

**Memory and Low-Swing Interconnect for Energy Harvesting
Circuit Applications**

By

JAMIE S. COLLIER

B.S. (Case Western Reserve University) May 18, 2003

THESIS

Submitted in partial satisfaction of the requirements for the degree of

MASTER OF SCIENCE

in

Electrical and Computer Engineering

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Chair, Dr. Rajeevan Amirtharajah

Member, Dr. Stephen Lewis

Member, Dr. Venkatesh Akella

Committee in charge
2005

© Copyright by Jamie S. Collier 2005
All Rights Reserved

Abstract

The recent explosion of small scale, embedded sensing and processing electronics has made the growing disparity between battery and circuit technologies even more pronounced. Energy scavenging systems have shown to be a promising supplement or replacement for battery technology in many cases. However designing these systems involves many challenges, including integrating the scavenger's behavior into the circuit design, redesigning standard systems such as memory to accommodate the variable power supply output, and reducing the energy dissipation of the electronics wherever possible. This thesis explores three areas aimed at reducing power and simplifying the development of circuits technologies for wireless sensor nodes. The first is a description of circuit based modeling for mechanical vibration energy harvesters. This work enables models of the generators to be built into circuit simulators, which will in turn allow for more integrated design of electronics for energy scavenging power supplies. Next, a memory design for a Distributed Arithmetic based filter is presented. This memory is designed to work with both an AC and a DC power supply to decrease the system dependence on voltage regulation for correct operation. The memory cells can hold data for up to 1.2 *ms* without refreshing and the sensing circuits require no DC bias currents. Finally, two methods of reducing interconnect power are considered. The range of operation and power estimates for six low-swing driver and receiver combinations are determined. The lowest power implementation was able to send data at 10 *Mb/s* across a 1 *cm* wire using less than 2.5 μW average power. A preliminary investigation into a pulse-width modulation system is also initiated with some promising results. Data rates of 10 *Mb/s* were achieved on a 1 *mm* wire with average power of approximately 2.5 μW . This power includes the power used to transmit the reference clock. In a larger bus, the power per bit would decrease with each additional wire because the reference clock could be shared.

Acknowledgments

I would like to thank my advisor, Professor Rajeevan Amirtharajah, for his guidance over the last year and a half. You have helped me to become a better circuit designer, a more thoughtful and critical researcher, and a better leader. I have learned more doing research with you than I thought possible, about everything from electronics to investing to sports. Thank you. I would also like to thank the members of my committee, Professors Stephen Lewis and Venkatesh Akella. You have both provided excellent insight into the inner workings and the applications of my project. I have greatly appreciated working with you and hearing your thoughts and suggestions. In addition, I would like to thank Professors Vojin Oklobdzija and Jamal Tuqan for many interesting conversations and ideas relating to this project and the field in general, and Lori in the ECE department office for her helpfulness and friendship.

I can honestly say that much of my enjoyment of graduate school has come from the students I have had the pleasure of working with. Albert, Bicky, Jeff, Liping, and Zulfi, thank you for creating a lab environment where ideas and concerns can be openly discussed without competition and where everyone feels supported. More importantly, thanks for making work fun, even at 10PM on Friday night. I will never forget the fish with laser beams, that things on tv must be ok, or that rhinos have ups. Mackenzie and Justin, I have not worked with you for long but I am sure you will both add to and enjoy the MCSG. VCL, particularly Mike, you have been a constant source of comraderie and noise in the CERL lab. Thank you for being fun and working all the time so we didn't feel so bad. I would also like to thank the first year posse of the out-of-staters plus Oliver, Albert and occasionally Morgan. You made life here much more enjoyable. Special thanks to Jeff for being an excellent research partner and an even better friend. You have made me a better person in the four years we've known each other. I wouldn't be nearly as good at snowboarding without you! I look forward to hearing your work and life stories, and to meeting you in cool locations to snowboard, even when we're both old and have families. Best wishes to you all.

Finally I would like to thank my family. Even from 2,400 miles away you have

always managed to be supportive and a little crazy. Jolene and Justine, I have enjoyed watching both of you grow into adults. I look forward to spending more time with you, even if it means going to drag races or being pelted with ham! Mom and Dad, you have taught me so many of the skills that have guided me through not only school but life so far. Your strength, honesty, and hard work inspire me every day. Thank you for your love and support. Lastly, I would like to thank my grandfather, Robert Zoller. Some of my earliest memories are of your strong belief in education and the importance of attending college. Although you did not live to see me receive any of my degrees, I am intensely aware of your influence on my life and career. I am proud to join you as the only college graduates in our family, and even more proud that soon my sisters will join both of us. I dedicate this thesis and all of the work it holds to you.

Contents

| | |
|--|-------------|
| Abstract | iii |
| Acknowledgments | iv |
| List of Figures | viii |
| List of Tables | xi |
| 1 Introduction | 1 |
| 1.1 Contributions of the Thesis | 2 |
| 1.2 Organization of the Thesis | 3 |
| 2 Energy Harvesting Generator Models | 4 |
| 2.1 Mechanical Energy Harvesting Mechanism | 5 |
| 2.2 Types of Energy Harvesters | 7 |
| 2.3 Piezoelectric Model | 9 |
| 3 Low-Power Memory for AC Power Supplies | 13 |
| 3.1 Memory Architecture | 13 |
| 3.1.1 Distributed Arithmetic | 15 |
| 3.1.2 DA Look-Up Table | 17 |
| 3.1.3 Controller for LUTs | 18 |
| 3.2 Memory Circuits | 25 |
| 3.2.1 DRAM Cells for AC Supply | 25 |
| 3.2.2 Register Cells | 46 |
| 3.2.3 Individual Cell and Row Decoders | 46 |
| 3.2.4 Circuit Implementation of LUT Controller | 48 |
| 3.3 Full Memory System Implementation | 53 |
| 4 Low-Swing Interconnect and Modulation | 56 |
| 4.1 Interconnect Power | 56 |
| 4.1.1 Dynamic Switching Power | 56 |
| 4.1.2 Static Power of Peripheral Circuitry | 57 |
| 4.1.3 Possible Power Reduction | 57 |
| 4.2 Low-Swing Interconnect | 58 |
| 4.2.1 Driver Circuits | 58 |
| 4.2.2 Receiver Circuits | 59 |

| | | |
|----------|---|-----------|
| 4.2.3 | Implementation and Comparison | 64 |
| 4.3 | Data Compression through Pulse-Width Modulation | 72 |
| 4.3.1 | Modulation Circuit | 73 |
| 4.3.2 | Demodulation Circuit | 74 |
| 4.3.3 | Implementation and Comparison | 75 |
| 5 | Conclusions and Future Work | 78 |
| 5.1 | Conclusions | 78 |
| 5.2 | Future Work | 79 |
| 5.2.1 | Dual-Polarity Architecture | 79 |
| 5.2.2 | Supply Voltage Monitoring and Adaptation | 79 |
| 5.2.3 | Adaptive Body Bias of DRAM Cell Write Path | 80 |
| 5.2.4 | DRAM Read and Write Improvements | 81 |
| A | Layout | 82 |
| | Bibliography | 83 |

List of Figures

| | | |
|------|--|----|
| 2.1 | Generic model for a mechanical vibration to electric energy converter [1] | 6 |
| 2.2 | Generic electrical model for a mechanical vibration system | 6 |
| 2.3 | Reported power outputs for energy harvesters, both simulated and measured | 8 |
| 2.4 | Mechanical schematic drawing of piezoelectric generator | 9 |
| 2.5 | Equivalent circuit model for piezoelectric generator | 10 |
| 3.1 | Highest and lowest mechanical vibration frequencies explicitly investigated, arrows indicate time where supply is unusable | 14 |
| 3.2 | 16–Tap DA Filter Block Diagram | 16 |
| 3.3 | Detailed view of DA unit | 17 |
| 3.4 | Lookup table block diagrams | 19 |
| 3.5 | Master FSM state diagram | 20 |
| 3.6 | Read FSM state diagram | 22 |
| 3.7 | Write FSM state diagram | 23 |
| 3.8 | DRAM cell with NMOS write path | 26 |
| 3.9 | DRAM cell with transmission gates in write path | 27 |
| 3.10 | DRAM cells utilizing alternative threshold voltages | 28 |
| 3.11 | Twenty percent fall times for various storage node sizes over temperature | 29 |
| 3.12 | Testbench for testing individual memory cells | 30 |
| 3.13 | Command and internal voltages for writing a 1 to the storage cell and reading out the 1 after a delay. | 31 |
| 3.14 | Command and internal voltages for writing a 0 to the storage cell and reading out the 0 after a delay. | 32 |
| 3.15 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , single write transistor, length of 180 nm | 33 |
| 3.16 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , single write transistor, length of 180 nm | 34 |
| 3.17 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , single write transistor, length of 270 nm | 34 |
| 3.18 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , single write transistor, length of 270 nm | 35 |
| 3.19 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two series write transistors, length of 180 nm | 36 |
| 3.20 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two series write transistors, length of 180 nm | 36 |

| | | |
|------|--|----|
| 3.21 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two series write transistors, length of 270 nm | 37 |
| 3.22 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two series write transistors, length of 270 nm | 37 |
| 3.23 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two transmission gates in series | 38 |
| 3.24 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two transmission gates in series | 39 |
| 3.25 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , transmission gate tied to store, NMOS tied to bitline | 39 |
| 3.26 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , transmission gate tied to store, NMOS tied to bitline | 40 |
| 3.27 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , NMOS tied to store, transmission gate tied to bitline | 41 |
| 3.28 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , NMOS tied to store, transmission gate tied to bitline | 41 |
| 3.29 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , high threshold write transistor | 42 |
| 3.30 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , high threshold write transistor | 42 |
| 3.31 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , write transistor body bias of -100 mV | 43 |
| 3.32 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , write transistor body bias of -100 mV | 44 |
| 3.33 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , write transistor body bias of -300 mV | 44 |
| 3.34 | Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , write transistor body bias of -300 mV | 45 |
| 3.35 | Static D flip-flop used in register LUT | 46 |
| 3.36 | Read and write sense circuit for DRAM memory cell | 47 |
| 3.37 | LUT 4-bit binary to 16-bit one-hot address decoder with high impedance switch | 48 |
| 3.38 | Read and write line decoders for both DRAM and register LUT | 49 |
| 3.39 | LUT controller top level schematic | 50 |
| 3.40 | Master FSM circuit schematic | 50 |
| 3.41 | Read FSM circuit schematic | 51 |
| 3.42 | Write FSM circuit schematic | 52 |
| 3.43 | 7-bit counter used for generating the LUT addresses and select bits, as well as the update end bit | 52 |
| 3.44 | Example output of the LUT controller, shows signals necessary to write one row of memory cells | 53 |
| 3.45 | Full schematic of DRAM LUT | 54 |
| 3.46 | Full schematic of register LUT | 54 |
| 3.47 | Sample waveforms showing voltage sequence necessary to write and read from DRAM LUT | 54 |
| 3.48 | Sample waveforms showing voltage sequence necessary to write to and read from register LUT | 55 |

| | | |
|------|---|----|
| 4.1 | Low-swing driver schematics | 60 |
| 4.2 | Conventional low-swing receiver schematic | 60 |
| 4.3 | Pseudodifferential low-swing receiver schematic | 62 |
| 4.4 | Dynamic ratioed low-swing receiver schematic | 63 |
| 4.5 | Simulation testbench example for conventional driver, pseudodifferential receiver | 64 |
| 4.6 | Sample input and output for conventional-conventional transmission pair, low-swing voltage of 1.4 V | 66 |
| 4.7 | Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, conventional driver, conventional receiver | 67 |
| 4.8 | Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, conventional driver, pseudo-differential receiver | 67 |
| 4.9 | Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, conventional driver, dynamic ratioed receiver | 68 |
| 4.10 | Modified low-swing test bench circuit to account for NMOS driver V_{tn} drop in pull-up path | 69 |
| 4.11 | Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, NMOS driver, conventional receiver | 69 |
| 4.12 | Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, NMOS driver, pseudodifferential receiver | 70 |
| 4.13 | Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, NMOS driver, dynamic ratioed receiver | 71 |
| 4.14 | Average total power dissipated by driver, receiver and wire for 1 mm and 5 mm wires using typical-typical transistor models at 27°C | 71 |
| 4.15 | Average total power dissipated by driver, receiver and wire for 1 cm wire using typical-typical transistor models at 27°C | 72 |
| 4.16 | Pulse width modulator schematic | 73 |
| 4.17 | Pulse width modulator timing diagram | 74 |
| 4.18 | Pulse width demodulator schematic | 75 |
| 4.19 | Waveforms at the input and within the pulse width modulator showing timing of decoding operation | 76 |
| 4.20 | Pulse width modulation power compared to low-swing drivers and receivers at 10 Mb/s | 77 |
| A.1 | Single DRAM memory cell | 82 |

List of Tables

| | | |
|-----|---|----|
| 2.1 | Acceleration magnitude and fundamental frequency for common vibration sources [1] | 7 |
| 2.2 | Physical dimensions of piezoelectric generator for model [8] | 11 |
| 2.3 | Mechanical parameters for generator model | 11 |
| 2.4 | Calculated electrical parameters for equivalent circuit model | 12 |
| 3.1 | Time between nearest 400 <i>mV</i> crossings for rectified waveforms | 15 |
| 4.1 | Wire parameters used in interconnect simulations | 65 |
| 4.2 | Range of operational low-swing voltages for each driver and receiver pair . . | 70 |

Chapter 1

Introduction

The recent explosion of small scale, embedded sensing and processing electronics has made the growing disparity between battery and circuit technologies even more pronounced. Energy scavenging systems have shown to be a promising supplement or replacement for battery technology in many cases. This is particularly true in the case of sensing and monitoring systems unable to be easily connected to a power source and where a long lifetime in the field is desired, such as wireless sensor networks. In these situations, the harvesting system can convert energy from sunlight, mechanical vibrations, or heat into electrical energy that can in turn be used to power small electrical systems. However, new challenges come with these new opportunities to harvest power from previously unavailable sources. The power output of an energy scavenger can range from $15000 \mu W/cm^3$ from a solar cell outdoors in direct sun to $3 nW/cm^3$ from an acoustic noise harvester in the presence of 75 dB noise [1]. These power densities are reasonable when we consider that the power density of a rechargeable lithium battery is only $7 \mu W/cm^3$ for a one year lifetime, but the battery may have a size of several cubic centimeters whereas the harvester will likely be limited to only one cubic centimeter. This reduced peak power is offset by the harvester being able to produce the same power output throughout its life while the battery's power tapers off with time. Another challenge comes from the large variation in output power from the energy harvesters. As mentioned above, solar cells may output $15000 \mu W/cm^3$ outdoors on a sunny day, but their output may only be $6 \mu W/cm^3$ in an office setting.

Likewise we would expect a mechanical vibration based harvester mounted on a bridge to output completely different amounts of power depending on whether it was rush hour or the middle of the night. These variations in possible power present a strong argument for maintaining flexibility in the electrical system so that performance can be increased when the power is available and decreased to a base level or even turned off when it is not. Another technique to handle this variation in energy availability is to combine several scavenging mechanisms into a single power system for the electronics. Designing circuits around this level of variation while minimizing the energy wasted in complex power regulation circuitry requires a thorough understanding of the behavior of the energy scavengers as well as new ideas and methods of approaching circuit design in a low power yet highly flexible way.

1.1 Contributions of the Thesis

This thesis describes three primary contributions related to circuit and system design for energy harvesting systems in particular and low power design in general. First, a circuit model of an energy scavenging generator that can be incorporated into standard circuits simulations is developed. This model allows a circuit designer to test the feedback relationship between the electronic and mechanical systems in a design incorporating a piezoelectric generator. Understanding this relationship is necessary for systems where some or all of the power regulation electronics have been removed. Next is the development of a memory system designed to work with an AC power supply as well as a DC supply. This memory structure is a key component in designing energy scalable processors in an environment where regulated and/or constant power supplies are not available. Currently, a test chip is under development to implement this memory structure as well as several other circuits designed for energy harvesting AC power supplies. Finally, the exploration of several low-swing signaling and modulation techniques for long wires targets the general field of low power interconnect for circuit design. Each of these techniques is considered as a way of potentially reducing the power necessary to transmit data over long wires.

1.2 Organization of the Thesis

Chapter 2 contains a detailed discussion of the energy scavengers, specifically the mathematical models for the various mechanical devices. These models are then used to develop an equivalent circuit model that can be included in standard circuit simulations. Chapter 3 considers the design of a memory system capable of functioning with both an AC and a DC power supply. Architecture as well as circuit implementation of the memory are described, along with the use of the memory in the test chip currently being implemented. Chapter 4 explores two possible methods of reducing on chip interconnect power: reduced voltage swing on bus lines and pulse width modulation of data bits. Several low-swing driver and receiver circuits are presented, including transmission power and regions of operability. A possible architecture and circuit implementation for the pulse width modulation system are also investigated. Chapter 5 summarizes the key results of this work and briefly describes several possible areas of future work. A layout of the primary memory cell is shown in Appendix A.

Chapter 2

Energy Harvesting Generator Models

Energy harvesting generators have proven to be a viable option for powering electronics. They have been used in applications including powering an RF beacon for a low power radio [2] and providing power for a digital signal processor aimed at biomedical monitoring [3]. With the growing popularity of wireless sensor networks driving toward low-power, self-sufficient electronics, energy scavenging will only become more prevalent with time. One of the challenges for circuit designers trying to build electronics to take advantage of these generators is the disconnect between models for the energy harvesters and typical CAD tools for circuit design. Energy harvesters do not function like the typical power supply used in electronics design. The electronic circuit will provide extra damping to the harvester, which may in turn change the harvester's operation. If we compare this feedback loop to a modern processor, this would mean that every time you tried to run a large simulation on your computer the supply voltage to the chip would decrease and the chip would have less power to run, but when your computer was idle or doing an easy task such as reading email, the power supply would generate power relatively undamped and the chip would have plenty of power. While this would be bad for a desktop computer, it provides a unique opportunity for a standalone sensor node. If the behavior of the circuit can be designed correctly to match the mechanical generator, the system speed and power usage

will automatically adjust to the output power capability of the energy harvester. This feedback offers the circuit designer the option of eliminating most or all of the power electronics that could be used to regulate power to the load electronics. However, to correctly design circuits using these principles, the power supply behavior including the feedback must be taken into account. In this chapter a model for including energy harvesting generators in standard simulation tools for circuit design such as Cadence is described. This model allows the circuit designer to incorporate the structure of the energy harvester into the electronics design for more complete system development.

2.1 Mechanical Energy Harvesting Mechanism

Mechanical vibration-based energy harvesters are some of the most commonly used types of energy scavenging devices. They all follow the general model shown in Figure 2.1. The spring, mass, dampers and $z(t)$ measurement represent the oscillations internal to the generator while the $y(t)$ oscillations represent the movement of the macroscale vibration source on which the harvester is mounted. This system can be described by Equation 2.1 where z is the spring deflection, y is the input displacement, m is the mass, b_e and b_m represent electrical and mechanical damping respectively, and k is the spring constant of the mechanical system [1].

$$m\ddot{z} + (b_e + b_m)\dot{z} + kz = -m\ddot{y} \quad (2.1)$$

The electrical damping in this equation refers to the energy removed from the mechanical system by the electrical system. Thus, the power used by the electronics is represented as extra damping to the generator. This equation also shows how the changing load of the electrical system will affect the behavior of the mechanical system. As the loading of the electrical system changes, perhaps when the system leaves a wait state to begin a heavy computation, the damping will increase and the amplitude of the mechanical vibrations will decrease. An equivalent electrical system with this same behavior is shown in Figure 2.2. The parallel equation to Equation 2.1 for this system is $V_s = LC\ddot{V}_c + (R_m + R_e)C\dot{V}_c + V_c$.




Figure 2.1: Generic model for a mechanical vibration to electric energy converter [1]




Figure 2.2: Generic electrical model for a mechanical vibration system

Many possible vibration sources have been examined for use in energy harvesting systems. The two most important characteristics of the vibrations are the peak frequency of oscillation and the acceleration magnitude at that frequency. If the natural frequency of the oscillating portion of the scavenger is at or near the peak frequency of the vibration source, then more energy can be harvested. Thus the dimensions of the mechanical system can be optimized if the target vibration frequency is known. The acceleration magnitude is important because it directly determines how much mechanical energy is available in the system. This is equivalent to the $y(t)$ term in Figure 2.1 or the power supply V_s in the electrical analogy. Table 2.1 lists the acceleration and fundamental frequency for several vibration sources as characterized in [1].

| Vibration Source | Acceleration Magnitude (m/s^2) | Fundamental Frequency (Hz) |
|-----------------------------------|------------------------------------|--------------------------------|
| Car engine compartment | 12 | 200 |
| Base of 3-axis machine tool | 10 | 70 |
| Blender casing | 6.4 | 121 |
| Clothes dryer | 3.5 | 121 |
| Person tapping heel | 3 | 1 |
| Car instrument panel | 3 | 13 |
| Door frame after door closes | 3 | 125 |
| Small microwave oven | 2.5 | 121 |
| HVAC vents in office building | 0.2 – 1.5 | 60 |
| Windows next to busy road | 0.7 | 100 |
| CD on notebook computer | 0.6 | 75 |
| Second story floor of busy office | 0.2 | 100 |

Table 2.1: Acceleration magnitude and fundamental frequency for common vibration sources [1]

2.2 Types of Energy Harvesters

There are three primary types of energy scavengers. To date the most widely used scavenger is made from piezoelectric material. The piezoelectric material is mounted as a cantilever beam with electrodes plated onto some part of the beam. A shim can be added to the center of the beam to increase the stiffness and a mass can be added to change or increase the vibration response as well. A second style of energy harvester is

the electrostatic harvester, which is based on variable MEMS capacitors. In this system, charge is placed on the plates of a capacitor when they are close together and then removed when the plates have been pushed further apart by vibration. Because charge, capacitance, and voltage are related as $Q = CV$, when Q is held constant and C decreases, V must increase. Thus charge is removed from the capacitor at a higher potential and is then used to power a circuit [4]. There have been several styles of electrostatic harvesters proposed based on the different dimensional possibilities of fabricating the variable capacitors [4] [5]. The third type of harvester is based on electromagnetics and functions similar to a common loudspeaker. A magnet and a coil, one fixed and the other attached to a spring or flexible membrane, are necessary for this generator. Vibrations cause the coil to move relative to the magnet, causing different amounts of magnetic flux to be linked as a function of time. This creates a voltage on the coil that can then be used as a power supply.

Figure 2.3 shows several reported powers for all three types of energy harvesters [4], [5], [6], [7]. Many of these figures are reported powers based on simulation but a few are measured power from physical devices. We are primarily concerned with oscillations in the 60 Hz to 1 kHz range with amplitudes on the order of 2.5 m/s^2 . This limits us to devices in the $\leq 300\text{ }\mu\text{W}$ range.

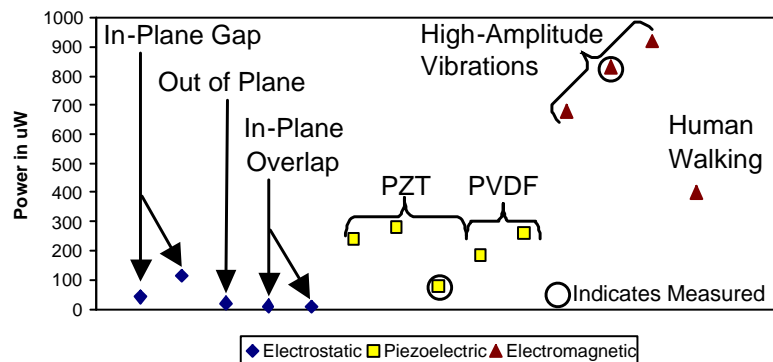


Figure 2.3: Reported power outputs for energy harvesters, both simulated and measured

2.3 Piezoelectric Model

Piezoelectric generators are currently the most widely available and most easily adaptable to power electronics. The basic structure of a piezoelectric generator is shown in Figure 2.4. The piezoelectric beam with effective length L_b extends from a fixed structure and has a mass mounted on the opposite end. Two layers of piezoelectric material form the beam and a stiff shim is shown between them, although the shim is not required for operation. Electrodes must be plated onto the piezoelectric material to collect the charge built up as the beam flexes.




Figure 2.4: Mechanical schematic drawing of piezoelectric generator

Equation 2.2 describes the operation of a piezoelectric generator [8]. In this equation, c_p is the elastic constant for the piezoelectric material, m is the mass, b_m is the mechanical damping coefficient, d_{31} is the strain coefficient in the direction of the strain, a indicates whether the piezo layers are wired in series or parallel, t_c is the thickness of a piezo layer, and \ddot{y} is the acceleration of the macro-scale physical system. Equations 2.3 and 2.4 describe geometric constants for the beam where b is the distance between the center of the piezo layer to the center of the shim, l_b is the length of the beam, l_m is the length of the mass, l_e is the length of the electrode, and I is the beam's effective moment of inertia.




Figure 2.5: Equivalent circuit model for piezoelectric generator

$$\ddot{S} = \frac{-c_p}{k_1 k_2 m} S - \frac{b_m}{m} \dot{S} + \frac{c_p d_{31} a}{2k_1 k_2 m t_{pz}} V + \frac{\ddot{y}}{k_2} \quad (2.2)$$

$$k_1 = \frac{b(2l_b + l_m - l_e)}{2I} \quad (2.3)$$

$$k_2 = \frac{l_b^2 (2l_b + \frac{3}{2}l_m)}{3b(2l_b + l_m - l_e)} \quad (2.4)$$

$$I = 2\left[\frac{wt_{pz}^3}{12} + wt_{pz}b^2\right] + \frac{\eta_s wt_{sh}^3}{12} \quad (2.5)$$

The format of this equation matches that of the series RLC circuit shown in Figure 2.5, which can be described by Equation 2.6. In this equation the parallel load resistance R_e and capacitance C_e are not explicitly included, though they are accounted for through the term V representing the voltage across them.

$$\ddot{q} = \frac{-1}{L_m C_k} q - \frac{R_b}{L_m} \dot{q} + \frac{n}{L_m} V + \frac{V_{in}}{L_m} \quad (2.6)$$

By equating the like terms in the two equations and solving to find the electrical parameters we arrive at the relations below.

$$L_m = k_1 k_2 m \quad (2.7)$$

$$C_k = \frac{1}{c_p} \quad (2.8)$$

$$R_b = k_1 k_2 b_m \quad (2.9)$$

$$n = \frac{-c_p d_{31} a}{2t_{pz}} \quad (2.10)$$

$$V_{in} = k_1 m \ddot{y} \quad (2.11)$$

These relations can be used to build an equivalent electrical model for the piezoelectric generator. The physical sizes of the generator are based on dimensions provided in [8] and are shown in Table 2.2.

| Physical Dimensions (mm) | l_m | h_m | w_m | l_b | l_e | w | t_{pz} | t_{sh} |
|--------------------------|-------|-------|-------|-------|-------|-----|----------|----------|
| | 8.5 | 7.7 | 6.7 | 6.5 | 6.5 | 3.2 | 0.14 | 0.1 |

Table 2.2: Physical dimensions of piezoelectric generator for model [8]

Several additional mechanical parameters are needed to calculate the electrical modeling values. These include the elastic constant for both the piezoelectric material and the shim, the density of the mass material, the mechanical damping coefficient b_m , and the strain coefficient d_{31} . The generator discussed in [8] uses PZT-5H as the piezoelectric material and brass as the center shim, with a Tungsten mass. Piezo Systems Inc. manufactured the piezo-shim beam. Several of the necessary parameters were provided in [8], [9] or [10]. A few parameters were not directly specified, but were able to be calculated using the relations $k_{31} = \frac{d_{31}^2 c_p}{\epsilon_r \epsilon_0}$ and $2\zeta\omega_n = \frac{b_m}{m}$, as well as the provided values. The piezoelectric layers are assumed to be connected in series, making a equal to 1. The beam's effective moment of inertia I is $1.5 \times 10^{-14} \text{ m}^4$, dimension parameter k_1 is $60 \times 10^6 \text{ m}^{-2}$, and dimension parameter k_2 is 0.201 m . All other parameters are shown in Table 2.3.

| Mechanical Parameters | | |
|--------------------------|--------------|--|
| Elastic Constants | c_p | $4.6 \times 10^{10} \text{ N/m}^2$ [8] [9] |
| | c_{sh} | $11 \times 10^{10} \text{ N/m}^2$ [10] |
| Elasticity Ratio | η_s | 2.39 |
| Coupling Coefficient | k_{31} | 0.14 [8] |
| Dielectric Coefficient | ϵ_r | 3800 [9] |
| Strain Coefficient | d_{31} | $-320 \times 10^{-12} \text{ m/V}$ [9] |
| Damping Ratio | ζ | 0.015 |
| Damping Coefficient | b_m | 0.03 [8] |
| Tungsten Density | W_{dens} | 19250 kg/m^3 |
| Mass | m | 8.44 gm |
| Natural/Driven Frequency | ω_n | 120 Hz |
| Acceleration Magnitude | \ddot{y} | 2.5 m/s^2 |

Table 2.3: Mechanical parameters for generator model

Using these parameters and the relations developed previously, values for the equivalent electrical model were calculated. The resulting values are provided in Table 2.4.

| Calculated | L_m | C_k | R_b | n | V_{in} |
|--------------------------|----------|-----------|---------------|----------|-----------|
| Electrical Parameters | 102 kH | 21.7 pF | 362 $k\Omega$ | 52.6 k | 1.27 MV |

Table 2.4: Calculated electrical parameters for equivalent circuit model

Simulations indicate that the model's qualitative behavior is consistent with previous experiments described in the literature. Future work will involve calibrating the model against physical data from a piezoelectric generator and simulating the generator model with the test chip as the load. The anticipated electrical load is approximately 300 $k\Omega$ [11] of resistance in parallel with 9.2 nF [8] of capacitance from the generator and the final test chip capacitance.

Chapter 3

Low-Power Memory for AC Power Supplies

Memory is an important piece of nearly any system design. Ideally it should function transparently, allowing information to be shifted around, accessed, and stored without slowing down other system functions or using much of the system's power budget. However, in more and more digital systems memory circuitry is responsible for an increasingly large amount of the power usage. For the system we chose to build, reducing this memory power was an integral part of the design. Aside from having relatively low power, our memory also needed to be able to function with a varying power supply voltage, even a power supply that falls to ground for some amount of time. In this chapter we will describe the design of the memory system chosen to meet these requirements. Section 3.1 will describe the overall memory architecture, including the Distributed Arithmetic system being implemented and its memory needs, as well as the controller needed to successfully read and write these particular memories. Section 3.2 discusses the circuits used to implement the memory architecture. The final section looks at the full memory system.

3.1 Memory Architecture

There are two unique concerns in designing this memory system. The first is that the power available is on the order of tens to hundreds of microwatts. Having such

a low power budget limits the inclusion of typical analog sense amplifiers because of their prohibitively large power usage. Likewise, our power budget also discourages using any memory cells that require frequent refreshing.

The second and most pressing concern involves use of an AC power supply. We are concerned with supply oscillations on the order of 60 Hz to 1 kHz where the electronics are functioning when V_{DD} is between approximately 400 mV and 1.8 V . In this range, the power supply may be below an operational level for up to 1.2 ms . This time indicates the length of time between a 400 mV falling and rising edge at the lowest possible oscillation frequency, 60 Hz . Figures 3.1(a) and 3.1(b) illustrate this waveform as well as the waveform showing the shortest time the memory will have to hold its value.


(a) Rectified 60 Hz power supply wave(b) Rectified 1 KHz power supply wave

Figure 3.1: Highest and lowest mechanical vibration frequencies explicitly investigated, arrows indicate time where supply is unusable

While the power supply is below a usable voltage, the memory system must be able to retain its values until the voltage rises again and the system continues operating normally. Data being processed should not be lost, and more importantly, filtering coefficients or system configuration parameters must be preserved. The required hold times for several frequencies in the range of interest are shown in Table 3.1. Despite this absolute need to maintain some system variables, using only ROM was undesirable because of the inherent

permanence of the values it stores. While this may seem counterintuitive at first glance, one of the benefits of using the system we propose is its ability to adapt and scale both coefficients and word lengths when energy scalable operation is desired. Thus the system can adapt to both low and high energy modes of operation. Details of the scalable system are considered in the next section which describes distributed arithmetic.

| Frequency | Hold Time |
|---------------|--------------------|
| 60 <i>Hz</i> | 1.2 <i>ms</i> |
| 120 <i>Hz</i> | 0.6 <i>ms</i> |
| 300 <i>Hz</i> | 240 μ <i>s</i> |
| 1 <i>KHz</i> | 72 μ <i>s</i> |

Table 3.1: Time between nearest 400 *mV* crossings for rectified waveforms

3.1.1 Distributed Arithmetic

Distributed arithmetic (DA) is a bit-serial method of computing the dot product of a constant vector and a variable vector [12]. It works by precomputing the outputs using each value of the variable vector and storing them in a lookup table (LUT). The variable vector is then used to address the lookup table and the memory output provides the correct dot product of the two vectors. This operation is particularly useful for implementing energy scalable filters for use in digital signal processing (DSP) [13] [14]. Because our system is targeted for not only low power but also variable power sensing applications, DA is a natural choice. We chose to implement a 16–tap FIR filter using a DA system. This filter length was selected to build on a test chip to allow a reasonable amount of circuit complexity for testing several interesting new circuits. Because DA requires precomputation of all of the possible vector combinations, building this filter using a single LUT would require 256 16–bit words to be stored. A memory this size would be unnecessarily large, so instead we chose to build four 4–tap filters and use three additional adders to accumulate their results. A block diagram of the system is shown in Figure 3.2.

New data enters the Data 1 block and is shifted through the four data blocks in the first DA unit. The data being shifted out of the first DA unit flows into and through




Figure 3.2: 16-Tap DA Filter Block Diagram

the subsequent DA units in the same fashion. Outputs from each of the DA units are added together in an adder tree to provide the full filter output.

A more detailed view of a single DA unit is shown in Figure 3.3. The LUT is illustrated as a table of values written in terms of the fixed vector A , which can be thought of as filter coefficients. These values are indexed by the variable vectors shifting through the data memory blocks. The bit length of the shifted data determines the number of iterations required to have a valid output. In our system the nominal data width is sixteen bits. Thus the DA output can be ignored for fifteen of every sixteen clock cycles. This provides an opportunity to reduce the system power and only allow the adder inputs to update every sixteen clock cycles. More details about the adder and clocking design can be found in [11].

There are two types of memory necessary to build a DA unit. The data memory must shift data around and through the DA units. In our system this was built with variable length shift register chains. An SRAM memory structure has also been proposed in other work for this purpose. Details of both designs can be found in [14]. The second memory

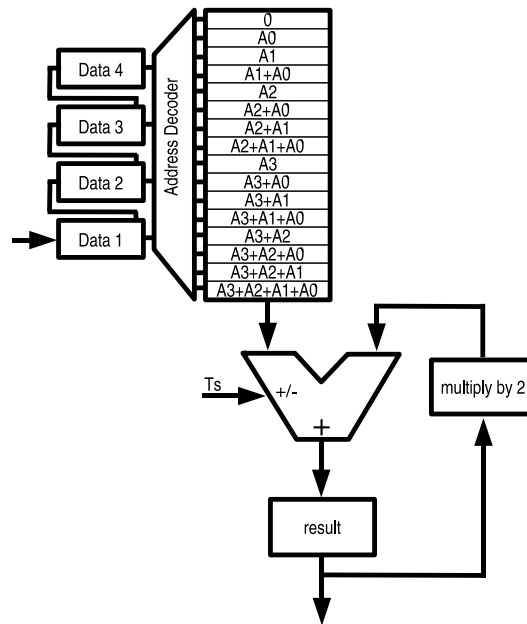


Figure 3.3: Detailed view of DA unit

block needed must implement the LUTs. This will be the subject of the rest of this chapter, including the architectural, circuit, and controller design.

3.1.2 DA Look-Up Table

Two types of memory were used to implement the lookup tables (LUTs) for this system. The first was an array of DRAM cells which were sized to hold their value without refreshing for up to 1.2 *ms*. This memory array is designed to work in cases where only an AC supply is present as well as when a DC supply is available. The second memory system implemented was an array of CMOS registers. These were chosen to be used in two of the LUTs because the memory cells were more robust and more likely to work correctly in case the DRAM array did not operate as planned.

Each LUT was implemented as a 16×16 array of memory cells. Wordlines are shared horizontally and bitlines vertically. Incoming addresses are in four-bit binary format and are decoded into sixteen-bit one-hot format such that a single wordline is active at any given time. The DRAM LUT requires separate read and write signals that are created by

ANDing the wordline with read and write enables. These lines are labeled $WL_n \bullet RdEn$ and $WL_n \bullet WrEn$ in Figure 3.4(a). Both data to be written to the memory and values read out of the DRAM use the same physical bitline, labeled $Data_n$ in Figure 3.4(a).

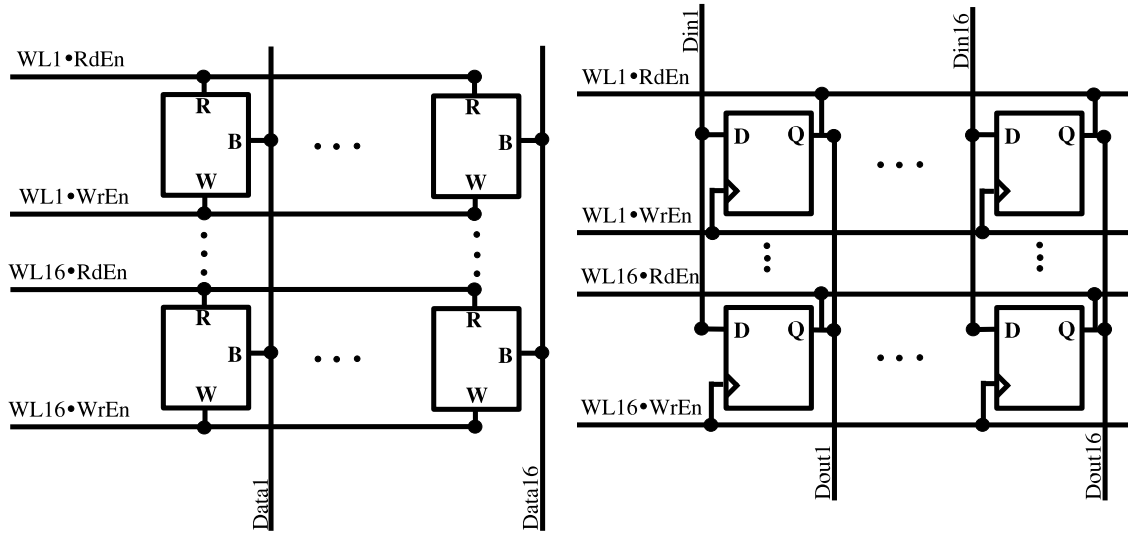
The DRAM is read and written through a bitline sense circuit that connects to the bitline. Data is written to the sense circuit and driven onto the bitline when the write enable signal is asserted, taking a single clock cycle to complete. Read operations use two phases: a precharge followed by a read enable. The bitline is precharged high during the precharge phase and then evaluated when read enable is asserted. Precharge and read enable are each asserted for half of a clock period.

In the register LUT the clock lines are shared horizontally and controlled by the AND of wordline and write enable, labeled as $WL_n \bullet WrEn$ in Figure 3.4(b). Data being written to the cells is asserted on the Din_n lines which connect to the D pin of the flip-flop. The Din_n line is shared vertically in the array. The values stored in the cells are read out through the $Dout_n$ lines, which are also shared vertically. This line is driven by the Q output of the flip-flop. However, because the Q outputs are actively driven, they cannot all be connected simply to the $Dout_n$ line. Transmission gates controlled by the AND of read enable and wordline must be used to connect the output being read to the $Dout_n$ line and disconnect the other outputs. This control signal is indicated in Figure 3.4(b) by the label $WL_n \bullet RdEn$. Bitline sharing is not a concern for the DRAM because the line is not actively driven by the memory cells.

A single controller is used to write and refresh the four LUTs and will be considered in the next section.

3.1.3 Controller for LUTs

The LUT controller is responsible for writing the correct coefficient values to the memory cells of the LUTs. All cells need to be written from external memory when the system first powers up. This operation is indicated by an initialize signal. After every power supply cycle, the register LUTs must be rewritten from the external memory and the DRAM cells must be read and rewritten. The DRAM must also be refreshed periodically, although DRAM reads are non-destructive, allowing for longer operation between refreshes.



(a) DRAM lookup table design

(b) Register lookup table design

Figure 3.4: Lookup table block diagrams

To simplify the controller design, the functions were implemented as five finite state machines (FSMs). Three of the machines, Master, Read, and Write, were implemented as Moore machines and the other two were implemented as multibit counters. The overall state of the system is controlled by the Master FSM, which has five states: “read”, “read wait”, “write”, “write wait”, and “wait”. Figure 3.5 illustrates the state diagram.

The read state was made the zero state so that clearing the system registers would automatically trigger the controller to write all of the LUTs. In the read state the read output is asserted. This output is responsible for starting the Read FSM. While the Read FSM is operating, the Master FSM sits in the read wait state. Once the Read FSM has completed its task, it asserts the readEnd input to the Master FSM and the Master moves to the write state. Now the write output is asserted to start the Write FSM and the Master moves into the write wait state. When writeEnd is asserted, indicating that the Write FSM is finished, the Master moves back to the read state if there are more memory cells to write or moves to the wait state if all of the cells have been written. Completion of the full write is indicated by the updateEnd input. The clear output is asserted continuously

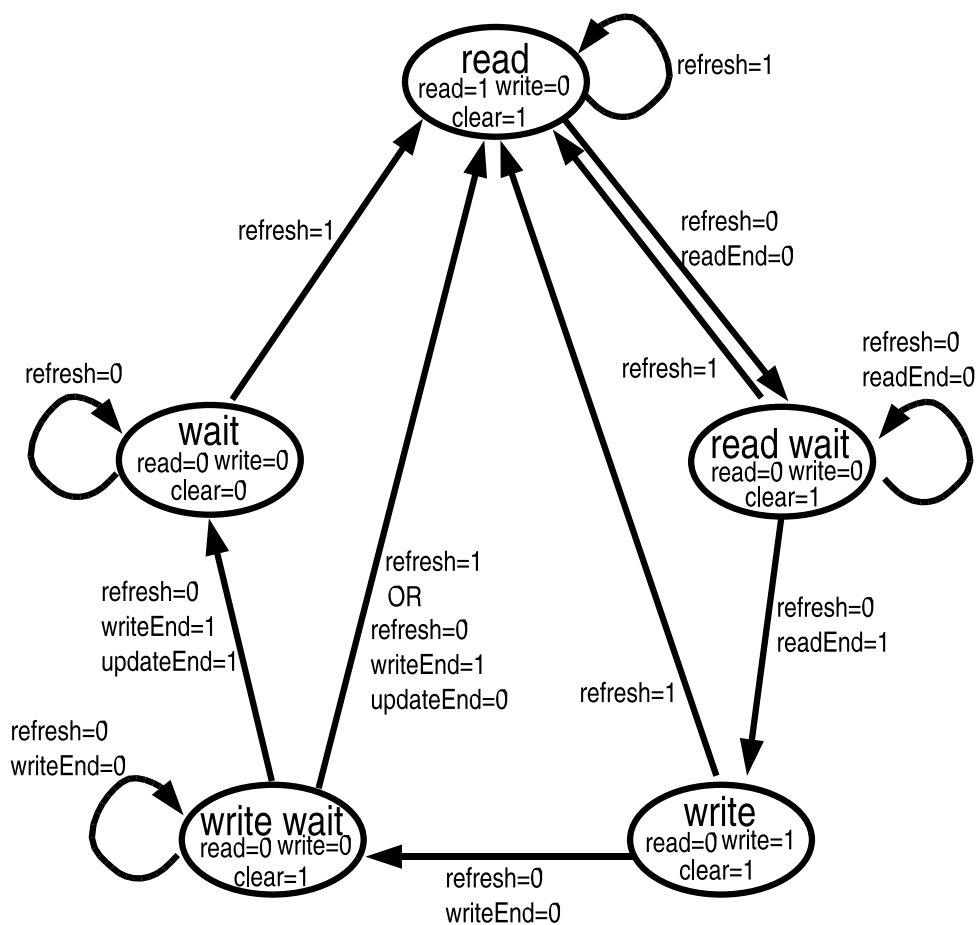


Figure 3.5: Master FSM state diagram

when the Master is in any state other than wait. This is necessary to gate the clocks to the data shift registers and prevent them from trying to access the LUTs while the controller is writing the coefficients. In the state diagram, the signal labeled refresh is a generic signal to indicate any of the following actions: initialize the memory, reset after power supply cycles, or refresh the DRAM cells. The differences between these operations are handled in the other FSMs and will be described next.

The Read FSM is the first state machine called by the Master FSM. It has five states, three inputs, and three outputs. The inputs are a start signal from the Master FSM, a dram signal indicating when the DRAM LUTs are being written, and an initialize signal, which is asserted during the system's initialization. A state diagram of the Read FSM is shown in Figure 3.6.

Wait is the zero state of the FSM. After the Master FSM pushes the Read FSM out of the wait state, the next state is determined by the dram input. If a DRAM LUT is being written then the next state is precharge, whereas the next state is read exterior when a register LUT is being written. The precharge state asserts the precharge output, which tells the DRAM sense circuits to precharge the bitlines when DRAM reads occur. From the precharge state, the FSM transitions into either read external or read internal based on whether the system is initializing or not. Note that the precharge is not required when the DRAM is being written but the values are coming from register memory. Thus, in the initialize state when all of the coefficients are being initialized from separate register memory, the precharge is not necessary and can be thought of as a dead state. However, this state was not removed from the Read FSM because of the possibility of reading data from a possible secondary DRAM memory or other memory using the same cell sense circuit. When in either the read external or read internal states the read enable output is asserted. This output is used to write 16 bits of data from a selected memory location into 16 static latches internal to the LUT controller. Once the data read is completed, read enable is deasserted and the FSM moves to the read end state. In this state a signal is sent back to the Master FSM to indicate the read is completed. Finally, the Read FSM moves back into the wait state and awaits another call by the Master FSM.

The next FSM called by the Master is the Write FSM. It has four states, one input,

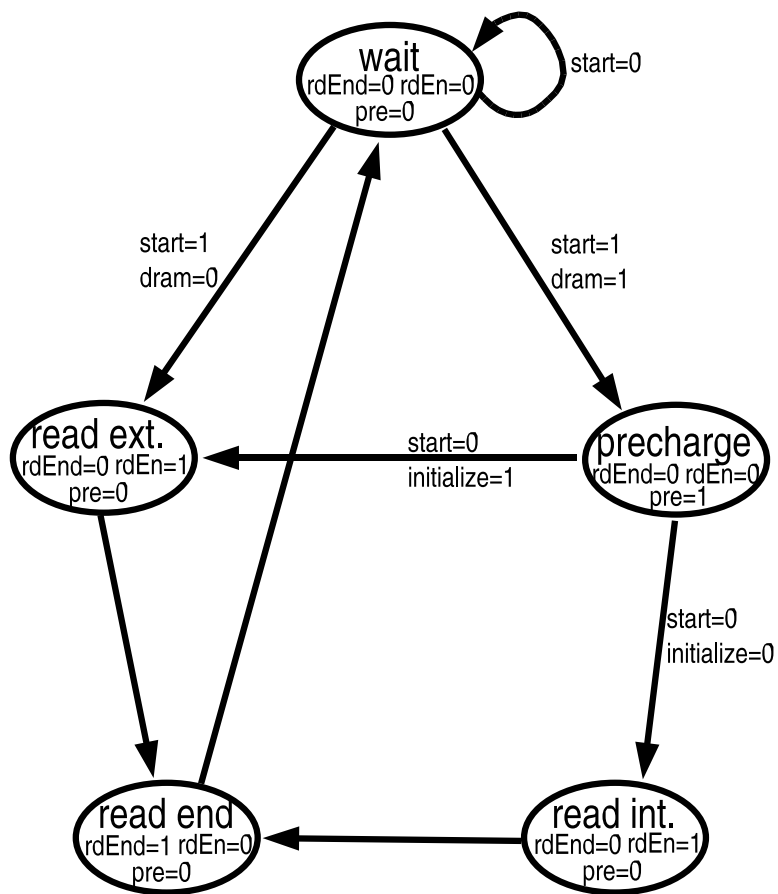


Figure 3.6: Read FSM state diagram

and three outputs. A start signal from the Master is the only input. The three outputs are responsible for enabling the write to the LUT memory cells, incrementing the counter that stores the address and LUT select bits, and signaling back to the Master FSM that operation is complete. This is illustrated by the state diagram in Figure 3.7.

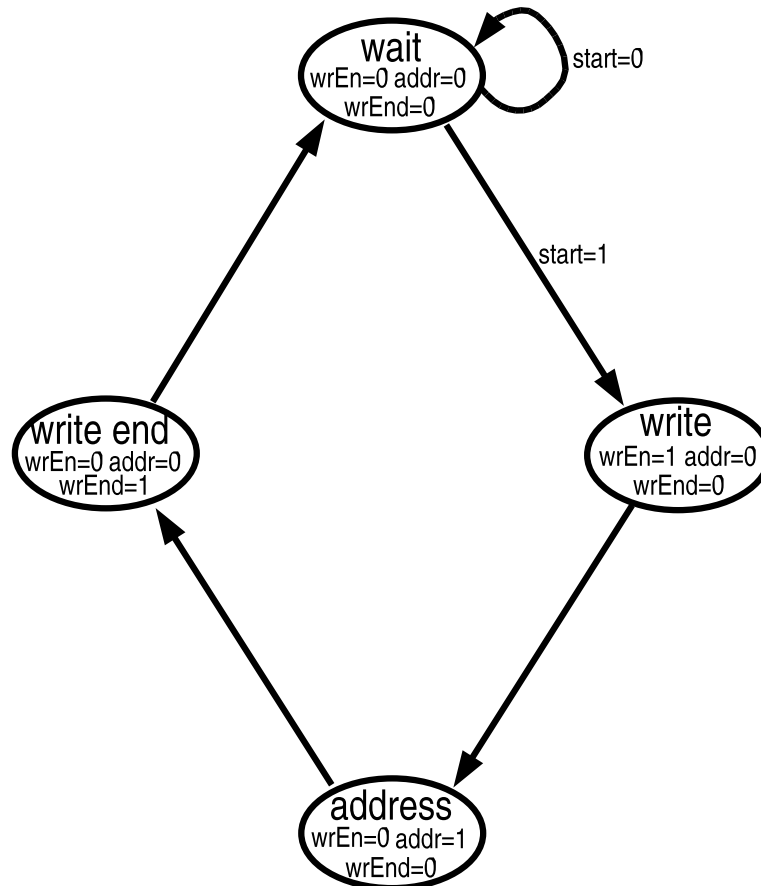


Figure 3.7: Write FSM state diagram

Similar to the Read FSM, this state machine starts in the wait state. Upon receiving a start pulse from the Master, the FSM moves into the write state. While in this state the write enable output is asserted and the values that were previously read into static latches in the controller by the Read FSM are written to the LUT in the selected DA unit. The state machine then moves to the address state and asserts the address output. This

output increments the seven bit counter that provides the address, LUT select, and update finished bits. The write end state follows the address state and signals the Master FSM that the write is finished before moving back into the wait state.

Addressing, LUT selection, and indicating when all LUTs have been written is done by a seven bit counter as mentioned above. The four LSBs, bits zero through three, make up the address used by the LUT to select a row. Bits four and five select which LUT to write and read from. Bit six, the MSB, indicates when all of the four LUTs have been updated. This MSB is used as the update end input to the Master FSM and is responsible for moving the Master into the wait state and deasserting the clock hold signal used by the data shift memory. The MSB is then cleared by the released clock hold signal. At this point the data shift memory begins clocking data through the DA unit and the entire system functions as a 16-tap FIR filter.

Because the DRAM must be refreshed, a second counter is necessary. The clock to this counter is held while the LUT controller is writing each of the LUT cells and the counter is cleared when the LUT update finishes and the DA resumes normal operation. While the DA runs, the counter is clocked by the system clock. Once a specified number of clock cycles passes, the counter output is used to indicate a refresh to the Master FSM. This initiates a full LUT update and again holds the clock for the data shift memory. The number of clock cycles between refresh values can be easily changed by the size of the counter or the combinational logic used to determine which count value should start the refresh. In the current system, the refresh happens every 16 clock cycles. This number was chosen to allow a full data word to enter the DA between each refresh and as a precaution to prevent the DRAM from discharging. In future implementations, it might be useful to implement this refresh adaptively so that the cells are refreshed often when the power supply is low and leakage time is a concern and prolong the refresh times when the power supply is high and leakage is unnoticeable. Details of a potential implementation of this will be considered further in the future work chapter.

3.2 Memory Circuits

3.2.1 DRAM Cells for AC Supply

The fundamental DRAM cell implemented consists of three transistors and is shown in Figure 3.8(a) [15]. Charge is stored on the gate of transistor M0 labeled in the schematic. Writing of the storage node occurs through transistor M2, which has its drain connected to the bit wire and its source connected to the store node. The gate of this transistor is asserted and the data to be stored is driven onto the bitline. To read the stored value, the bitline is first precharged high and then the gate of transistor M1 is asserted. If the storage node is high then both transistor M1 and the storage transistor are on and the bitline charge is cleared to ground. If the storage node is low, then the storage transistor is off and the bitline remains charged. Thus the stored data can be read by inverting the value of the bitline during a read operation.

The two largest design decisions for this DRAM cell are how to implement the write path and how to size the storage transistor. Two possible write paths are shown in Figure 3.8. Both circuits use only NMOS transistors in the write path, which limits the voltage on the storage node to the bitline voltage minus a threshold drop. The four transistor design restricts the write voltage of the storage node to a lower value than the three transistor design, but is expected to have lower charge leakage and thus be able to maintain the charge for a longer time. Leakage may also be reduced by increasing the length of the write transistors at the expense of lower storage node voltage. Each of these write configurations was simulated with the minimum length of 2λ as well as a length of 3λ . The results will be described later.

Several other write path implementations are also possible. Two examples are shown in Figure 3.9. Both of these circuits use transmission gates to remove or reduce the threshold drop between the bitline and the storage node. Two devices are used in series to reduce leakage and balance the additional leakage path provided by the PMOS transistors in series with the NMOS transistors. One problem with using two transmission gates in the write path is that when V_{DD} falls to ground, both of the PMOS transistors will be on if the storage node is charged high. In that case the storage node will act as the source and charge

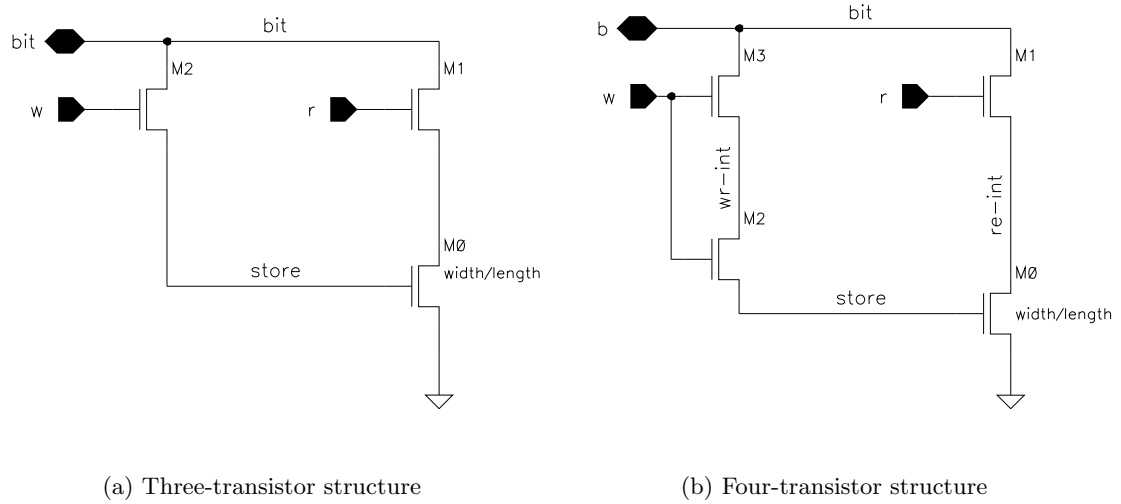


Figure 3.8: DRAM cell with NMOS write path

will drain out through the PMOS transistors. The hybrid NMOS-transmission gate cell was considered as an alternative where the NMOS transistor will decrease the cell leakage. The full transmission gate cell was simulated as well as the hybrid cell with the NMOS transistor in either the position connected to the storage node or the position connected to the bitline for comparison and will be described later.

The final two memory cells examined are shown in Figure 3.10. The cell in Figure 3.10(a) uses a high threshold transistor in the write path. These transistors are designed to work with 3.3 V power supplies and thus have much thicker gate oxide. For a given supply voltage, the high threshold transistor will have less current than a standard transistor, and for low power supplies it will operate fully in subthreshold. This reduces the leakage current through the write path significantly. However, it also makes writing the storage node high much more difficult. The memory cell in Figure 3.10(b) uses standard NMOS devices in the write path but drives their bulk voltage below ground. By lowering the bulk voltage, V_{SB} is increased for the transistor and the threshold voltage increases as well. The increased threshold voltage reduces the leakage current but again makes the cell more difficult to write. Implementing either of these circuits would require an advanced process with p-wells or complicated layout to fit in two oxide thicknesses while satisfying the design

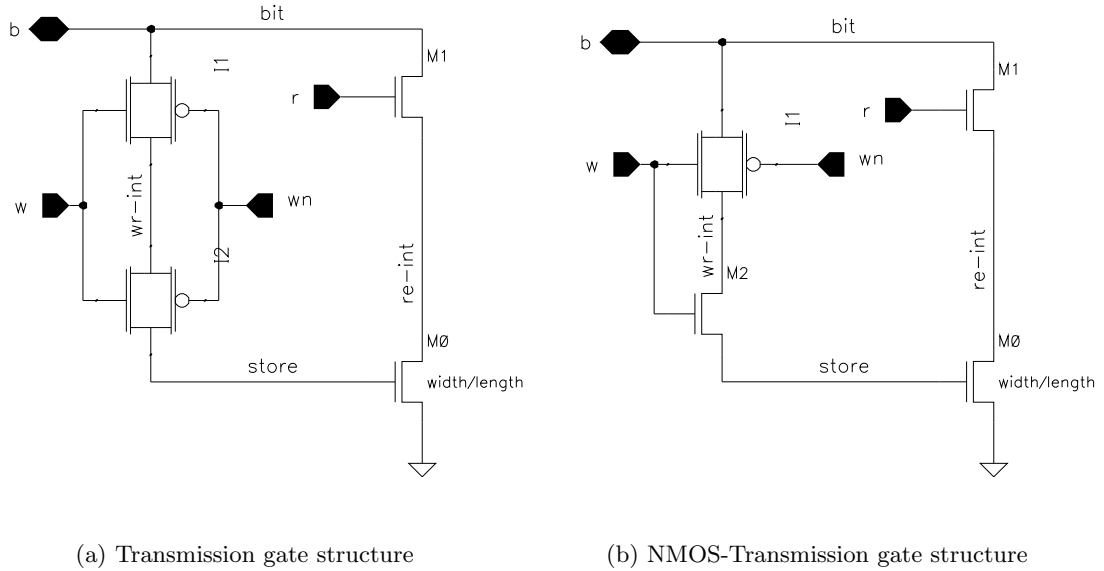


Figure 3.9: DRAM cell with transmission gates in write path

rules. This cost is unreasonable for the current design. However, the cells were simulated for comparison. Simulation results will be described later.

Before each of the cells described above could be tested, a preliminary estimation of the size of the storage node had to be made. The write path was built as two series NMOS transistors as shown in Figure 3.8(b) with length of 2λ . This setup was chosen because it was expected to have average leakage and ability to write the storage node high. Using this cell, the storage transistor width and length were set equal and swept from 450 nm to $19.35\text{ }\mu\text{m}$ in increments of 10λ or $0.9\text{ }\mu\text{m}$. For each storage transistor size, a high voltage was written to the node and the time for the maximum voltage on the storage node to fall 20% was measured. This was done using typical-typical transistor models with a V_{DD} of 400 mV and a write time of $2.7\text{ }\mu\text{s}$. These conditions were selected to emulate a write just before the supply voltage drops below a usable value [11]. Temperature was swept from 0°C to 85°C in 5°C increments as well. The simulation results are shown in Figure 3.11. A line indicating the hold time for the 60 Hz power supply frequency is also drawn to show the worst-case required time for the storage node to maintain its value.

Simulation results show that the smallest storage node necessary to stay within

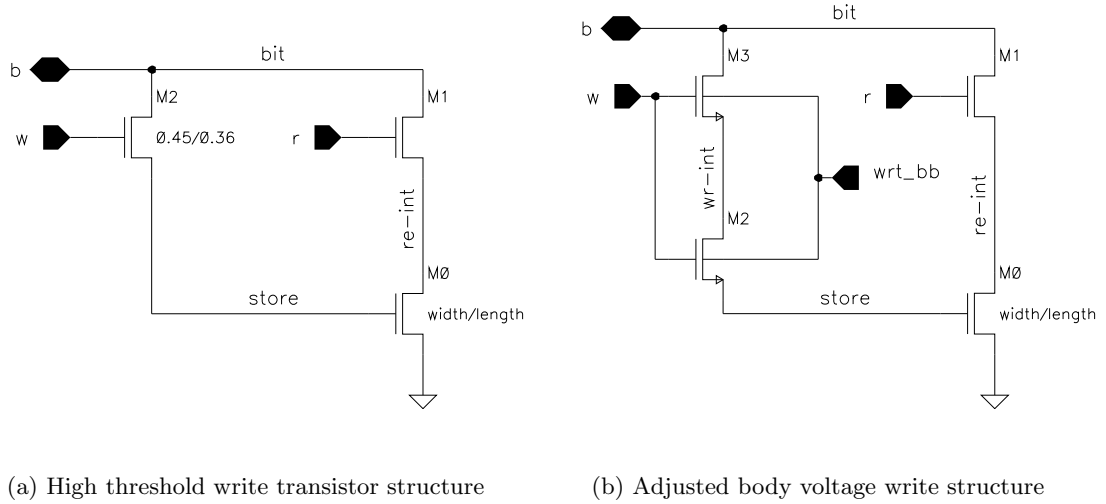


Figure 3.10: DRAM cells utilizing alternative threshold voltages

20% of the written voltage at all temperatures has an area of $147.6 \mu\text{m}^2$ or over $18000\lambda^2$. This is quite large, even compared to one of the largest cells in the design, the static flip-flop, which has an area less than $16000\lambda^2$. Storage node area this large is neither practical nor acceptable for our system. To reduce this size, we restricted the temperature range to between 0°C and 50°C . Considering that one of our primary applications would be to use these circuits in wireless sensor nodes, this reduced temperature range is reasonable. In addition, the hold time required for a power supply frequency of 1 kHz is only $72 \mu\text{s}$, and that of the 300 Hz power supply is only $240 \mu\text{s}$. Thus, higher temperatures can be sustained if the power supply frequency is increased.

If we limit our temperature range to values below 50°C , the smallest cell size that satisfies the worst case hold time constraint has an area of $46 \mu\text{m}^2$ or $5625\lambda^2$. Keeping this storage node area, the length was set to $9.9 \mu\text{m}$ or 110λ , making the width $4.6 \mu\text{m}$ or 51λ . The length was chosen to be longer than the width to prevent the transistor from turning on due to charge coupling from the read transistor onto the storage node when the node was holding a zero. The cost of increasing the transistor length is that when the cell is holding a one, it will take longer to discharge the bitline and read the one out. A value of $9.9 \mu\text{m}$ was chosen for the length because this value was close to one of the dimensions

of the static flip-flop, which would be used to build the register memory. This dimension matching is not necessary, but was chosen somewhat arbitrarily to simplify the large scale system layout. At this size, the storage node transistor has an equivalent capacitance of approximately 300 fF . This also accounts for a minimum sized source diffusion capacitance to ground, which is added by the write transistor touching the storage node.





Figure 3.11: Twenty percent fall times for various storage node sizes over temperature

Using the $9.9\ \mu\text{m} \times 4.6\ \mu\text{m}$ storage node, a test bench was built to look at the performance of each of the memory cells. The test bench schematic is shown in Figure 3.12. Each cell is tested by writing a one to the storage node and reading the cell some time later as well as by writing a zero to the storage node and reading the cell at a delayed time. The power supply voltage V_{DD} is swept from 1.8 V to 400 mV and the delay between write and read times is swept to match the hold times for power supply frequencies of 60 Hz , 120 Hz , 300 Hz , and 1 KHz as shown in Table 3.1. Two sets of simulations are run using this

bitline, and storage node voltages during a write operation, where the value being written is a one. Note that the bitline is cleared following the write to test the worst case leakage out of the storage node. The corresponding read of the storage node through the DRAM sense circuit is shown in Figure 3.13(b). Here we see the precharging of the bitline followed by the read enable pulse. During the read enable pulse the read output of the sense circuit rises to indicate a one was stored on the node.




(a) Write one to storage node

(b) Read one at output of sense circuit

Figure 3.13: Command and internal voltages for writing a 1 to the storage cell and reading out the 1 after a delay.

Figure 3.14(a) shows the same write sequence for writing a zero to the storage node. The node starts high and is discharged through the bitline when the write enable pulse occurs. Following the write the bitline is charged to again account for the worst case leakage. Figure 3.14(b) shows the cell being read. The bitline is precharged and the read enable pulse occurs. However, now the bitline is not fully cleared and the sense circuit outputs a zero. Although the bitline should hold at approximately V_{DD} when read enable is asserted, we see that it falls by nearly a volt. This drop is caused by charge sharing between the bitline and the drain diffusion of the storage transistor. In this simulation, only a single memory cell is connected to the bitline and schematic-level views are used rather than extracted layouts, which would include wire and other parasitic capacitances. Adding the other memory cells and the wire capacitance to the bitline should reduce the

magnitude of this voltage drop. Precharging the drain diffusion capacitance of the storage node before asserting the read enable would also reduce or potentially eliminate this voltage drop. By turning on the read transistor during the precharge cycle we can precharge the diffusion capacitance to $V_{DD} - V_{tn}$. This can be implemented by allowing the word address and either precharge or read enable to control the wordline select for the row of cells. This modification will be included in the test chip. These sample waveforms were taken from the typical-typical simulation of the memory cell with a single, minimum length write transistor with temperature of $27^{\circ}C$, V_{DD} of $1.2 V$ and a delay time between write and read of $600 \mu s$, corresponding to a power supply frequency of $120 Hz$.




(a) Write zero to storage node

(b) Read zero at output of sense circuit

Figure 3.14: Command and internal voltages for writing a 0 to the storage cell and reading out the 0 after a delay.

The first cells examined are the four variations of the NMOS only write path. Figures 3.15 and 3.16 show the read one and zero values for the single write transistor cell with length of $180 nm$. Figure 3.15 shows that the cell can safely read a one over temperature and process variation for $V_{DD} \geq 1.2 V$. However, as V_{DD} falls to $1.0 V$ and below, nearly all of the cells fail to correctly store and read out a one regardless of the required hold time. This is true for all temperatures and all processes except for the slow-NMOS, fast-PMOS process corner. Figure 3.16 shows that in nearly all cases the cell can

correctly store and output a zero. Here the slow-NMOS, fast-PMOS process corner is the only one that gives an incorrect read output. This happens when V_{DD} is 1.8 V, giving a read output voltage value around 400 mV. It is possible that this voltage may be read as a zero into the following stage. However, it is around the NMOS threshold voltage which may cause a problem.




(a) Read one, temperature variation

(b) Read one, process variation


Figure 3.15: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , single write transistor, length of 180 nm

Figures 3.17 and 3.18 show the simulation results for the single write transistor cell with 270 nm length. These plots show similar behavior to the plots of read values for the 180 nm length write transistor. The greatest difference is that more cases fail when the length increases to 270 nm. This indicates that the higher storage node voltage may be more valuable than a lower leakage current.

Figures 3.19 and 3.20 show the read out values for a memory cell with two series write transistors where each of their lengths are 180 nm. Again these results show similar characteristics to both of the single NMOS write transistor simulations. However, in these simulations the output is often incorrect. When reading a one the cell fails at even higher supply voltages than previously seen, and now some failures occur at the highest supply voltage. Reading a zero value is more consistent than reading a one, but more high voltages fail in this simulation compared to the previous results as well. This set of results shows the




(a) Read zero, temperature variation

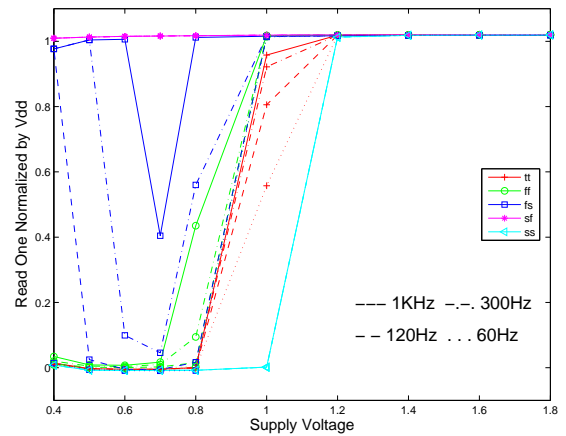


(b) Read zero, process variation

Figure 3.16: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , single write transistor, length of 180 nm



(a) Read one, temperature variation



(b) Read one, process variation

Figure 3.17: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , single write transistor, length of 270 nm

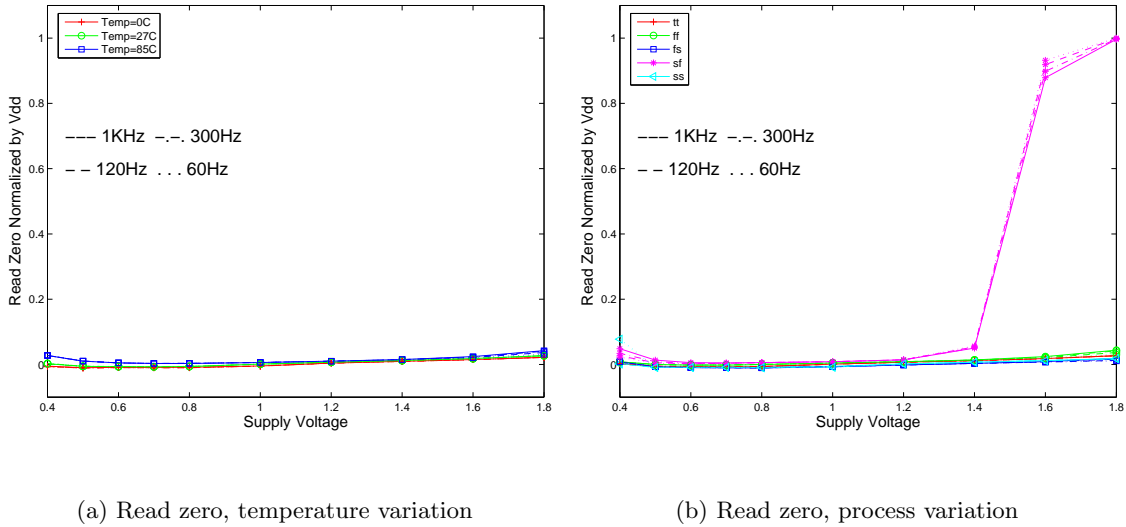



Figure 3.18: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , single write transistor, length of 270 nm

first failure for the typical-typical simulation when temperature is being swept. All of these results strengthen the evidence that increasing the peak write voltage is more important than reducing the leakage current.

Simulation results from the final NMOS write cell are shown in Figures 3.21 and 3.22. This cell has two write transistors in series, each with a length of 270 nm. The read outputs when a one was written to the cell are again worse than those shown previously, as are the read outputs when a zero was written.


The next set of memory cells to be simulated were the cells which used transmission gates. Figures 3.23 and 3.24 show the read output values for the memory cell with two transmission gates in series to form the write path. As mentioned earlier, using only transmission gates in the write path is not possible because of the problems when V_{DD} drops to zero volts. However, the results of this cell when a one is written to the storage node are quite promising. This is the first cell that can correctly read out a one in all cases. The main difference between the transmission gate cell and the NMOS write cell is the removal of the threshold drop between the bitline and the storage node. These results confirm the theory developed earlier that writing a high value to the storage node is more important for reading the value out than the leakage off of the node. This is an interesting result and



(a) Read one, temperature variation

(b) Read one, process variation


Figure 3.19: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two series write transistors, length of 180 nm



(a) Read zero, temperature variation

(b) Read zero, process variation


Figure 3.20: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two series write transistors, length of 180 nm



(a) Read one, temperature variation

(b) Read one, process variation

Figure 3.21: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two series write transistors, length of 270 nm




(a) Read zero, temperature variation

(b) Read zero, process variation

Figure 3.22: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two series write transistors, length of 270 nm

is somewhat counterintuitive. The only failures of this memory cell occur for high supply voltages when a one is written to the storage node.



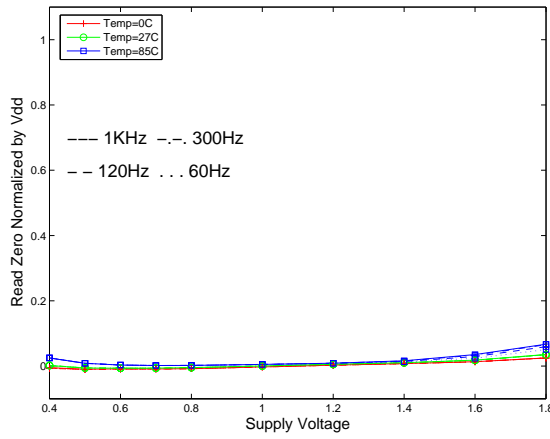
(a) Read one, temperature variation

(b) Read one, process variation

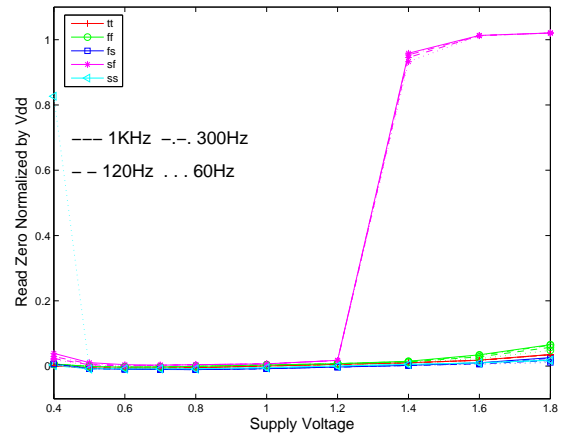
Figure 3.23: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two transmission gates in series

Figures 3.25 and 3.26 show a hybrid memory cell where a transmission gate is in series with an NMOS transistor. The transmission gate connects to the storage node and the transistor connects to the bitline. The read out values show that this cell has fewer failures than the all transmission gate cell when trying to read out a zero, but far more failures when reading out a one. This cell's performance when reading a one is similar to the NMOS only write cells.

A second hybrid cell was also tested. In this cell the transmission gate connects to the bit line and the NMOS transistor connects to the storage node. Figures 3.27 and 3.28 show the simulation results for this cell. It shows nearly the same performance as the other NMOS-transmission gate cell when reading out a zero. However, the ability to read out a one is improved. A possible cause for this is that connecting the NMOS transistor's source to the storage node allows the source-bulk voltage to be lower than when the source is connected through a transmission gate to the storage node. This lowers the threshold voltage of the NMOS transistor and allows it to write a higher one to the storage node. This trend is further evidence that the peak write voltage on the storage node is the key to




(a) Read zero, temperature variation

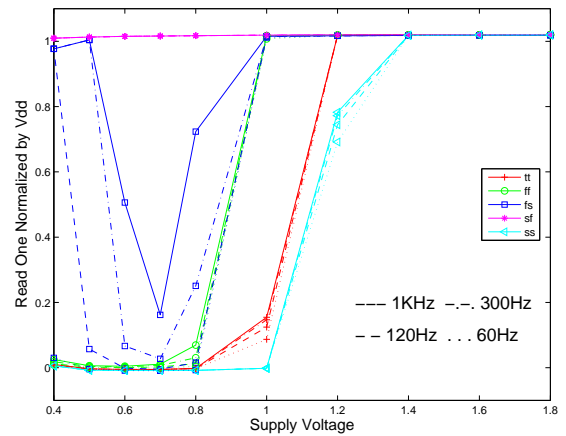


(b) Read zero, process variation

Figure 3.24: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , two transmission gates in series



(a) Read one, temperature variation



(b) Read one, process variation

Figure 3.25: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , transmission gate tied to store, NMOS tied to bitline

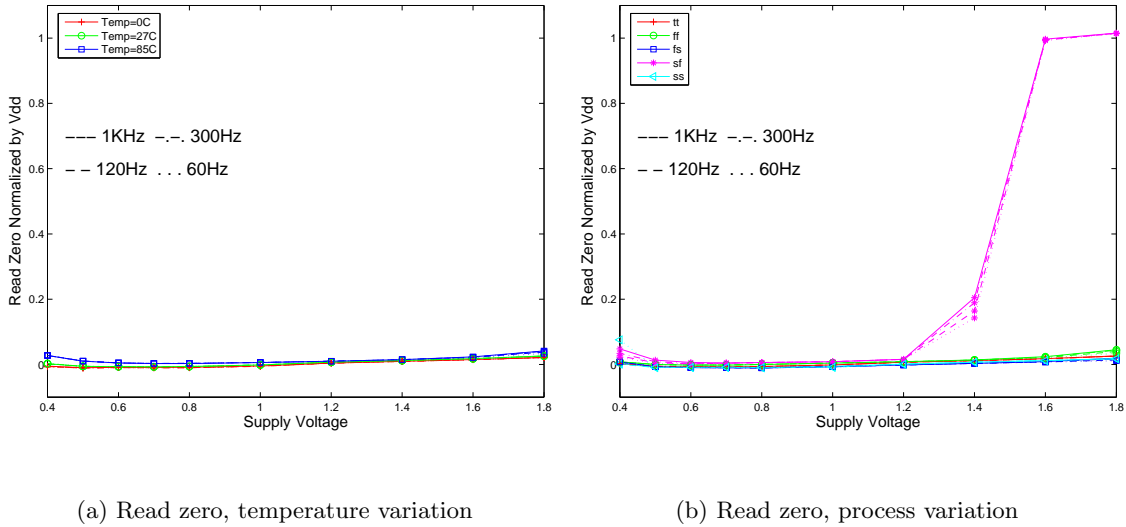



Figure 3.26: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , transmission gate tied to store, NMOS tied to bitline

the memory working for low power supply voltages.

The last type of memory cells examined use alternative devices or biasing to increase the threshold voltage of the write path. The first cell tested uses a high threshold transistor in the write path. This transistor is designed for a power supply of 3.3 V and has a minimum channel length restriction of 4λ . Both will severely reduce any charge leakage off of the storage node, but will also make it difficult to write a high voltage to the node. Our expectation is that this inability to write a high voltage to the node will cause this cell to fail when reading a one for much lower supply voltages than previously seen. Figure 3.29 shows the read out values when the cell should read as a one. These plots do in fact show that the cell fails for much higher supply voltages than previously seen. Figure 3.30 shows the read out values when the cell has been written to zero. For the slow-NMOS, fast-PMOS case this cell is never able to read a zero. These results indicate that this cell is definitely not a good choice for our system.


The final cell structure tested uses two NMOS transistors in series as the write path with their body voltages tied to a separate bias voltage. Simulations were run with body voltages of -100 mV , -200 mV , and -300 mV . The negative body voltage was used to increase the V_{SB} of the write transistors, increasing the threshold voltage and decreasing



(a) Read one, temperature variation

(b) Read one, process variation


Figure 3.27: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , NMOS tied to store, transmission gate tied to bitline



(a) Read zero, temperature variation

(b) Read zero, process variation


Figure 3.28: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , NMOS tied to store, transmission gate tied to bitline



(a) Read one, temperature variation

(b) Read one, process variation

Figure 3.29: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , high threshold write transistor



(a) Read zero, temperature variation

(b) Read zero, process variation

Figure 3.30: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , high threshold write transistor

the leakage current. Again, the expectation is that this will hinder the cell's ability to read out a one correctly. Figures 3.31 and 3.32 show the read outputs for the cell with a -100 mV body bias and Figures 3.33 and 3.34 show them when the bias is decreased to -300 mV . As expected, the more negative body bias, which should write lower voltages to the storage node, fails in more cases than the cell with a lower source to bulk voltage. These results indicate that using a lowered body voltage is not a good choice for our system. However, it is possible that this may still have a place in future system designs. The body bias could be determined adaptively to create a low write transistor threshold when the cell is being written and a high write transistor threshold when the cell is holding its value. This will be considered further in the future work chapter.

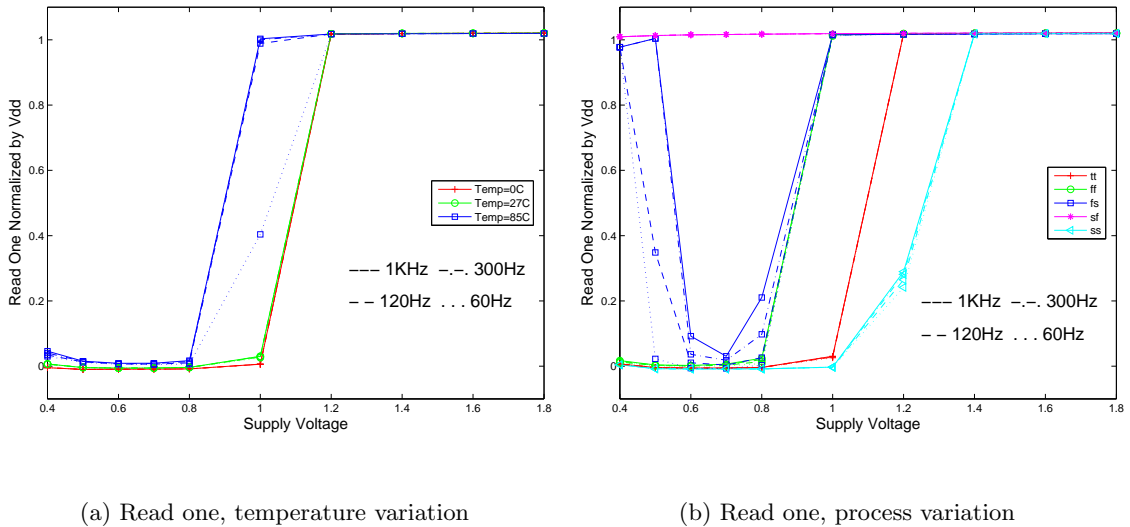




Figure 3.31: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , write transistor body bias of -100 mV

Based on the simulation results, a three transistor memory cell using an NMOS write path with minimum length was chosen to build the DRAM LUT. This memory cell was built using the Virtuoso layout software in the Cadence Design System. The resulting layout is shown in Figure A.1 in Appendix A.

Two concerns arise from these simulation results. The first regards the non-monotonicity of some of the voltages read from the output as V_{DD} is decreased. This occurs only with the fast-NMOS, slow-PMOS process corner around the supply voltages




(a) Read zero, temperature variation




(b) Read zero, process variation

Figure 3.32: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , write transistor body bias of -100 mV



(a) Read one, temperature variation



(b) Read one, process variation

Figure 3.33: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , write transistor body bias of -300 mV

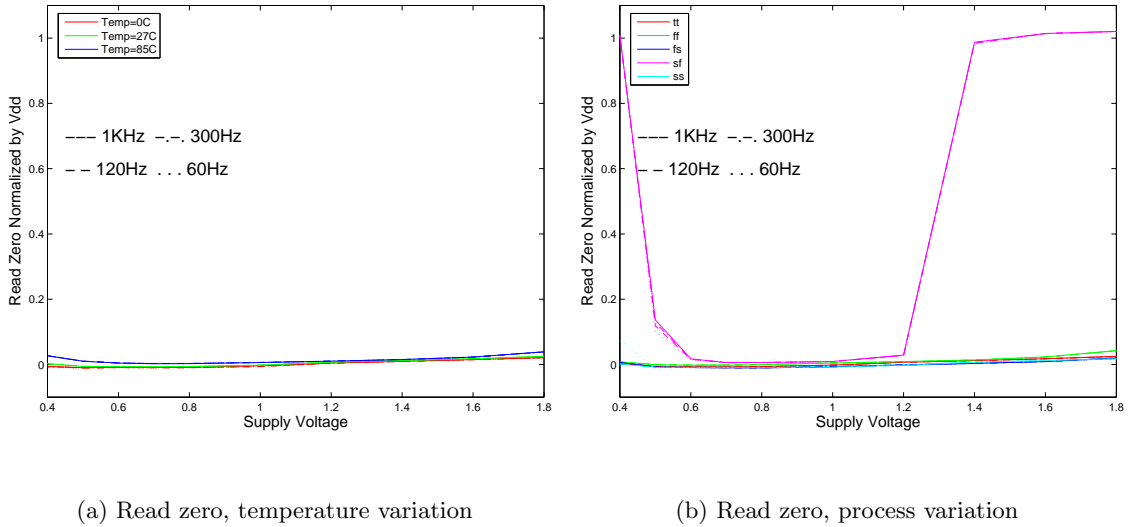


Figure 3.34: Voltage on rdOut node of DRAM sense circuit normalized by V_{DD} , write transistor body bias of -300 mV

where the storage node is being charged to a voltage near the NMOS threshold voltage. Though the reasons for this behavior are not yet completely understood, a possible explanation is that the NMOS transistor is rising out of subthreshold faster than the PMOS transistor as the supply voltage is increased from 400 mV to approximately 700 mV . This behavior may also be somewhat rooted in the inability of the ring oscillator providing the timing reference to track the performance of the DRAM cells. Currently the ring oscillator is designed to track the performance of an adder over process variation and temperature. However, the relationship between the DRAM and the timing are not fully understood. A possible improvement for future systems would be to include a path to track the memory performance in the self-generated clock. The second concern arising from the simulation results is the inability of the memory cells to function correctly for supply voltages below approximately 1 V . Three possible solutions have been proposed for this problem. The first and simplest is to increase the clock period for the read operation so the bitline has a longer time to clear. The second solution is to fold a replica of the memory cell into the power-on-reset circuit responsible for resetting the circuit after each power supply cycle [11]. Both of these solutions are currently being implemented and tested to be included on the test chip. In the case that these options do not provide an adequate solution, a third

option is to add a bootstrapping circuit to the write lines to pull them above the supply voltage and force the storage node to charge to a higher voltage. Writing the storage node to a higher voltage would both leave a larger margin of error for charge to leak and increase the V_{SB} of the write transistor, further reducing the leakage.

3.2.2 Register Cells

The base of the register memory cell is a static D flip-flop shown in Figure 3.35. A transmission gate was added to the output of the flip-flop to prevent all of the outputs on a single bitline from driving the line at once. This transmission gate is controlled by the read wordline.

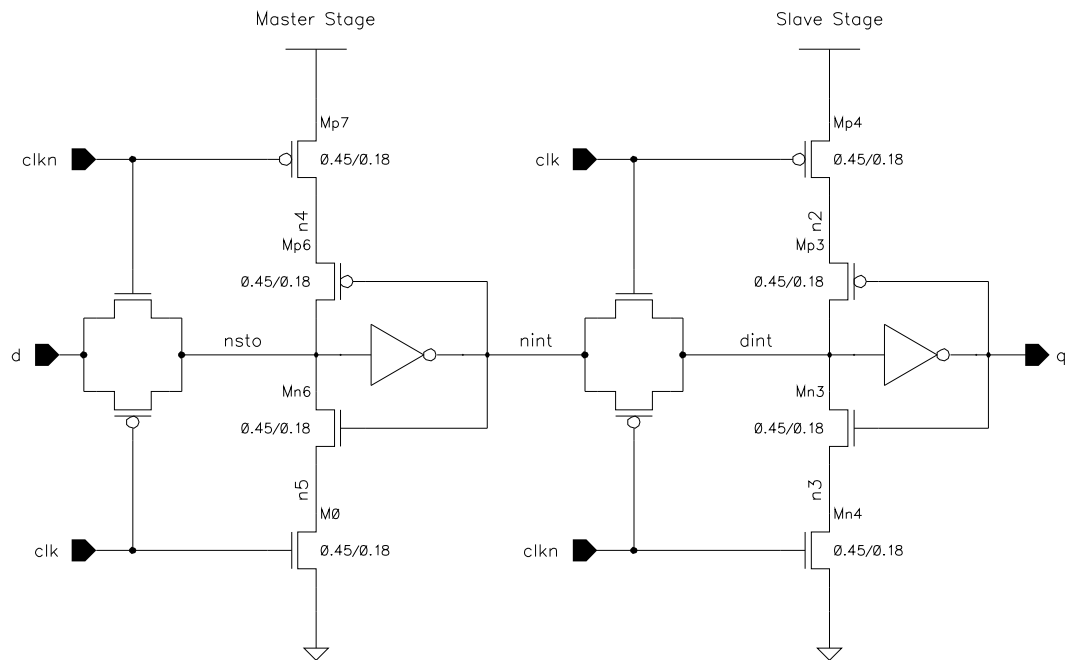


Figure 3.35: Static D flip-flop used in register LUT

3.2.3 Individual Cell and Row Decoders

The DRAM cell is written and read through a cell sense circuit shown in Figure 3.36. To write the memory cell, data is driven onto the write input line through the pin labeled W . The write enable signal, labeled we , is then asserted, turning on the transmission

gate connecting the write input to the bitline. At this time the row decoder which will be described next has already selected a single word line and the value on the bitline is written to the storage node. When the write enable signal is deasserted, the bitline is connected to the drain of the precharge PMOS transistor and the input of the tri-state buffer. To read the memory cell the precharge signal must first be asserted to precharge the bitline. This signal is active low in the sense circuit and active high in the LUT controller so an inversion stage is needed but not shown here. Once the bitline is precharged, the read enable signal is asserted. This activates the tri-state buffer and provides a path for the selected memory cell to clear the bitline depending on the charge on the storage node. The bitline will be cleared if the storage node is charged and should remain charged if the storage node voltage is around 0 V . The output of the tri-state buffer is the inversion of the bitline. This is the read output of the memory cell.

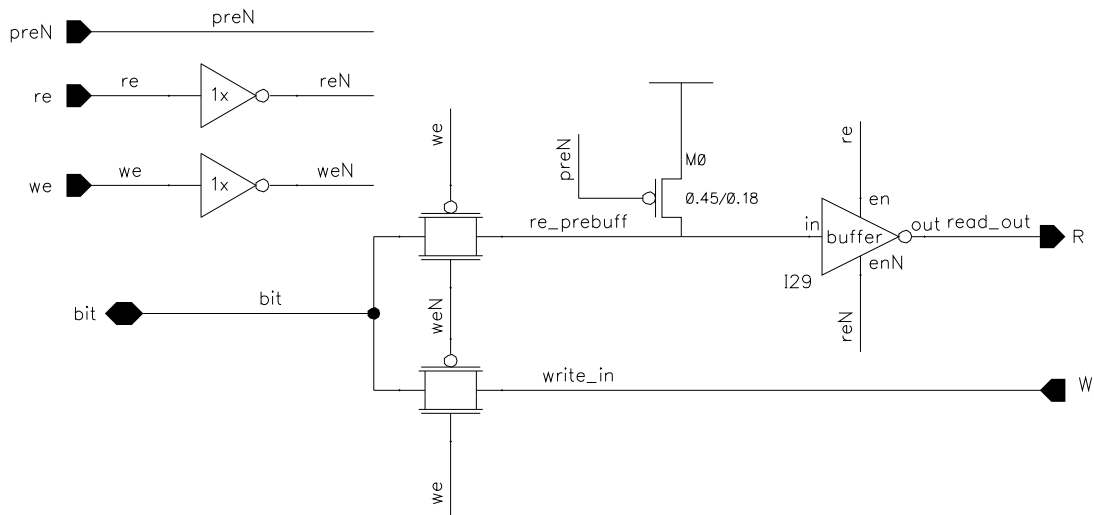


Figure 3.36: Read and write sense circuit for DRAM memory cell

The address decoder and the read/write line decoders are responsible for selecting a single word line to activate and sending the read and write control signals to that word line. Both of these circuits are used by the DRAM LUT as well as the register LUT. The address decoder takes a four bit binary number as its input and converts it into a sixteen bit output with only one bit high at a time. There is also an option to force all of the output lines to ground. This is used when one LUT is being written to prevent word lines on the

other LUTs from having floating voltages. A schematic of the address decoder is shown in Figure 3.37.

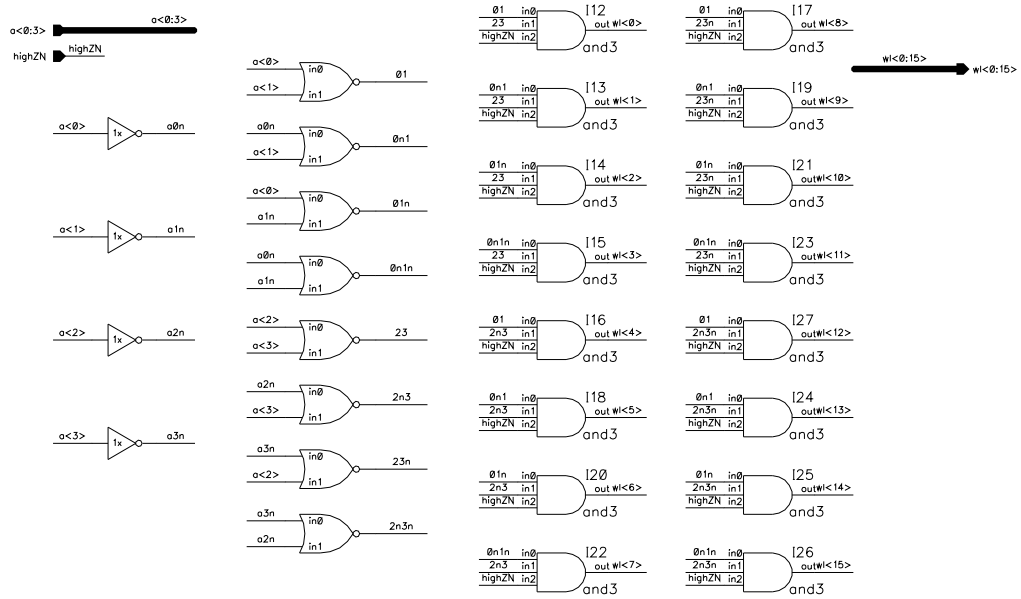


Figure 3.37: LUT 4-bit binary to 16-bit one-hot address decoder with high impedance switch

Each memory cell has a separate read and write control signal. These signals are generated by the AND operation between the word line selection for each row and the read enable and write enable signals. To add extra drive strength for the long wordline wires, the AND gate was split into a minimum size NAND gate followed by a $4\times$ minimum size inverter. This schematic is shown in Figure 3.38.

3.2.4 Circuit Implementation of LUT Controller

As described above, the LUT controller was built as three FSMs and two counters. Figure 3.39 shows the top level schematic of the controller. There are three additions to the block diagram for the controller that were not described previously. The first is the RS-latch connected into the initialize pin of the Read FSM. This latch is set when an initialize pulse is received from an off-chip source and is reset when the update end signal indicates that all cells have been written. The latch is used to allow the off chip initialize signal to be pulsed once to indicate the initialize sequence as opposed to having a handshaking or

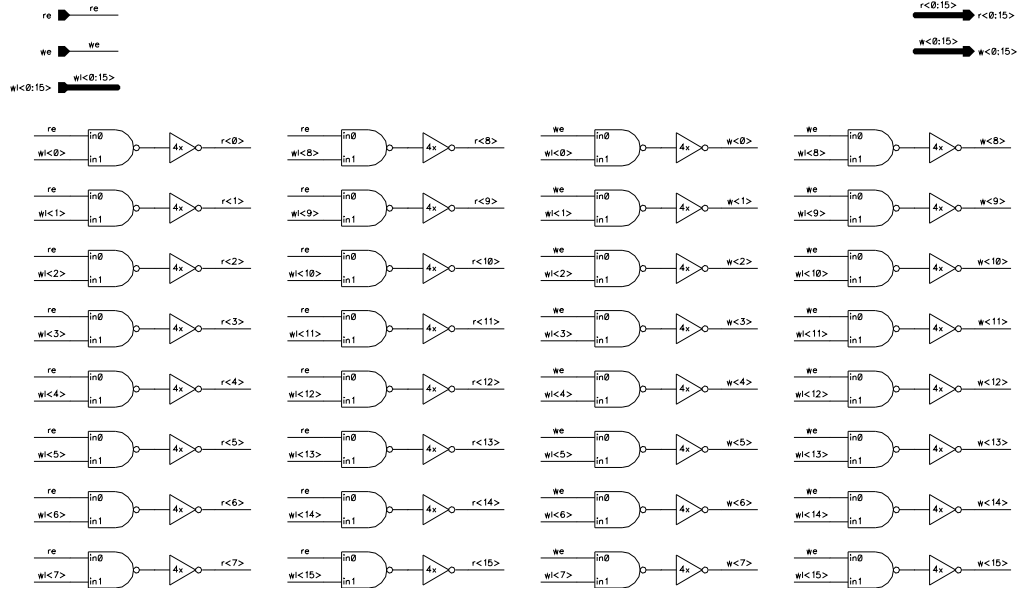


Figure 3.38: Read and write line decoders for both DRAM and register LUT

other indicator circuit to tell the off chip source to release the line when the cells have been written. Second, an OR gate is added to make the write exterior signal the OR function of initialize and the LUT select MSB. Cell write values should be read from the external memory every time the chip is initialized, as well as when the register LUTs are being written. The OR gate implements this function. Finally, a block of sixteen static latches is shown in this schematic. These latches hold the data being read in by the Read FSM until it is read out to the LUT by the Write FSM. These latches are only transparent when read enable is asserted.

The circuit schematic for the Master FSM is shown in Figure 3.40. Again, an RS-latch is used to hold a pulsed signal labeled clear constant throughout the memory refresh. The clear signal is used to hold the clocks to the data shift memory throughout the LUT write sequence and thus must be asserted until the update has ended.

Figure 3.41 shows the schematic for the Read FSM. We expect to have a delay between the clock edges and the assertion and deassertion of the precharge signal. This is because of the combinational logic between the flip-flop outputs and the precharge generation. However this delay should not cause problems with the circuit operation because

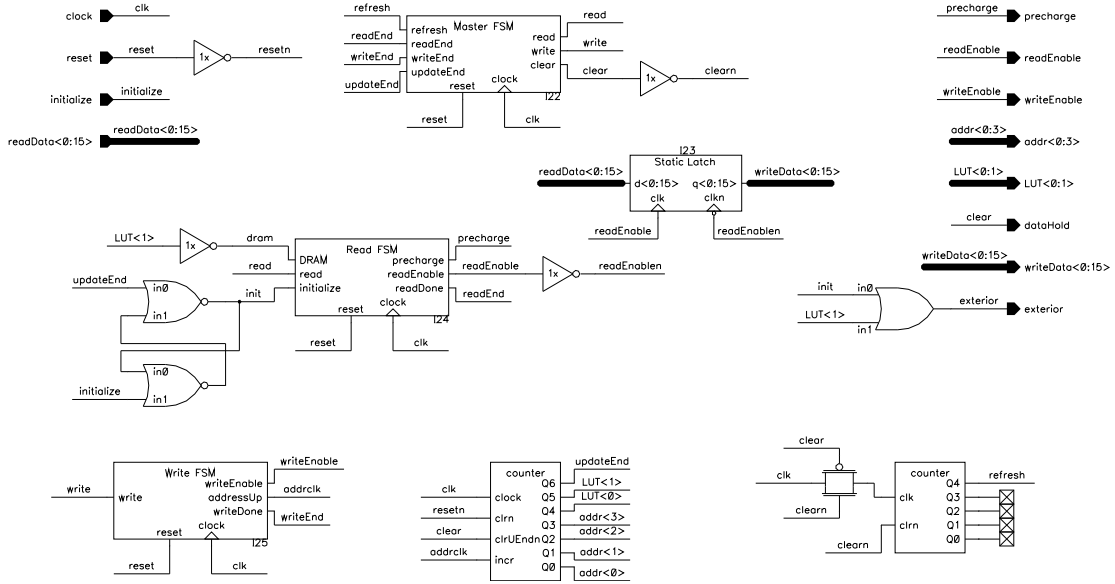


Figure 3.39: LUT controller top level schematic

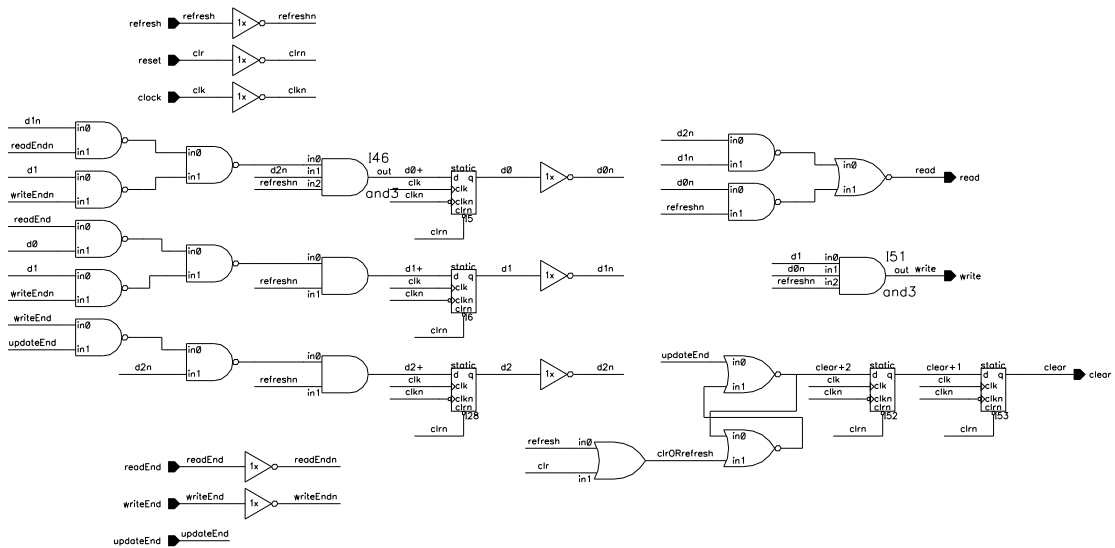


Figure 3.40: Master FSM circuit schematic

neither the precharge nor the read enable signals can cause a glitch in operation.

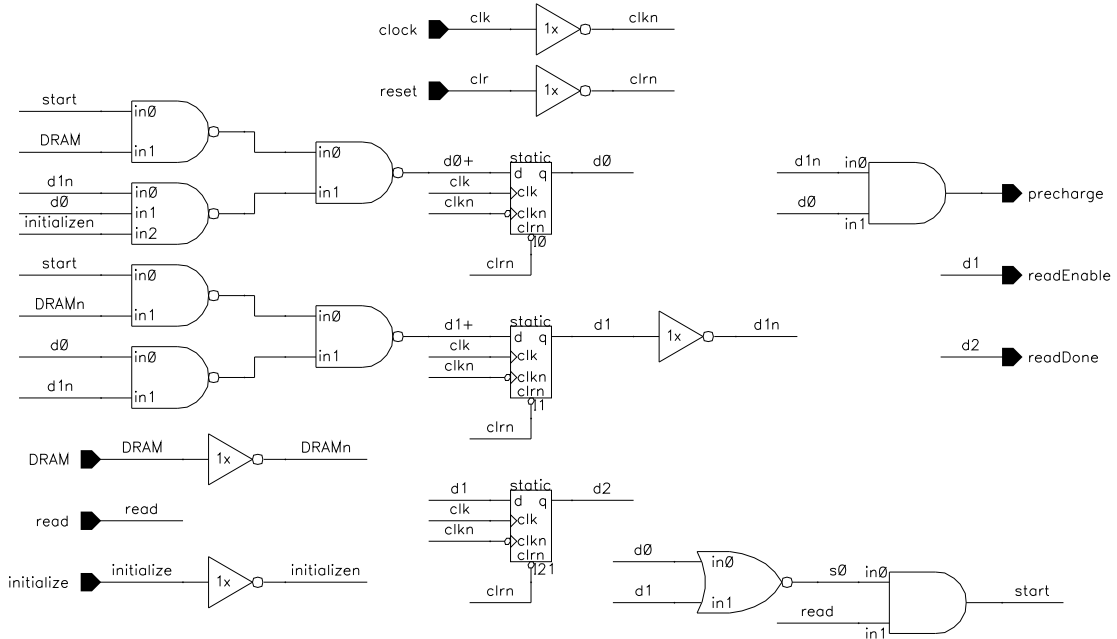


Figure 3.41: Read FSM circuit schematic

The Write FSM schematic is shown in Figure 3.42. Two extra inverters have been added to delay the $d1n$ signal that is input to the AND gate generating write enable. This was done to ensure that the write enable signal did not overlap with the address increment signal, possibly causing data to be written to two cells while the address was changing.

The counter used to provide the address and LUT select bits is shown in Figure 3.43. To prevent glitching in the counter, two-input MUXes were used to gate the flip-flop inputs to either the next count value or the current count. Selection of the flip-flop input was made by the count increment input, driven by the Write FSM. Bits zero through three are used to select the address to be written, bits four and five select the LUT being written, and bit six indicates that all LUTs have been written. Once the Master FSM has received the signal that the write sequence is complete, bit six is reset. All bits will be reset by a system reset. A replica of this counter truncated to five bits is used to initiate the refresh sequence.

Figure 3.44 shows a sample set of outputs from the LUT controller. The address

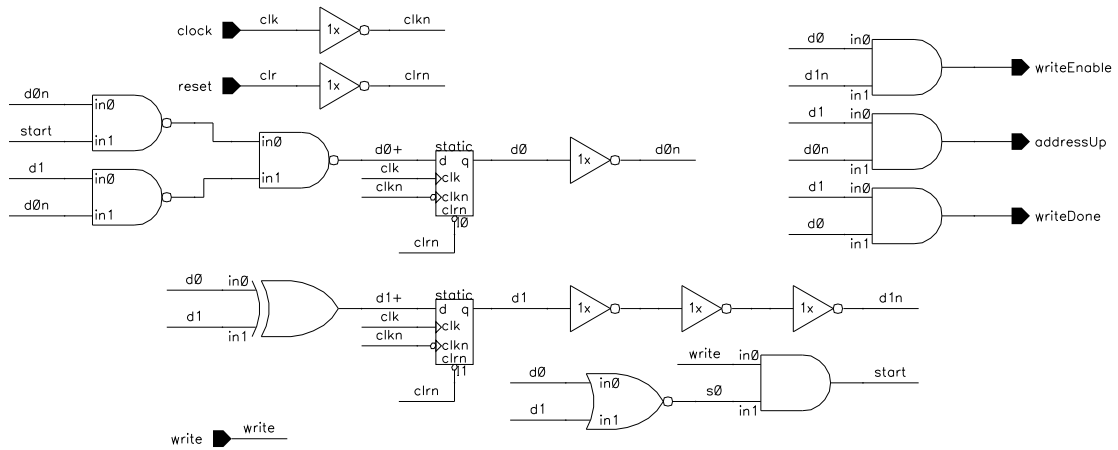


Figure 3.42: Write FSM circuit schematic

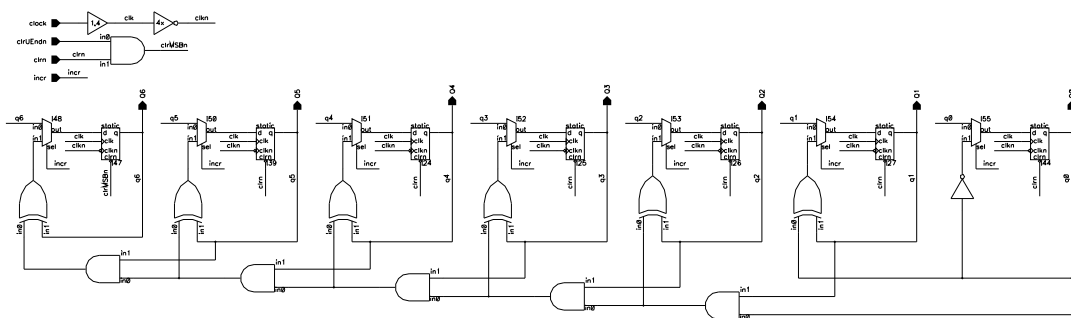


Figure 3.43: 7-bit counter used for generating the LUT addresses and select bits, as well as the update end bit

changes on the rising edge of the clock and two clock cycles later the precharge is asserted, followed by the read enable signal. After two more clock cycles the write enable is asserted. There is a clock cycle of timing margin between the falling edge of the write enable and the address change. The clock cycles where no signals are asserted are times where the finite state machines are transitioning between states and machines. These provide some timing margin to ensure that the address and thus the wordline selection is steady when the read and write events occur.




Figure 3.44: Example output of the LUT controller, shows signals necessary to write one row of memory cells

3.3 Full Memory System Implementation

The full LUT memory system for the 16–tap DA filter being built on the test chip requires four LUTs, the LUT controller, and several pieces of support circuitry to allow the single controller to write all four LUTs. The final design for the DRAM LUT is shown in Figure 3.45. The corresponding design for the register LUT is shown in Figure 3.46.

Waveforms showing write and read sequences to the DRAM LUT are shown in Figure 3.47. These were simulated at V_{DD} of 1.8 V for typical-typical transistor models.

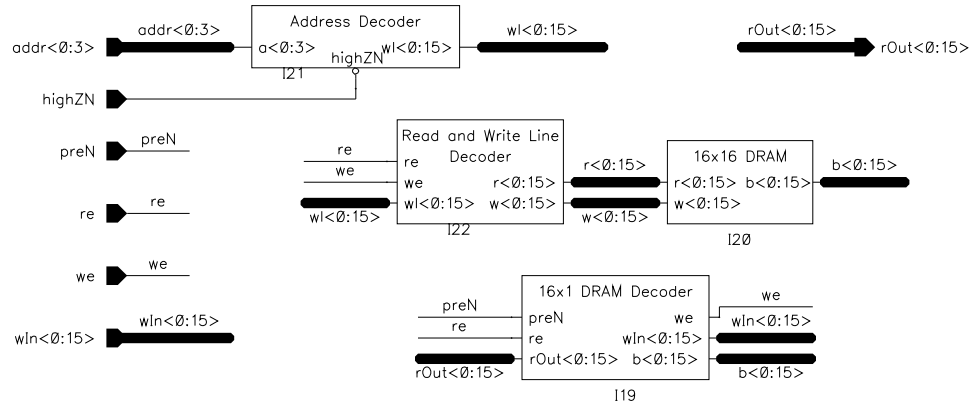


Figure 3.45: Full schematic of DRAM LUT

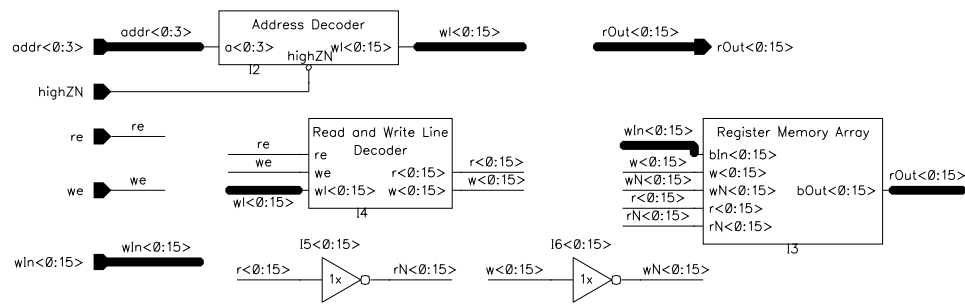



Figure 3.46: Full schematic of register LUT




(a) Read and write one to DRAM LUT

(b) Read and write zero to DRAM LUT

Figure 3.47: Sample waveforms showing voltage sequence necessary to write and read from DRAM LUT

The same sample waveforms for the register LUT are shown in Figure 3.48. Again these were simulated at V_{DD} of 1.8 V for typical-typical transistor models. The delay between a value being present on the input to it being read out is not caused by delay in the LUTs. This is a product of first having to write the value into the LUT before it can be read back out.



(a) Read and write one to register LUT

(b) Read and write zero to register LUT

Figure 3.48: Sample waveforms showing voltage sequence necessary to write to and read from register LUT

Chapter 4

Low-Swing Interconnect and Modulation

On-chip signaling, particularly over relatively long wires and at high speeds, is taking an ever increasing percentage of a system's power budget. Long wires present an increasing problem as process sizes scale down. Wire widths scale faster than lengths, causing the wire resistance to increase. The capacitance of the wire decreases, but not as quickly as the gate and parasitic capacitance. Thus, greater percentages of a system's power are being used to drive long wires. In this chapter we describe some of the sources of power dissipation in interconnect as well as some possible methods of reducing that power. A few circuit implementations from the literature are considered along with some new ideas.

4.1 Interconnect Power

There are two main sources of interconnect power: dynamic switching and static power arising from several sources. Both of these will be described as well as some possibilities for reducing interconnect power.

4.1.1 Dynamic Switching Power

Dynamic power can be described by Equation 4.1, where α is the switching probability, C_L is the load capacitance that must be charged or discharged, V_{DD} is the power

supply and the voltage swing of the load capacitance, and f is the switching frequency. In signaling systems the switching probability will largely be determined by the data being transmitted as well as any bus encoding architecture implemented in the system. The load capacitance will depend on the physical dimensions and layout of the wires being driven as well as the internal parasitic capacitance in the driver and the capacitive load presented by the receiver. The maximum voltage swing of the signal will depend on the largest supply voltage of the system and the lowest voltage swing will depend on the sensitivity of the receiver as well as the availability of a lower voltage supply. Finally, the signaling frequency is largely determined by the speed of the surrounding circuits that are using the data being transmitted.

$$P_{dynamic} = \alpha C_L V_{DD}^2 f \quad (4.1)$$

4.1.2 Static Power of Peripheral Circuitry

Static power is another source of power dissipation in interconnect circuits. This power comes from any DC current flowing from the power supply to ground and can be described by Equation 4.2. Traditionally analog bias circuits in the driver or receiver have been the primary source of any static power dissipation. However, many drivers and receivers can be built using static CMOS, so this source of DC current has become much less significant. Presently, leakage current accounts for much of the static power consumed. This leakage current term accounts for reverse-biased diode leakage, subthreshold leakage, and tunneling through the gate oxide. As threshold voltages have decreased with each process generation, subthreshold leakage has taken an increasingly large role in the total leakage current.

$$P_{static} = V_{DD}(I_{bias} + I_{leakage}) \quad (4.2)$$

4.1.3 Possible Power Reduction

Each of the power equations offers some opportunity for power reduction. Common techniques for reducing dynamic power include reducing load capacitance, switching activ-

ity, clock frequency, and power supply voltage. Static power can be reduced by decreasing the power supply voltage or the static current. We chose to integrate several different techniques into our design to target these various opportunities to reduce interconnect power. The drivers and receivers implemented do not require any analog biasing, which eliminates all but the leakage component of static current. Minimum sized transistors are used when possible to reduce load capacitance. Modulation techniques were investigated to reduce clock frequency and switching activity. Finally, reduced voltage swings on the greatest capacitive load, the wire itself, were examined to reduce both dynamic and potentially static power.

4.2 Low-Swing Interconnect

The basic principle of low-swing interconnect is that if a signal line's swing can be reduced, it will be easier to switch. Initially this was used to make faster interconnections. However, more recently this has been used to decrease switching power. The tradeoff in power reduction is that more complicated drivers and receivers and possibly additional power supplies are needed. These extra circuits and supplies will require additional power, making low-swing interconnect inefficient for some situations. Several driver and receiver circuits are presented here and the signal pairings are compared.

4.2.1 Driver Circuits

Two low-swing driver circuits were designed and simulated. The first low-swing driver is a static CMOS circuit, which is conventionally used in low-swing implementations[16]. In the conventional low-swing driver, input data is first inverted using a full-swing inverter and then it is inverted again through a second CMOS inverter with a reduced power supply voltage. The resulting output data has the same logical value as the input value but the high voltage is limited to the value of the reduced power supply of the second inverter. A schematic of this driver is shown in Figure 4.1(a).

An all NMOS low-swing driver described in [16] was also considered. Again the input data is inverted by a full-swing CMOS driver but then $input$ and \overline{input} are used to

drive an NMOS inverter with a potentially reduced supply voltage. Because of the NMOS pull-up transistor, even if the full supply voltage is used for the NMOS inverter, the output will not be able to rise above $V_{DD} - V_{tn}$. Note that V_{tn} will not be constant and may be large because of the body effect. As with the conventional driver, the NMOS driver output data has reduced high voltage swings. This circuit is shown in Figure 4.1(b).

Each of the drivers is sized to have the smallest area possible to still achieve reasonable drive strength. In the conventional driver, the pull-up PMOS transistor was sized to be three times the minimum width to give the driver approximately symmetrical pull-up and pull-down strength. Initially both the PMOS and NMOS had minimum size but this caused the driver to perform poorly with data rates around 100 MHz. The NMOS driver does not require a larger pull-up transistor. However, it does need to use an additional inverter that is not required in the CMOS driver. Another benefit of the NMOS driver at the architectural level is that it can achieve reduced voltage swings on the output without requiring an additional power supply. The tradeoff is that this threshold voltage drop will vary due to the body effect, making the signal swing less predictable. Some of the power costs or savings from these small differences in the drivers will depend on the resources available in the large scale system, such as additional power supplies or logic that provides the input and its complement. The testing method for these drivers accounts for the power drawn from the full supply as well as the reduced supply and any inversion is included in the power usage. However, power estimates for using a DC to DC converter to implement the reduced supply are not reflected in these simulations.

4.2.2 Receiver Circuits

Three low-swing receiver circuits were investigated. A conventional low-swing receiver was the first examined [17]. The circuit schematic is shown in Figure 4.2. Low-swing input data is used to drive a pair of NMOS transistors that are connected to two cross-coupled pull-up PMOS transistors.

The second receiver considered is a pseudodifferential implementation that was described in [16] and is shown in Figure 4.3. In this receiver the input data comes into a parallel pair of PMOS transistors. The sources of these transistors are connected through

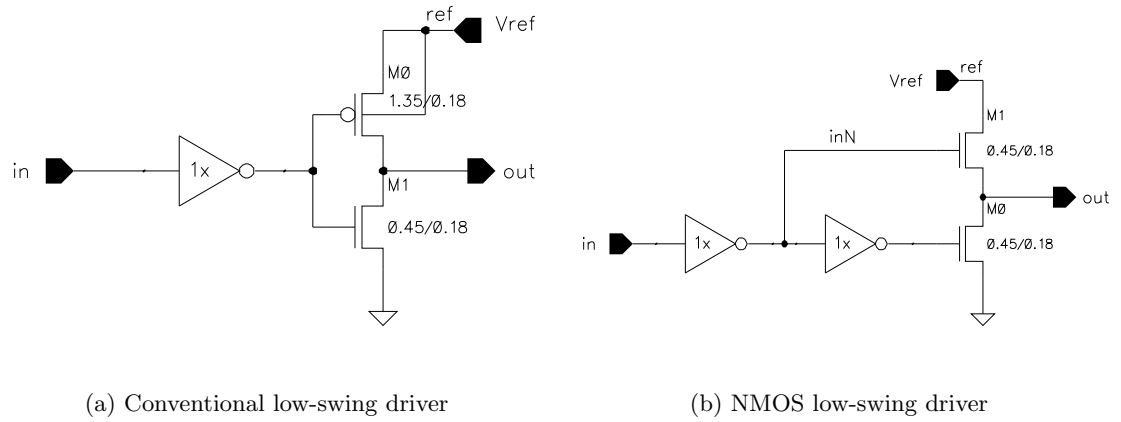


Figure 4.1: Low-swing driver schematics

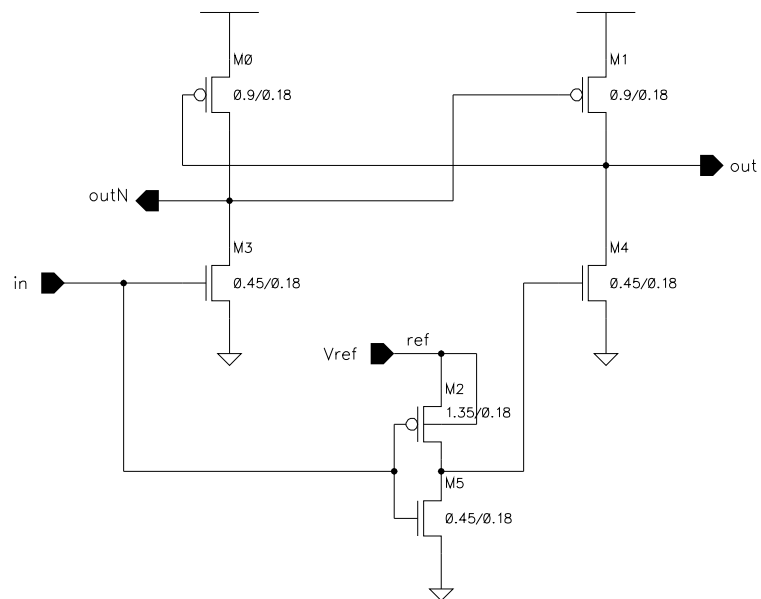


Figure 4.2: Conventional low-swing receiver schematic

a clocked PMOS transistor to the full supply, as well as to another pair of parallel PMOS transistors with one gate grounded and the other connected to the low-swing supply. Both pairs of PMOS transistors have their drains connected to a pair of cross-coupled NMOS transistors, which are connected to ground through a second cross-coupled NMOS pair. The drains of the lowest pair of NMOS transistors in the stack are connected as inputs to an RS-latch. The clock signal is used to turn on the topmost PMOS transistor for half of the new data period, allowing the incoming data to be evaluated. Clock then turns the NMOS transistors at the bottom of the stack on for the other half of the data period, driving the RS-latch inputs to ground and holding the evaluated data values at the output. In simulation, the clock was timed such that the incoming signal was sampled during the second half of the bit period. This caused the edges of the output data to be skewed by half a clock cycle from the edges of the input data. However, sampling the input wave later in the bit period improves the accuracy of the receiver.

A key aspect of this design is the potentially reduced lower functional limit of the low swing supply as compared to the conventional receiver. The conventional receiver (Figure 4.2) relies on low-swing inversion of the signal to drive the mirrored pull-down path. When the low-swing voltage is below an NMOS threshold voltage, the inverter will not be turned on strongly and we would expect neither pull-down path to be strongly off or on. This would slow down the receiver and potentially limit the minimum low-swing voltage to no lower than an NMOS threshold voltage. However, the pseudodifferential receiver connects the low-swing input to the gates of two pull-up PMOS transistors and relies on current ratios rather than inversion to flip the comparator. This should allow the pseudodifferential receiver to operate at voltages much below an NMOS threshold voltage. Operating at lower voltages would further reduce the dynamic power of driving the long wire and further offset the added power of this more complicated receiver. In addition, this receiver uses two positive feedback (PFB) loops to evaluate the input as opposed to the single PFB loops used in the other two receivers. The second PFB loop provides additional gain, allowing the receiver to switch faster.

The final receiver considered is an adaptation of the dynamic comparator from the DC/DC converter described in [18], which we will call the dynamic ratioed receiver. The

pseudodifferential receiver, we would expect to be able to operate reliably at much lower voltages than the conventional receiver. Second, this receiver builds the current ratioing circuits out of NMOS transistors as opposed to PMOS transistors. Because the input signal is compressed toward ground, the NMOS transistors will have reduced overdrive voltages and will thus be conducting less current. However when the input voltage enters into the gates of PMOS transistors those transistors should have larger average current because the conduction path will never be fully cut off. This current will only increase as the reduced supply decreases relative to the full supply. In addition to power savings because of the reduced ratioing current, this receiver also requires many fewer transistors than the pseudodifferential receiver, 12 vs. 20.

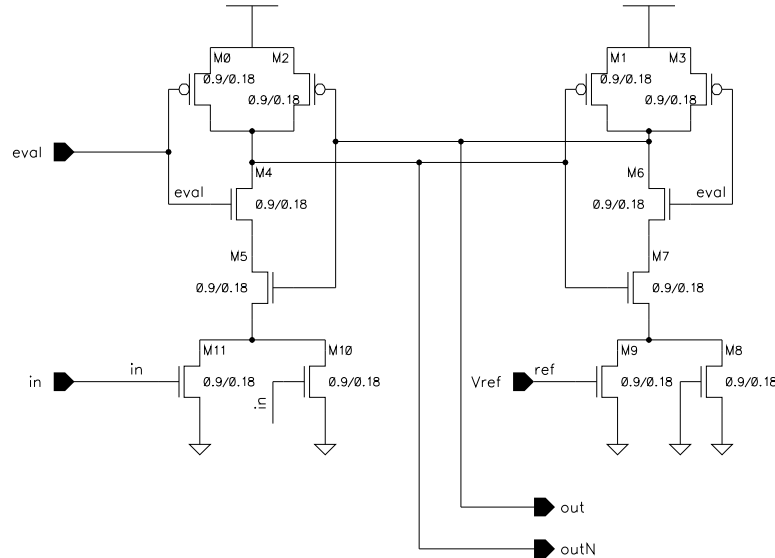


Figure 4.4: Dynamic ratioed low-swing receiver schematic

The true test of each of these circuits is in their performance when driving and receiving signals on a long wire. Each of the possible transmit and receive pairs were simulated using Spectre in the Cadence Design System. The details and results of these simulations are presented in the next section.

4.2.3 Implementation and Comparison

Each of the low-swing driver and receiver pairings were simulated using the same test bench. The low-swing voltage was swept from 1.8 V to 100 mV in 100 mV increments. Temperature was swept over 0°C , 27°C and 85°C using the typical-typical transistor models. The process corners were swept at a temperature of 27°C as well. One of the test benches used for these simulations is shown in Figure 4.5. The conventional driver and pseudodifferential receiver are shown. However, all of the pairings were simulated in the same configuration excluding the clock or evaluate signals where necessary for specific receivers.

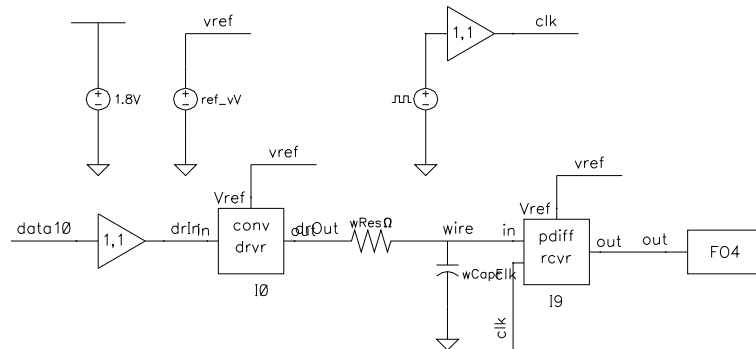


Figure 4.5: Simulation testbench example for conventional driver, pseudodifferential receiver

Wire length and thus wire resistance and capacitance have a large effect on the amount of dynamic power used by the circuit, so wire lengths of 1 mm , 5 mm and 1 cm were all tested. The long wires were assumed to be signal wires in the metal two or metal three layers. Based on our current chip layouts, there should be little or no signal wiring in the metal layers above metal three. Therefore the primary interconnect capacitance is between the signal lines in a given metal layer and from the signal lines to the more densely packed layers below, such as metal one, poly, and the substrate and active regions. The worst case capacitance between the possible signal lines and the layers below them is between metal one and metal two. This approximation is realistic although somewhat pessimistic because the power supply rails are routed in metal one. In this case the capacitance per unit length is $0.219\text{ fF}/\mu\text{m}$ for a metal width of $0.28\text{ }\mu\text{m}$ and the metal two sheet resistance is

$0.078 \Omega/sq$ [19]. The equivalent lumped capacitances and resistances for wires with lengths of 1 mm , 5 mm and 1 cm are shown in Table 4.1.

| Wire Length | Lumped Capacitance | Lumped Resistance |
|----------------|--------------------|-------------------|
| 1 mm | 219 fF | 279Ω |
| 5 mm | 1.095 pF | 1393Ω |
| 1 cm | 2.19 pF | 2786Ω |

Table 4.1: Wire parameters used in interconnect simulations

A single random bit stream was used as the test input to each of the transmission pairs. The values were generated using the random number generator provided in Excel and are listed here: $11B04820E5E89$. The data source was implemented as a piecewise linear voltage source built by the Bitgen program developed at North Carolina State University. Bit rates of 10 MHz were tested and the average total power as well as the average power dissipated in the wire were measured.

In addition, the input bitstream provided an opportunity to examine the response of the system to a single high bit surrounded by zeros as well as a single zero bit surrounded by ones. Figure 4.6 shows sample driver and receiver waveforms using the sample bit stream input for nominal process and temperature parameters. At time equal to $2.50 \mu s$, the input wave rises to a one for a single bit period and then falls back to zero. This one is surrounded on both sides by several consecutive zeros. Thus, the wire and any parasitic capacitance should be fully discharged when this one occurs and will have to charge up to the output high noise margin in a single bit period. By examining the receiver output near the end of the bit period containing the one we can see a good measure of the system's ability to drive the given wire length. The equivalent test can be done for driving the wire low for a single bit period with surrounding high bits at a time of $4.20 \mu s$. Both of these tests provide valuable information about the transmission circuits, which could not necessarily have been gained by only sending a clock signal through the channel. Measurements of the output voltages at these worst case points in the data stream are even more important than the total power necessary to send the data and will be considered first below.

The first signal pair considered used the conventional driver and the conventional

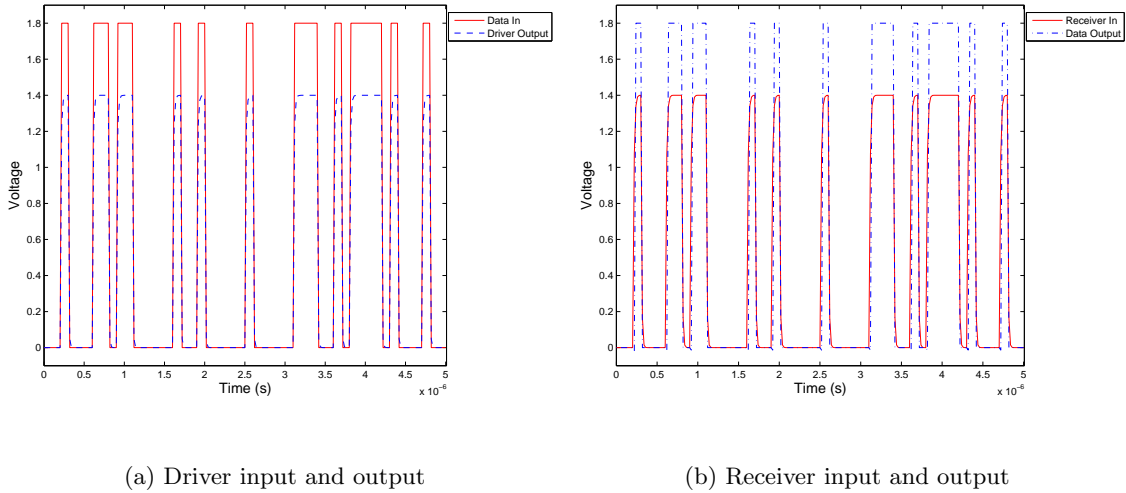


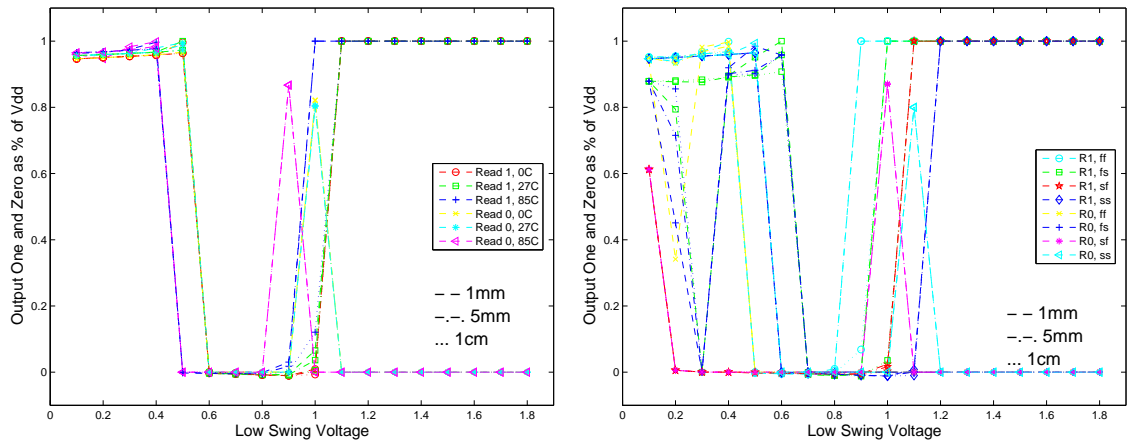
Figure 4.6: Sample input and output for conventional-conventional transmission pair, low-swing voltage of 1.4 V

receiver. Figure 4.7 shows the simulated output for the worst case one and zero as described above, normalized by V_{DD} . These results indicate that temperature limits the low-swing voltage to 1.1 V and some of the process corners fail to transmit correctly for low-swing voltages below 1.2 V. Thus the power measurements for this pair are only valid when the reduced supply voltage is ≥ 1.2 V for this process and these data transmission parameters.

The simulated worst case received one and zero values for the conventional driver and pseudodifferential receiver pair are shown in Figure 4.8. This pair receives all values correctly for reduced supply voltages of ≥ 0.9 V and is limited by the process corners rather than temperature. However, if the length of the wire being driven is limited to 5 mm the signal pair works down to 0.8 V as well, although the process corners are still the limiting factor.

Figure 4.9 shows the simulated output for the worst case one and zero transmissions with a conventional driver and the dynamic ratioed receiver. The output is correct for low-swing voltages ≥ 1.0 V and is limited by the slow NMOS, slow PMOS process corner. As with the conventional-pseudodifferential pair, if the wire length is limited to 5 mm the lower limit of the reduced supply voltage decreases by 100 mV.

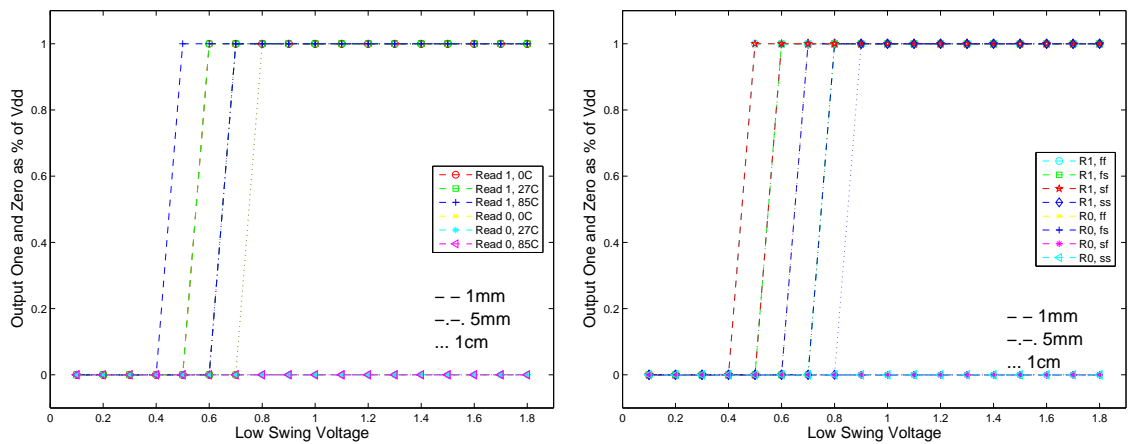
Next, each of the receivers was simulated with the NMOS driver, which is an



(a) Read one and zero over temperature, tt process

(b) Read one and zero over process, 27°C temperature

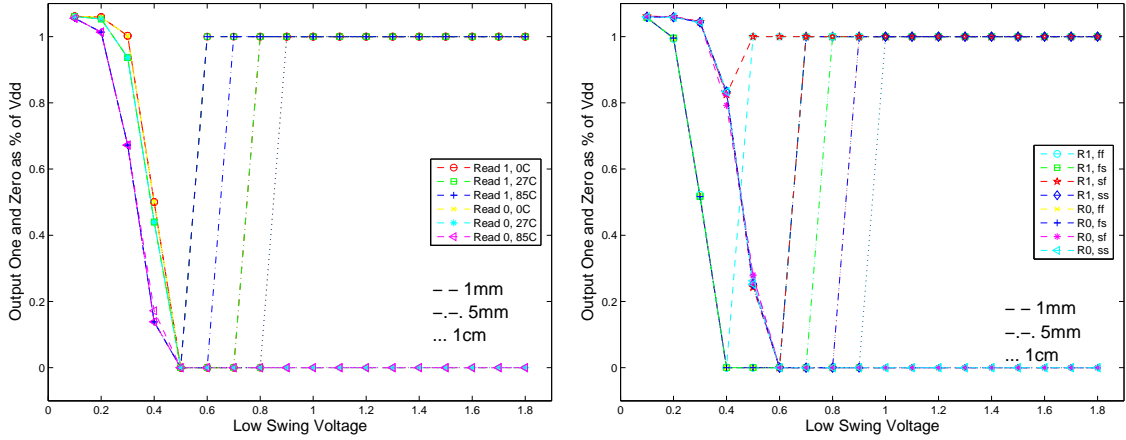
Figure 4.7: Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, conventional driver, conventional receiver



(a) Read one and zero over temperature, tt process

(b) Read one and zero over process, 27°C temperature

Figure 4.8: Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, conventional driver, pseudo-differential receiver



(a) Read one and zero over temperature, tt process

(b) Read one and zero over process, 27°C temperature

Figure 4.9: Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, conventional driver, dynamic ratioed receiver

inverting driver. The simulation testbench was modified to account for the threshold voltage drop across the NMOS pull-up transistor in the driver for some low-swing voltages. The modified testbench for the NMOS driver and conventional receiver is shown in Figure 4.10. The resulting outputs for the NMOS driver with the conventional receiver are shown in Figure 4.11. Again, these plots show the output of the receiver normalized by V_{DD} for the worst case one and zero transmission. This pair fails to correctly transmit a one over the 1 cm wire for a low-swing supply voltage of 1.7 V. If the wire length is restricted to ≤ 5 mm then this pair works correctly for low-swing supply voltages above 1.2 V.

The simulated worst case received one and zero values for the NMOS driver and pseudodifferential receiver pair are shown in Figure 4.12. These results show that the worst case one and zero are both received correctly and upconverted for low-swing voltages ≥ 300 mV.

Figure 4.13 shows the simulated output for the worst case one and zero transmissions with an NMOS driver and the dynamic ratioed receiver. In some cases, the receiver incorrectly outputs a zero rather than the expected one for low-swing voltages of 1.6 V to 1.8 V. This restricts both the maximum and minimum low-swing voltages for this driver,

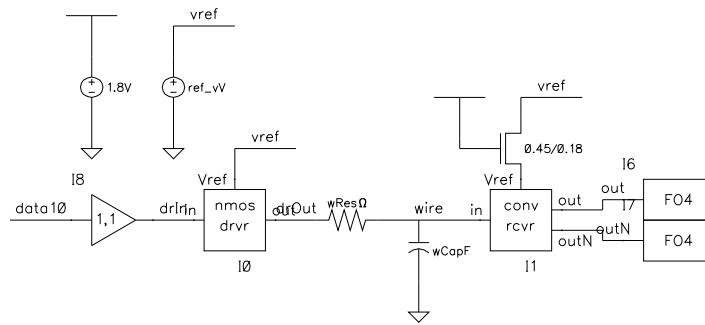
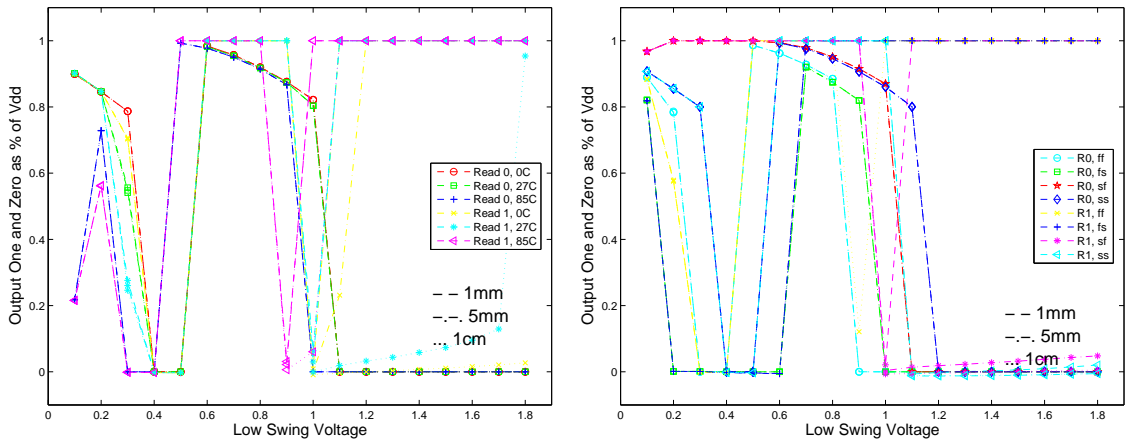


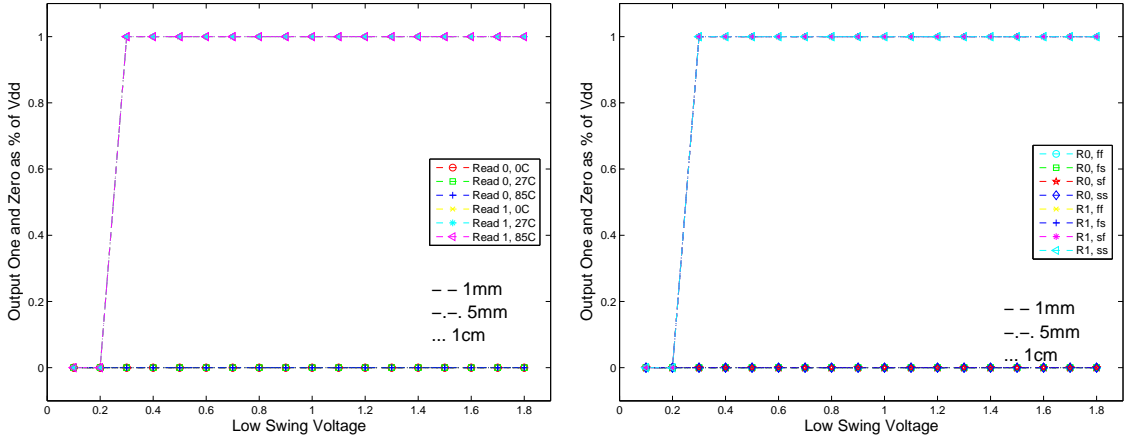
Figure 4.10: Modified low-swing test bench circuit to account for NMOS driver V_{tn} drop in pull-up path



(a) Read one and zero over temperature, tt process

(b) Read one and zero over process, 27°C temperature

Figure 4.11: Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, NMOS driver, conventional receiver



(a) Read one and zero over temperature, tt process

(b) Read one and zero over process, 27°C temperature

Figure 4.12: Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, NMOS driver, pseudodifferential receiver


receiver pair.

The simulation results described above indicate the range of low-swing voltages where each of the transmission pairs operate correctly. These results are summarized in Table 4.2.


| Driver and Receiver | Range of Low-Swing Supply Voltage | | |
|---------------------------------|-----------------------------------|-----------------|-----------------|
| | 1 mm | 5 mm | 1 cm |
| Conventional-Conventional | 1.2 V | 1.2 V | 1.2 V |
| Conventional-Pseudodifferential | 700 mV | 800 mV | 900 mV |
| Conventional-Dynamic ratioed | 700 mV | 900 mV | 1 V |
| NMOS-Conventional | 1.2 V | 1.2 V | 1.8 V |
| NMOS-Pseudodifferential | 300 mV | 300 mV | 300 mV |
| NMOS-Dynamic ratioed | 600 mV | 600 mV to 1.5 V | 600 mV to 1.3 V |

Table 4.2: Range of operational low-swing voltages for each driver and receiver pair

Based on these operating ranges, the average total power of each of the driver and receiver pairs was compared using each of the three wire lengths. Figure 4.14 shows the simulated power used to transmit data across 1 mm and 5 mm wires and Figure 4.15 shows the simulated power used to transmit across a 1 cm wire.




(a) Read one and zero over temperature, tt process




(b) Read one and zero over process, 27°C temperature

Figure 4.13: Voltage at receiver output normalized by V_{DD} for worst case one and zero transmission, NMOS driver, dynamic ratioed receiver



(a) Average total power for full transmission system with 1 mm wire



(b) Average total power for full transmission system with 5 mm wire

Figure 4.14: Average total power dissipated by driver, receiver and wire for 1 mm and 5 mm wires using typical-typical transistor models at 27°C




Figure 4.15: Average total power dissipated by driver, receiver and wire for 1 *cm* wire using typical-typical transistor models at $27^{\circ}C$

These simulation results indicate that the NMOS driver using the pseudodifferential receiver is lower power than the other combinations for nearly all cases. As the wire length increases, the dynamic ratioed receiver with either receiver approaches the lowest signaling power. However, this receiver has a large increase in power just before it reaches low-swing voltages too low to receive. This occurs because the transistors slow down due to reduced overdrive voltages and require longer to reach a digital steady state output. Because the circuit is in an intermediate state for a longer time, more current and thus more power is drawn from the supply. In addition, this suggests that the failure mechanism for the dynamic ratioed receiver is slow response time of the input path.

4.3 Data Compression through Pulse-Width Modulation

Pulse width modulation takes advantage of a system's ability to detect small delays or advances in the rising and falling edges of a pulse to encode more information in each transition of the bitline. In some cases, this can be used to increase data throughput [20]. However, we are more concerned with decreasing the transmission power by reducing the signaling frequency and potentially the number of transitions on the wire. The modulation scheme chosen uses three modulation bits encoded in a return-to-zero pulse and requires a forwarded reference clock. The bitrate can easily be divided by three because of the

way the pulses are generated and the inherent speed of the process being used, as well as the relatively low data rates being considered. Thus, adding the modulation allows us to divide the data rate by three. However, now for every three bits we will have two signal line transitions, as well as a significant amount of additional circuitry required to generate and receive the pulse. Some power savings are gained because the forwarded clock can be shared among the bus. The tradeoffs of this system are investigated in the next sections.

4.3.1 Modulation Circuit

The modulation circuit used was presented in [20] and is shown in Figure 4.16. The rising edge of the system clock is delayed various amounts depending on the input data and used to construct the data pulse. Each of the delay elements is implemented as a number of series inverters, four for delay-4 blocks. The minimum data pulse width is 24 inverter delays, which corresponds to a zero input. As the input is increased to one, two, or three, the pulse width increases by eight additional inverter delays each time. Thus two bits are incorporated into the pulse width. An additional bit is encoded by varying the position in time of the rising edge of the data pulse compared to the forwarded clock pulse. The data rising edge is either delayed eight inverter delays or not at all, whereas the forwarded clock is delayed four inverter delays. A timing diagram illustrating the delay at various nodes in the modulator and the data and clock outputs is shown in Figure 4.17.

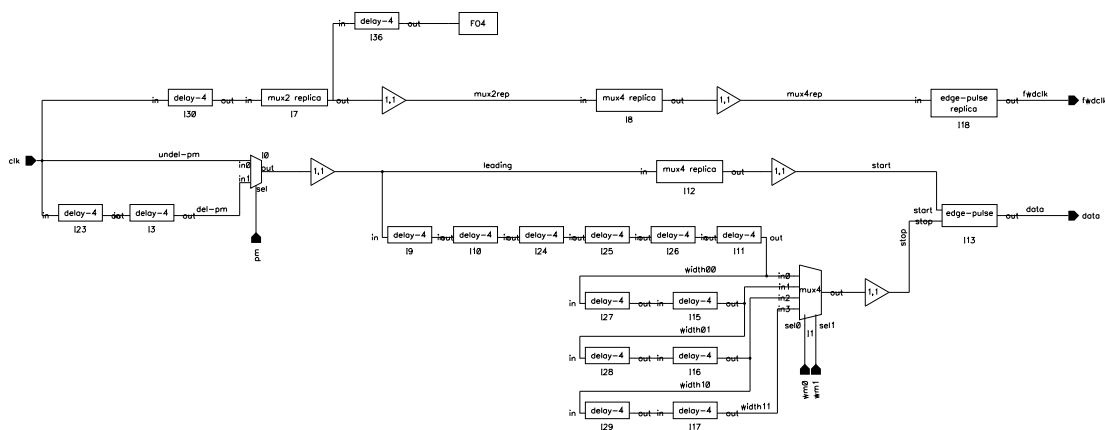


Figure 4.16: Pulse width modulator schematic

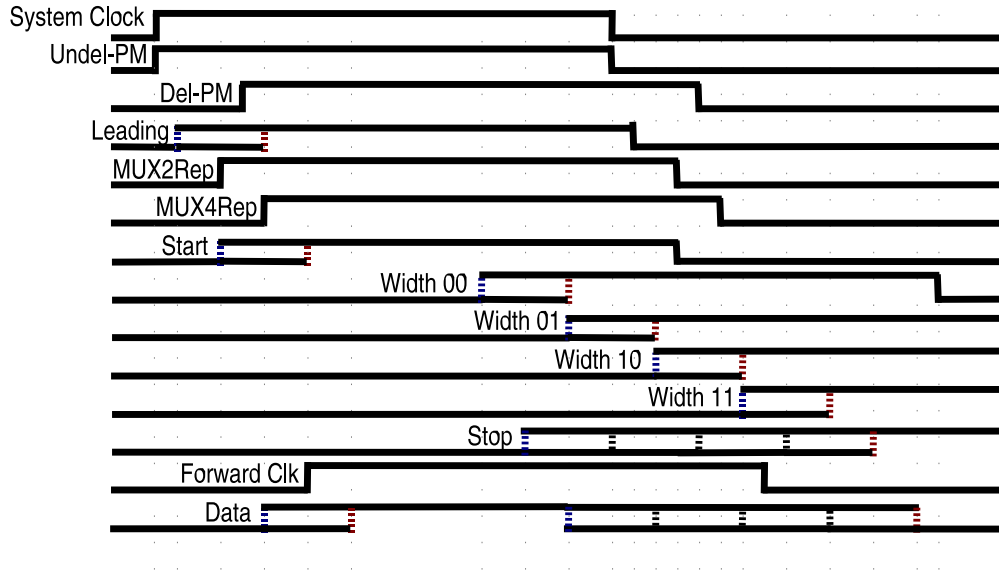


Figure 4.17: Pulse width modulator timing diagram

The width of the data pulse is entirely dependent on the amount of delay introduced in the inverter delay chain. This presents a serious challenge in advanced processes where the inverter delay is quite small. In our case, the inverter FO1 delay was simulated to be approximately 30 ps . Initially, the minimum pulse width was only twelve inverter delays. However, this pulse proved too narrow to reliably receive and decode. This problem was solved by doubling all of the delay paths, although this adds an unacceptable amount of additional power draw to the system. Using current starved inverters in the delay path is suggested for future implementations. In addition, even with increased delay the pulse was unable to pass through the 5 mm wire model because of the low pass filtering inherent in the wire. Again, increasing the delay of the inverters in the pulse chain or adding repeaters to break up the long wire should improve the transmission reliability.

4.3.2 Demodulation Circuit

The demodulation circuit for this system is shown in Figure 4.18 [20]. The position modulation bit is extracted by comparing the rising edge of the forwarded clock to the edge of the incoming data. This comparison was done using an arbiter circuit described in [20].

The width modulation bits are extracted by inverting the data pulse and comparing the falling edge (the rising edge in the inverted wave) to delayed versions of the non-inverted data pulse.

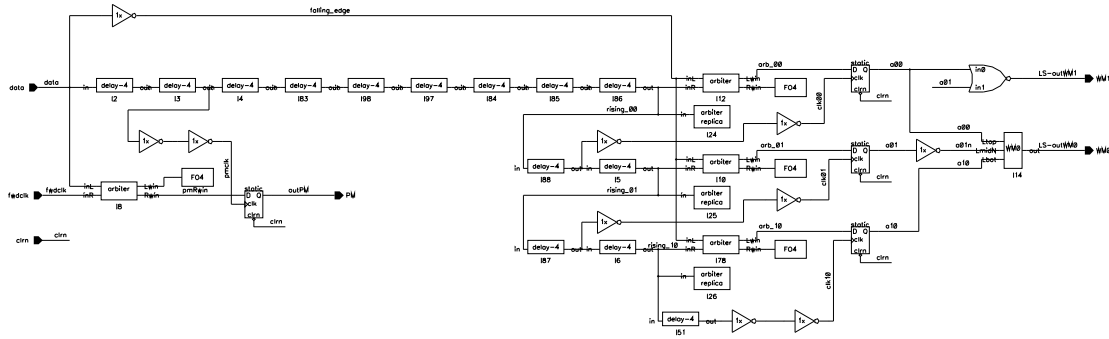
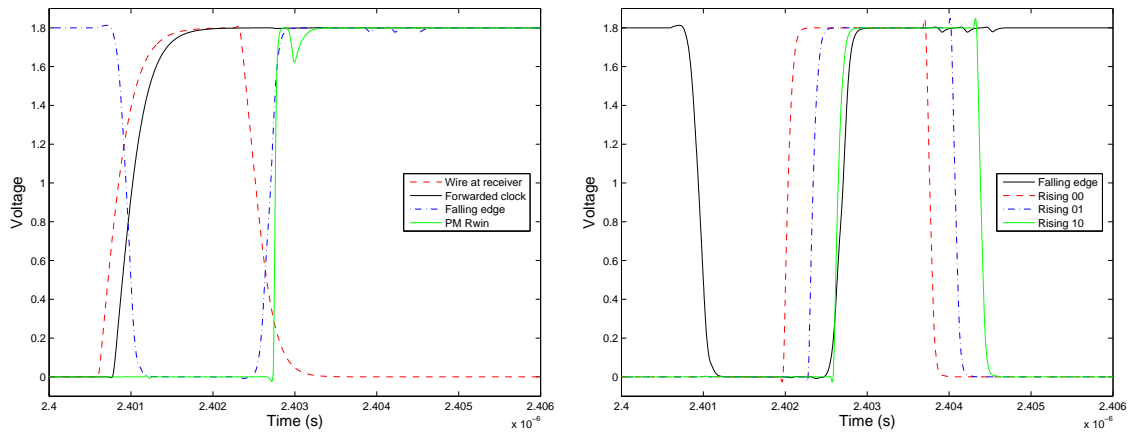


Figure 4.18: Pulse width demodulator schematic

4.3.3 Implementation and Comparison

The pulse width modulation system was simulated with the lumped capacitance for a 1 mm wire as the transmission path for the data pulse and the forwarded clock. A total data rate of 10 MHz was desired so the clock frequency was 3.3 MHz . Temperature was 27°C and the typical-typical transistor models were used. The power supply was swept from 1.8 V to 400 mV . The demodulator was able to correctly receive data for supply voltages down to 600 mV . At 400 mV , the demodulator could not correctly decode all of the data values. Figure 4.19 shows several simulated waveforms from the demodulator input and internal nodes. Figure 4.20 shows the average power of the pulse-width modulation system compared to the power of the low-swing driver and receiver pairs. This data indicates that combining pulse-width modulation with reduced supply voltages can potentially reduce the transmission power below the lower limits of low-swing interconnect alone.



(a) Forwarded clock and data waveforms at input to the demodulator

(b) Inverted data pulse compared to shifted versions of non-inverted pulse

Figure 4.19: Waveforms at the input and within the pulse width modulator showing timing of decoding operation




Figure 4.20: Pulse width modulation power compared to low-swing drivers and receivers at 10 Mb/s

Chapter 5

Conclusions and Future Work

5.1 Conclusions

Three areas were explored in this thesis. The first was a consideration of circuit based modeling for mechanical vibration energy harvesters. This work enables models of the generators to be built into circuit simulators, which will in turn allow for more integrated design of electronics for energy scavenging power supplies. Next, a memory design for a Distributed Arithmetic based filter was presented. This memory was designed to work with both an AC and a DC power supply to decrease the system dependence on voltage regulation for correct operation. The memory cells can hold data for up to 1.2 *ms* without refreshing and the sensing circuits require no DC bias currents. Finally, two methods of reducing interconnect power were described. The range of operation and power estimates for six low-swing driver and receiver combinations were determined. The lowest power implementation was able to send data at 10 *Mb/s* across a 1 *cm* wire using less than 2.5 μW average power. A preliminary investigation into a pulse-width modulation system was also initiated with some promising results. Data rates of 10 *Mb/s* were achieved in simulation on a 1 *mm* wire with average power of approximately 2.5 μW . This power includes the power used to transmit the reference clock. In a larger bus, the power per bit would decrease with each additional wire because the reference clock could be shared. Each of the technologies described can be directly applied to power reduction and design simplification for the development of wireless sensor nodes.

5.2 Future Work

In the process of working on this project, there have been many circuits and architectures that have seemed promising enough to merit further investigation. These will be described here as proposed future work.

5.2.1 Dual-Polarity Architecture

One of the earliest thoughts was that if it is better to get rid of the AC to DC converter and some of the power regulation electronics, would it not be best to get rid of all of the power regulation circuitry? The generator is connected into the electronics differentially and there is no explicit ground. It may be possible to have one set of electronics that operates on the positive phase of the harvester output, using the positive terminal as V_{DD} and the negative terminal as ground, as well as a copy of the electronics that operates on the negative phase of the harvester output, using the negative terminal as V_{DD} and the positive terminal as ground. The DRAM cells investigated in Chapter 3 could be modified such that the charge on the storage node is always relative to whichever power supply is in use. This would require some switches to shift the other connections, but otherwise minimal overhead. There would need to be a detection circuit that could tell when the power supply was close to switching polarities. At that point, any data or coefficients not being stored in DRAM would have to be written to DRAM so they could be passed to the opposite polarity circuits. Although the concept is fairly simple, the challenge would be in implementing the control needed to transition smoothly between the power supply phases without losing information.

5.2.2 Supply Voltage Monitoring and Adaptation

As mentioned in Section 3.1.3, it could potentially be very useful to adjust the refresh timing of the DRAM to account for the power supply voltage. This would allow the DRAM to be refreshed often or even continuously refreshed when the supply voltage is very low and to be refreshed much more sparsely when the supply voltage is high. The optimal refresh rate and the range of supply voltages that count as high or low will be system

dependent, but the trend of having weak writes with a low supply and stronger writes with a high supply will hold in all or nearly all cases. One possible way of determining the supply range would be to build a large counter that is continuously clocked by the system clock and calibrate the count values to match power supply voltages. It was suggested that a voltage reference, such as a circuit that conducts current for voltages above a certain value but not for those below the value, could be used in the first few power supply cycles to calibrate the counter to a few reference voltages and times in the supply waveform. If the system was calibrated once in the first few power supply cycles, it would probably provide a good reference for any circuits with a constant supply frequency, such as an energy scavenging sensor mounted on the base of an industrial machine that runs continuously. However, for any circuits that would work on a variable supply frequency, the calibration would have to be run periodically. This may not be problematic as the mechanical frequency will almost always be changing much more slowly than the electrical circuits are operating and the amount of data being processed would be fairly low, allowing for a relatively large calibration duty cycle. This type of operation may be ideal for sensing circuits placed in duct work or on bridges where the vibrations of the mechanical structure will change depending on when the furnace turns on or heavy traffic is crossing the bridge.

5.2.3 Adaptive Body Bias of DRAM Cell Write Path

The simulation results for the DRAM cell that uses a negative body bias to reduce leakage through the write path show that a constant negative bias hurts the cells ability to read out a one value. However, if the body bias were set adaptively, this cell may be able to achieve high write voltages to the storage node and then decrease the bias voltage to lock the charge onto the node. The easiest implementation could use two separate power supplies to set the body voltage: one for writing the cell and a more negative one for standby times or reading the cell. A more complex method may be to use a charge pump or some switching combined with the power supply oscillations to drive the bulk voltage below ground without needing an additional power supply. These circuits have not yet been designed in detail. However, implementing them may be quite reasonable by adapting some of the charge pump circuits used today.

5.2.4 DRAM Read and Write Improvements

Two possibilities for improving the read and write abilities of the DRAM cells would be to implement gate bootstrapping on the write transistor and to build a path that will track the DRAM's performance into the self-timing oscillator being used as the system clock. The bootstrapping circuit would be able to pull the write transistor's gate above V_{DD} and ensure that a strong one was written to the storage node. As shown by the performance of the transmission gate based memory cell, the strength of the write voltage has the greatest impact on the cell's ability to correctly output a one as the supply voltage decreases. When the memory cell is being read, the amount of time allowed for the output to be evaluated greatly affects the output voltage. Adding a path into the ring oscillator that will track the memory's read strength over supply voltage, process, and temperature could potentially make a large difference in the overall robustness of the memory system.

Appendix A

Layout

All layouts are drawn using MOSIS scalable CMOS rules for deep submicron in Cadence Virtuoso.




Figure A.1: Single DRAM memory cell

Bibliography

- [1] S. Roundy, P.K. Wright, and J. Rabaey. A study of low level vibrations as a power source for wireless sensor nodes. *The International Journal for the Computer and Telecommunications Industry*, 26(11):1131–1144, 2003.
- [2] S. Roundy, B. Otis, Y.H. Chee, J.M. Rabaey, and P. Wright. A 1.9ghz rf transmit beacon using environmentally scavenged energy. In *Proceedings of the 2003 International Symposium on Low Power Electronics and Design*, August 2003.
- [3] R. Amirtharajah and A.P. Chandrakasan. Self-powered signal processing using vibration-based power generation. *IEEE Journal of Solid-State Circuits*, 33(5), May 1998.
- [4] S. Meninger, J.O. Mur-Miranda, R. Amirtharajah, A.P. Chandrakasan, and J.H. Lang. Vibration-to-electric energy conversion. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 9(1), February 2001.
- [5] S. Roundy, P. Wright, and K. Pister. Micro-electrostatic vibration-to-electricity converters. In *Proceedings of IMECE2002*, November 2002.
- [6] R. Amirtharajah and A.P. Chandrakasan. A micropower programmable DSP using approximate signal processing based on distributed arithmetic. *IEEE Journal of Solid-State Circuits*, 39(2):337–47, February 2004.
- [7] J.M.H. Lee, S.C.L. Yuen, W.J. Li, and P.H.W. Leong. Development of an AA size energy transducer with micro resonators. In *IEEE International Symposium on Circuits and Systems*, pages IV–876–9, Thailand, May 2003.
- [8] S. Roundy and P.K. Wright. A piezoelectric vibration based generator for wireless electronics. *Smart Materials and Structures*, 13:1131–1142, August 2004.
- [9] Piezo Systems Inc. Datasheet for lead zirconate titanate, psi-5h-s4-enh. Website: <http://www.piezo.com>.
- [10] Engineering Fundamentals. Properties of common solid materials. Website: <http://www.efunda.com>.
- [11] J. Siebert. Pipelined datapaths for ac power supplies. Master’s thesis, University of California, Davis, 2005.
- [12] S.A. White. Applications of distributed arithmetic to digital signal processing: A tutorial review. *IEEE ASSP Magazine*, 1989.

- [13] R. Amirtharajah, T. Xanthopoulos, and A.P. Chandrakasan. Power scalable processing using distributed arithmetic. In *Proceedings of the 1999 International Symposium on Low Power Electronics and Design*, pages 170–175, 1999.
- [14] B. Zhou. Memory design for energy scalable reconfigurable logic. Master’s thesis, University of California, Davis, 2004.
- [15] T. Simon and A.P. Chandrakasan. An ultra low power adaptive wavelet video encoder with integrated memory. *IEEE Journal of Solid-State Circuits*, 35(4), 2000.
- [16] H. Zhang, V. George, and J.M. Rabaey. Low-swing on-chip signaling techniques: Effectiveness and robustness. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 8(3), June 2000.
- [17] Y. Nakagome, K. Itoh, M. Isoda, K. Takeuchi, and M. Aoki. Sub-1-v sing internal bus architecture for future low-power ulsi’s. *IEEE Journal of Solid-State Circuits*, 28(4), April 1993.
- [18] A.P. Dancy, R. Amirtharajah, and A.P. Chandrakasan. High-efficiency multiple-output dc-dc conversion for low-voltage systems. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 8(3), June 2000.
- [19] Tsmc 0.18um logic 1p6m salicide 1.8v/3.3v spice models. Technical report, Taiwan Semiconductor Manufacturing Co. LTD, 2004.
- [20] T. Simon, R. Amirtharajah, J.R. Benham, J.L. Critchlow, and T.F. Knight Jr. A 1.6gb/s/pair electromagnetically coupled multidrop bus using modulated signaling. *IEEE International Solid-State Circuits Conference*, 2003.