Figure C.8: Sampling a digital waveform at a jittered time T_S, with noise σ_L and σ_U.

probability of a misidentified bit can be calculated as a function of the sample time. When that function is multiplied by the sample jitter distribution and integrated over time, the total probability of error can be calculated.

C.6 Phase Noise

Phase noise is the frequency domain characterization of jitter. It is more commonly used in the context of analyzing physical circuits that generate signals. The introduction of a phase noise signal into an otherwise perfect sinusoid will result in the formation of side- bands of noise near the center frequency. The result of a periodic sinusoidal phase error is a modulation, forming frequency impulses above and below the desired center frequency. If the errors are gaussian or otherwise random, they generate a noise "skirt" around the central frequency.

Most phase noise analysis starts with device noise in circuits and the resulting generation of noise at the circuit output. Individual noise sources are quantified by normalized power density with a particular distribution over frequency. Thus, the units are usually $Volts^2/Hz$. This power spectral density can be multiplied by the circuit power transfer function to find noise power at the output. To make the noise values more meaningful, they are usually normalized once again relative to the power of the ideal carrier output. The noise is then expressed in terms of

dBc/Hz, which are decibels relative to the carrier per Hz. Specifications of phase noise have to indicate how far from the central peak that they are taken, so a typical specification might be, "The phase noise for oscillator X is -90 dBc/Hz at a 1MHz offset from the carrier." When integrated over the circuit's noise bandwidth, the total noise power at the output can be found. Resistive thermal noise is gaussian with its variance (or spectral power density) given by:

$$\overline{V_N}^2 = 4kTRdf \tag{C.5}$$

Many treatments of noise in LCR oscillator tank circuits use only the resistive noise in calculating the phase noise spectrum. The rest of the circuit is considered noiseless. More detailed analysis can incorporate shot noise as well as other noise effects. Final calculations usually end up with a spectrum that falls off as $1/f^2$. This is due to the fact that the bandwidth of the of the LCR circuit falls off as $1/f$, and the power bandwidth is equal to that value squared. In practice, there is always a "noise floor" present below which the noise density doesn't fall. Also, near the carrier peak, the drop-off is usually steeper than the predicted $1/f^2$, instead being $1/f^3$.

Noise calculations for bipolar ring VCOs are similar, but they require further assumptions because there are multiple locations were noise is generated. The noise is band limited by the low pass filter formed by the resistors and the capacitance attached to the outputs either intentional or parasitic. The noise generated is independent of the number of stages in the ring [28], because the "active edge" which is traveling through the ring is present in only one buffer at a time. This is very different from CMOS ring VCOs, where the number of stages can influence the noise in a differential implementation, but not in a single ended design [19]. This is because the primary noise generating components are different, and have different distributions(white .vs. 1/f). According to [29], the SiGe HBT has superior noise characteristics when compared to normal bipolar Silicon BJTs, so it is reasonable to assume the resistors are still the dominant source of noise in our ring VCOs.

Another means of analyzing noise performance is based on the impulse sensitivity function (ISF). In this type of analysis, the oscillator is considered to be a

linear time variant system as opposed to a time-invariant one. The rationale for this is that the noise generated by components at different times in the oscillator period causes different effects. For example, in a circuit that has amplitude limiting, (as all oscillators must), the addition of noise when the circuit is at a peak or minimum has little effect on the output phase of the signal. The circuit tends to restore the correct amplitude before starting the transition to the other extreme. On the other hand, a noise impulse present during a rising or falling edge directly advances or retards the phase. By determining a circuit's impulse sensitivity function, phase noise can be predicted. A full treatment of this subject is presented in [19].

C.7 Clock Recovery

Clock recovery is the process of extracting a clock from a serial bit stream. For NRZ data, the process usually involves detecting edges in the serial stream relative to those in a local oscillator and adjusting the local oscillator phase accordingly. The characteristics of the loop filter used in a clock recovery PLL are dependent on the quality of clock used in the transmission of the data. A small loop bandwidth is desirable if the remote oscillator is of high quality. Any variations in edge location will be random and generated enroute. If the transmit clock is particularly noisy, a larger loop bandwidth and quick response is needed to stay aligned. Using the correct bandwidth is important as a large filter bandwidth with a good transmit clock will cause the local oscillator to track the random variations introduced by the channel, increasing the likelihood of losing lock. Likewise, a small filter bandwidth coupled with a poor transmit clock will result in a system which cannot vary quickly enough to lock to the data. The loop bandwidth also has implications for the local VCO slew rate, affecting the speed at which the system can acquire lock.

C.8 Data Retiming

Data retiming is the process of taking digital data and allowing transitions to occur at only specific points in time, usually coinciding with a an edge from a local oscillator. For complete retiming of data, a clock that has a frequency of twice the bit rate is used A master-slave latch can then ensure that edges are present

only during the correct times. The requirement of a 2X clock is one of the most significant bottlenecks in monolithic transmitter design and is what prompted us to investigate lower frequency clock schemes. Normally, data to be transmitted is multiplexed to a serial stream, and then that re-clocked, ensuring accurate placement of edges. Because the output passes through the same circuit, variations in process parameters do not affect the inter-bit characteristics. However, the bandwidth of the final circuitry must be large enough to prevent data-dependent jitter effects from arising.

APPENDIX D
Source Code

D.1 bits31.c

The following program is used to generate next-state functions to allow a state machine to advance by several states per clock. It is easily modifiable for any state machine and for any number of states that need to be advanced by.

```
#include <stdio.h>

char *XSTR[31];

void initstr(void)
{
 char buf[50];
 int i;
 for(i=0;i<31;i++) {
  sprintf(buf,"Reg[%d]",i);
  XSTR[i]=(char *)malloc(sizeof(char)*(strlen(buf)+1));
  strcpy(XSTR[i],buf);
 }
}

void printstr(void)
{
 int i;
 for(i=0;i<31;i++) {
  printf("Reg[%d]=%s\n",i,XSTR[i]);
 }
}

void NextState(void)
{
 char *tmp;
 int i;
 tmp=(char *)malloc(sizeof(char)*(
    strlen(XSTR[30])+strlen(XSTR[27])+4));
 strcpy(tmp,"(");
 strcat(tmp,XSTR[27]);
 strcat(tmp,"^");
 strcat(tmp,XSTR[30]);
 strcat(tmp,")");
 free(XSTR[30]);
 for(i=30;i>0;i--) {
  XSTR[i]=XSTR[i-1];
 }
 XSTR[0]=tmp;
}

int main(void)
{
 int j;
 char *tmpstr;
 initstr();
 for(j=0;j<16;j++)  NextState();
```

```
 printstr();
 return(0);
}
```

APPENDIX E
Equipment Specifications

E.1 Signal Source SML01

Manufacturer: Rohde & Schwarz, Inc.

9 kHz to 1.1 GHz (3 models)

0.1 Hz frequency resolution

-140 dBm to +13 dBm (+19 dBm overrange)

SSB phase noise <-122 dBc/Hz (at f = 1 GHz, Delta f = 20 kHz)

Setting times <10 ms

High level accuracy (<0.5 dB at levels >-120 dBm)

High reliability through electronic attenuator

Digital frequency and level sweep

AM/FM/Phase Modulation

Optional pulse modulator with integrated pulse generator

3-year calibration cycle

GPIB and RS-232 interface

E.2 Probes

GGB Industries Inc. Picoprobe

10 pin Multi-Contact Wedge(MCW) MCW-8-2880

4 dedicated power and ground pins

6 signal pins, DC-40GHz

K-connectors

150μm pitch,

Pattern: S-S-P-G-S-S-G-P-S-S

E.3 Oscilliscope 11801C

Manufacturer: Tektronix

50GHz bandwidth

Vertical Resolution – 8 bits

Horizontal Resolution – 1 ps/div to 5ms/div

Max. Sample Rate 200ks/s

SD-24 Dual 20GHz Sample Heads

SD-32 50GHz Sample Heads