High Speed DPWM Switched Mode Supply and Control

By

TIMOTHY LOO
B.S. (University of California, Berkeley) 2005

MASTER'S REPORT

Submitted in partial satisfaction of the requirements for the degree of

MASTERS OF SCIENCE

in

Electrical Engineering and Computer Science

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

BERKELEY

Approved:

Chair, Dr. Seth Sanders

Member, Dr. Elad Alon

Committee in charge
2007

© Copyright by Timothy Loo 2007
All Rights Reserved

Contents

List of Figures

Chapter 1

Introduction

One growing market in electronics today is that of mobile communications. The desire of consumers to be constantly connected with one another drives the demand for wireless devices such as cell phones and wireless laptops. One important consideration when designing for mobile applications is power consumption. A mobile product, by its nature, has a limited supply of energy; therefore the less power it uses, the longer the device can operate. At the same time, however, these devices are inherently power hungry due to the large power needs for wireless communications.

Therefore, one large area of research is to improve the efficiency of these wireless transceivers, specifically that of the transmit power amplifier (PA). PAs are traditionally inefficient: Class A PAs can achieve 50% efficiency at best and on average achieve much less. A comprehensive overview of the inherent inefficiencies of power amplifers (PA) and past research to improve their efficiency can be found in [1]. Generally, PA's are inefficient due to their large constant quiescent current draw which leads to wasted power dissipation. Attempts have been made to reduce this quiescent current consumption, such as in Class B, C, and D PA topologies; however, they come at the expense of input-output linearity.

To complicate matters, PAs often operate well below their theoretical maximum efficiency because they do not always transmit at their maximum power output [1]. A PA will often back off its power output depending on its proximity to the receiver to minimize interference with other transceivers. Moreover, if the transmitted signal contains amplitude

modulation, then the voltage output will vary constantly, inherently leading to a varying power output. Because the power supplied to a Class A PA is independent of its output power in a constant supply topology, the effective PA efficiency will be much less than optimal.

Therefore, instead of modifying the PA, another area of research has focused on modifying the power supply to the PA to ensure that only the required amount of power is supplied in order to maximize efficiency. As summarized by [1], these techniques include adaptive current biasing [2], envelope tracking [3], average power tracking [4], and envelope elimination and recovery (EER) [5]. Adaptive current biasing varies the quiescent current drawn by the PA, according to the output, to maximize efficiency. Envelope and average power tracking solutions, on the other hand, attempt to vary the supply voltage to optimize the power supplied to the PA for a given power output. EER takes it a step further and takes away the amplitude modulation role from the PA entirely and rests it solely on the supply. The benefit is that alternative nonlinear, yet highly efficient, PA topologies can be used instead.

This project attempts to create an envelope tracking supply for the 802.11a standard. In this system, a feedforward control system detects the input RF amplitude signal and controls a high speed switched mode power converter so that the output supply precisely tracks the envelope amplitude of the RF output signal.

Chapter 2

System Overview

The goal of this project is to develop a high speed, high efficiency envelope tracking supply. To achieve this goal, a feedforward control algorithm is combined with a digitally controlled switched mode power supply. Unlike conventional regulators, which rely on feedback control, a feedforward approach was chosen due to the high bandwidths needed to ensure adequate tracking. Creating a feedback control system with enough loop bandwidth to track the anticipated 20MHz bandwidth of an 802.11a signal is difficult and may introduce complexities in regards to stability. A feedforward approach, on the other hand, does not suffer from this stability constraint; control relies entirely on the input signal and the predetermined load characteristics that the PA presents to the powertrain.

The second component of the project involves the design of a digital switched mode power supply. The supply itself consists of two major components: the digital pulse width modulator (DPWM) and the powertrain. The DPWM converts the digital envelope signal from the control algorithm into pulse widths that the powertrain, consisting of a half-bridge power network, can then use to convert back into a tracking supply voltage. Although class D supplies are highly efficient, they are not normally used for precise signal reproduction due to issues involving PWM precision and LC filter dynamics. To overcome this problem, this design incorporates techniques based on the work in reference [6]. In [6], Pascual designed a high fidelity class D audio amplifier with the use of oversampling, noise shaping, and a naturally sampled digital PWM to achieve accurate signal reproduction. First, Pascual

utilized an 8x oversampling and interpolation of a high resolution source audio data in order to support a sufficiently fast switching frequency at the powertrain to reject side band distortion caused by switching harmonics. Second, an 8 bit fourth order sigma delta modulation is used to reduce the signal resolution to a realizable level without raising the quantization noise floor in the band of interest [6] [7]. Finally, the uniformly sampled data is converted to a naturally sampled format through the use of a Lagrange approximation [6]. Not only does natural sampling more closely mimic the analog PWM process, but it also reduces significant in band distortion commonly found in uniformly sampled digital PWM [6]. To further increase signal fidelity, a dual edged sampling scheme may be used to transmit twice as much signal information without increasing the switching rate [6][7]. An illustration depicting the differences between a single edged uniformly sampled PWM and a dual edged naturally sampled PWM is shown in Figure 2.1. Using the above techniques, Pascual achieved a signal to noise ratio of over 60dB and a THD of 0.02% in a digital audio amplifier [6].

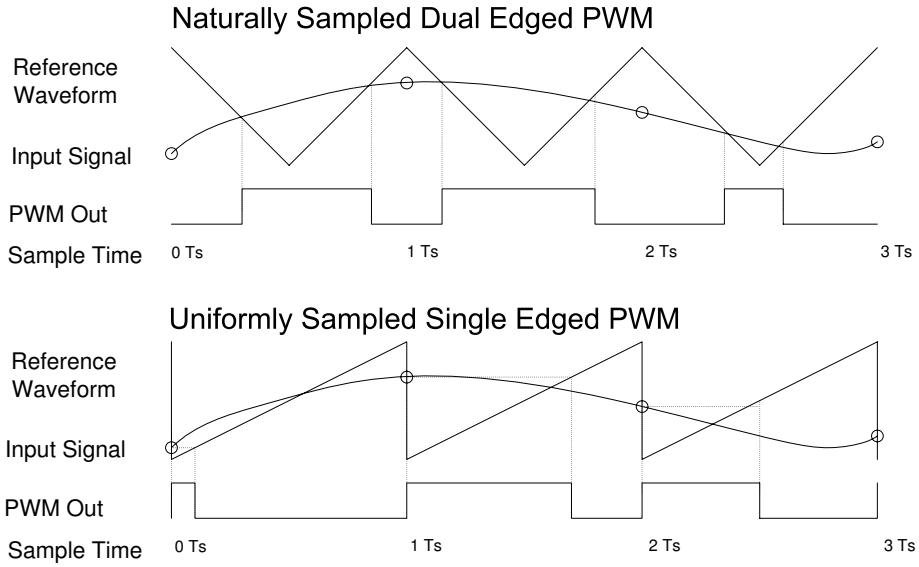

Figure 2.1: Comparison between Dual Edged Natural and Single Edged Uniform Sampling

A system level diagram of the proposed PA system is shown in Figure 2.2. It borrows from the data path and techniques introduced in reference [6], but is tailored

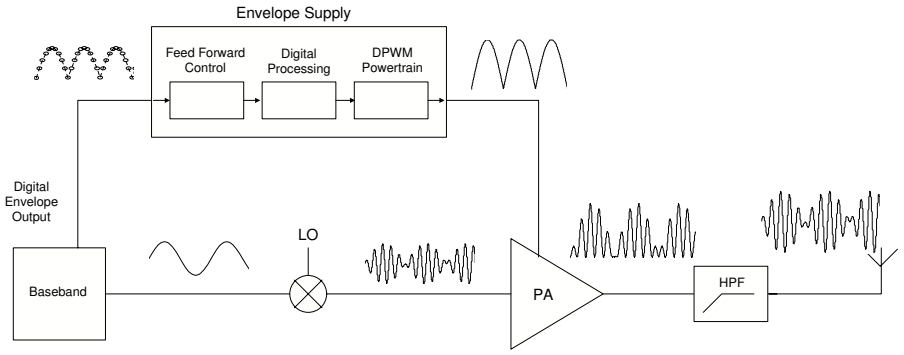

Figure 2.2: System Level Diagram of Proposed Envelope Modulation Scheme

for an RF modulation application. The modulation system begins by up sampling the amplitude RF data at 100MHz with 8 bits of resolution. With an estimated 802.11a signal bandwidth of 20MHz, a 100MHz sampling rate yields an oversampling ratio of 2.5. Although the oversampling ratio and resolution is not as high as in Pascual, supply tracking does not need to be as precise because the supply does not convey any transmitted data, the tracking envelope only sets the optimal PA supply voltage to enhance efficiency. The digital envelope signal is then preprocessed by the feedforward control block to ensure that the output is correct. After processing, the data is converted from amplitude modulation to digital pulse width modulation data through the use of a dual edged natural sampling block. The converted data is then down quantized from 8 to 5 bits via a 2nd order sigma delta modulator before being sent to the on chip power supply block. In the power supply, the digital pulse data is converted to analog pulse widths. These pulses are then used to drive a powertrain which translates the pulses back into a usable tracking supply.

In this system, the primary considerations are the DSP resolution, the DSP frequency, the DPWM resolution, and the switching frequency. The DSP resolution is the number of bits used in the digital control logic. This resolution determines the minimum quantization noise floor of the output signal. As discussed later, one should note that an N bit control system does not necessarily output a signal with a signal to quantization noise ratio (SQNR) of N bits: if an insufficient amount of bits is allocated, computation errors may appear at the output in the form of additional noise. In this work, a DSP resolution of

8 bits has been chosen as a trade-off between complexity and tracking accuracy. The DSP frequency is the rate at which the control system samples and computes the envelope data; for simplicity, it is set equal to the switching frequency. The DPWM resolution determines the number of supply voltages available, though this limitation may be overcome via noise shaping techniques mentioned above. Finally the switching frequency determines the time interval allocated for each pulse; ideally, this value is set such that a balance is made between minimizing distortion and minimizing switching losses. The DPWM resolution and the switching frequency are limited by the fabrication technology currently available. A 100 MHz 5 bit ring oscillator based DPWM requires an inverter delay of less than 150 ps; although this is twice as much as the approximate minimum inverter delay in a thick oxide 90nm CMOS process, this margin was chosen to guard against variations such as process and temperature. In this project, a thick oxide process was chosen over a standard one because the higher supply voltages supported by this process allowed for an extended frequency tuning range for the ring oscillator. With the inverter delay limit in mind, the selection of the DPWM resolution and the switching frequency is a tradeoff between minimizing the quantization noise and minimizing the distortion created by the switching harmonics.

One area of further research not addressed in this report is a realizable conversion block that translates the digital amplitude data representation to a digital PWM representation through a dual edged naturally sampling process. Pascual and others have explored this area and proposed several solutions. In [6], Pascual uses a Lagrange approximation equation to find the correct pulse rise and fall times. In this project, however, depending on the complexities of performing the approximation, it may be more efficient to perform the approximations a priori across the range of all possible inputs and store the precomputed values for later use in a look up table. The table size will vary greatly depending on the DSP resolution and the degree of accuracy required. A basic approximation using an 8 bit DSP and two sampling points will require $(2^8)^2$ or around 65000 entries while a highly accurate but memory intense 4 point approximation will require $(2^8)^4$ or over 4 billion entries. A middle ground between accuracy and resource usage may be found by varying the DSP resolution of the sample points used in the approximation; however, further investigation is still necessary.

Chapter 3

DSP Design

3.1 DSP Overview

The DSP portion of the system, shown in Figure 3.1, which includes the feedforward control, PWM conversion, and the sigma delta modulator, is implemented through Matlab's Simulink package. In future investigations, these Simulink blocks will be converted to hardware description language (HDL) code and programmed onto a Xilinx field programmable gate array (FPGA) via the BEE-XPS software designed by the Berkeley Emulation Engine (BEE) research group at the Berkeley Wireless Research Center (BWRC).

3.2 Feedforward Analysis

Because the power supply lacks feedback regulation, the envelope signal must be predistorted a priori to ensure that the desired output voltage can be recovered at the

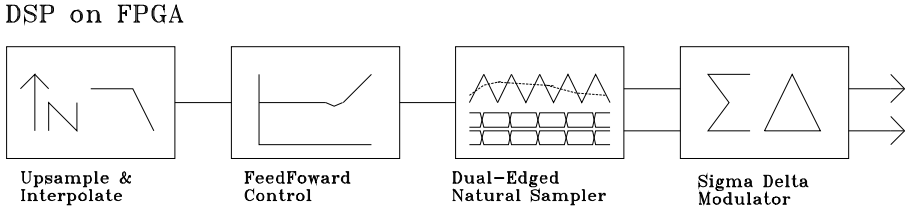

Figure 3.1: DSP Components

output of the LC filter. If an ideal linear DPWM and buck converter is assumed, the only component that requires compensation is the LC tank. For simplicity, the PA can be initially modeled as a simple resistor if we assume all the power supplied to the PA is transferred to the output load. The value of the resistor is chosen by following Ohm's law: $V = IR$, where V is the maximum output voltage and I is the average current supplied to the load. In this case, the transfer function to be compensated is a LRC tank, Figure 3.2, which is given in Equation 3.1, shown below:

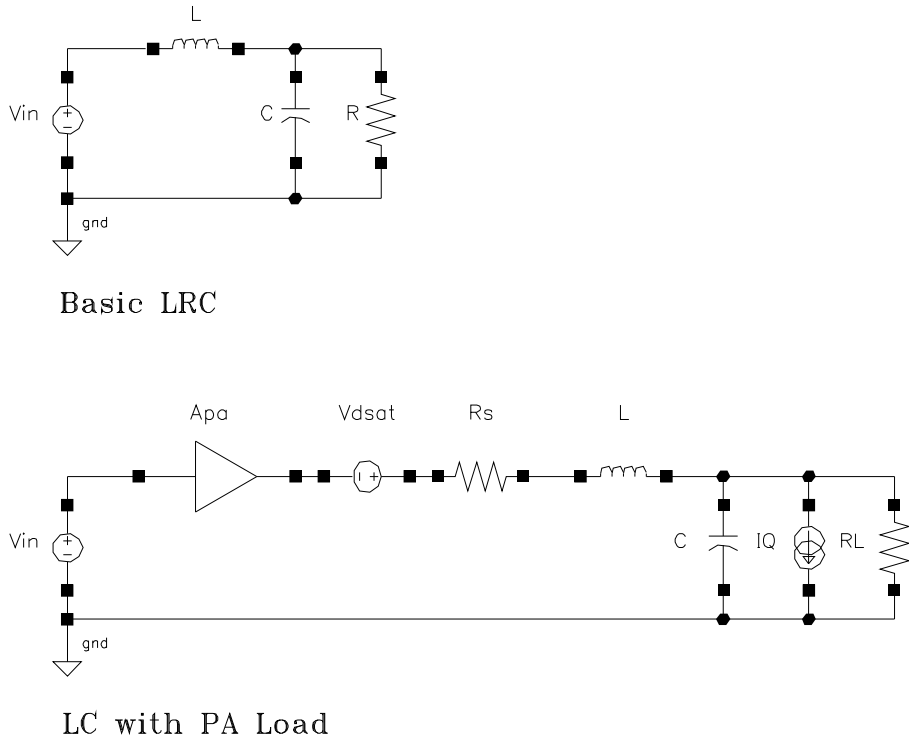

Figure 3.2: LCR Filter

$$H(s) = \frac{1}{1 + s\frac{L}{R} + s^2LC} \tag{3.1}$$

To cancel out the LRC transfer function, the inverse of Equation 3.1 is applied to the signal envelope:

$$H(s)^{-1} = 1 + s\frac{L}{R} + s^2LC \tag{3.2}$$

Figure 3.3 displays the transfer characteristics of Equation 3.1 and 3.2 as well as

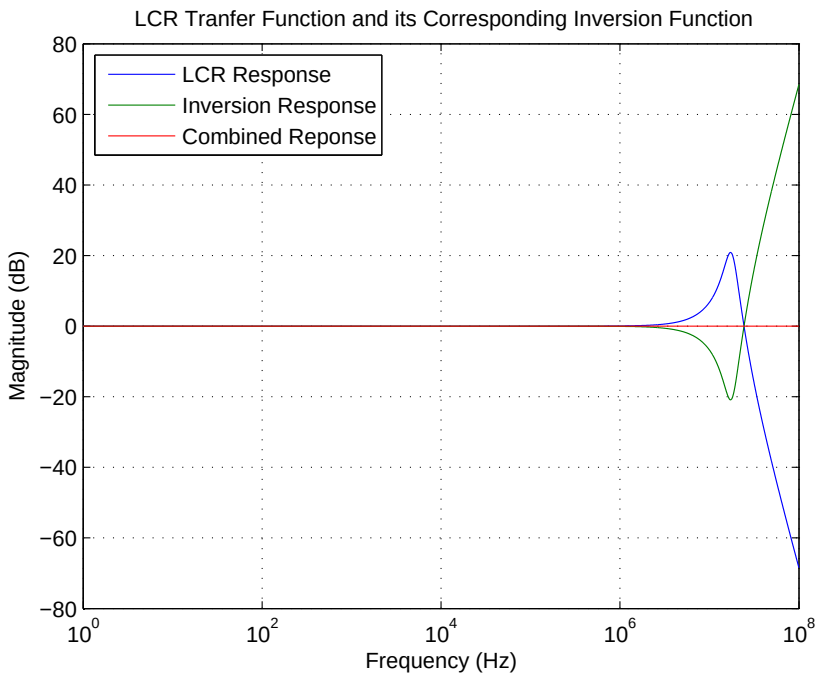

Figure 3.3: LCR Filter Inversion Plot

their combined output. Note that in this figure, the LC tank displays peaking due to the large Q value of 3 chosen for this filter. Normally, a low Q such as 0.5 value is chosen for the reasons stated in [1]; however, for illustrative purposes a larger Q has been used.

In an ideal inversion, the magnitude response of this function would approach infinity as the frequency approaches infinity due to the presence of multiple zeros, implemented through derivatives, as denoted by s. However, in a real implementation, this behavior is not possible to achieve; fortunately, implementing this infinite gain is not necessary either. The finite gain is due to the fixed supply rail. If the input signal is too large, the resulting output signal will saturate at the supply and the equivalent gain that the signal sees will fall, leading to an incomplete inversion. Fortunately, due to the band limited nature of the envelope signal, the transfer function only needs to be accurate up to 20 MHz; any signals above this band would have been previously filtered out by anti-aliasing filters. Although the system no longer needs to achieve infinite gain, care must still be taken to ensure that the required gain within the band of interest does not cause the resulting output signal to

exceed the rail.

This consideration affects the choice of the LC corner for the tank. Ideally, the LC corner would be set well below the bandwidth of the signal in order to gain additional ripple attenuation; the inversion system would then be used to gain the signal accordingly so that the output response is still unity. However, the minimum LC corner is now set by the supply rail: the corner cannot be so low that the maximum output voltage generated by the inversion function cannot be accommodated by the supply. Conversely, the Q of the LC tank is also limited by the performance of the switching supply. An excessively high filter Q will lead to large peaking that may require the control system to attenuate the output signal below the noise floor, thus clipping the signal on the low end.

A more accurate transfer function would better model the loading effects of the PA, also shown in Figure 3.2. As described in reference [8], to model the PA, two characteristics should be considered. First, a constant current source may be added to the output to represent the quiescent current draw necessary to properly bias the PA. The resistor now represents the load that the PA is designed to drive. Second, the voltage on the output of the PA is not necessarily the same as the input. Because the PA may have non unity gain, the envelope signal must be scaled to reflect this effect; otherwise, the supply may be too low and the output may rail. Moreover, the PA must be kept in saturation to maintain gain and minimize distortion. To ensure that this condition is met, the output envelope should be increased such that the output of the PA supply always remains one V_{DSAT} above ground. Finally, a source resistor should be added to model the high and low side on resistance of the powertrain switches. Using KCL to solve, the new inversion function is as follows:

$$Vx(t) = A_{PA}V_{IN}(t) + V_{DSAT} \tag{3.3}$$

$$V_{OUT}(t) = [(R_S I_Q + \frac{R_S}{R_L}Vx(t) + \frac{L}{R_L}\frac{d}{dt}Vx(t) + R_S C\frac{d}{dt}Vx(t) + LC\frac{d^2}{dt^2}Vx(t))] \tag{3.4}$$

A_{PA} represents the PA gain, V_{DSAT} represents the minimum drain source voltage to operate the PA in saturation, and I_Q represents the quiescent current of the PA. R_S, R_L, L, and C represent the switch resistance, load resistance, filter inductance, and filter capacitance, respectively. Because the input and output signal levels to the PA are not necessary small

signals, the voltage gain may become non-linear and the quiescent current may no longer be constant. To handle this case, as suggested in [8], a periodic steady state (PSS) analysis across PA output power may be performed in Spectre to characterize this nonlinear behavior. The nonlinear gain A_{PA} and quiescent current draw I_Q may then be modeled as a voltage controlled voltage and current source, respectively, which follows a transfer function characterized by the PSS analysis. Further research may still be needed to better understand PA loading behavior, for example, through a PSS analysis combined with a periodic S parameter (PSP) analysis, to further enhance the inversion equation and to implement a PA model into Matlab Simulink simulations.

3.3 Implementation Considerations

The conversion from a continuous time inversion transfer function to an implementable digital counterpart is not completely straightforward. Although operations such as addition and multiplication are easily realized, others such as the derivative (represented by the s in Equation 3.2), are not as easy. The most naive method to implement a derivative is through the following transfer function:

$$H(z) = Fs * (1 - z^{-1}) \tag{3.5}$$

where F_s is the sampling frequency. The equation mimics the operation of a continuous time derivative and the results are quite similar, especially if F_s approaches infinity. The results of an inversion simulation using this simple digital derivative is shown in Figure 3.4. In this simulation, a sample 802.11a waveform is fed into the inversion DSP block. The output of this block is then fed into a continuous time LRC filter, which the DSP block has been designed to invert, and the original and resulting signal are then compared. Ideally, the inversion will be perfect and the original and resulting signal will match. In Figure 3.4, the stepped waveform represents the sampled envelope signal that is fed into the inversion DSP block and the smooth waveform represents the continuous time output waveform from the LRC filter. Any constant delay between the input of the inversion block and the output of the LRC filter has been manually compensated. As shown in Figure 3.4, the inversion is

not very accurate, and an error of around 100mV between the two signals appears.

The primary cause of this error is the low sampling rate compared to the maximum bandwidth of the input signal: in this DSP, the oversampling frequency is only 2.5 times larger than the expected bandwidth of the envelope signal. The low oversampling ratio (OSR) results in two sources of degradation. First, the accuracy of the derivative shown in Equation 3.5 is highly dependent on the OSR: a higher OSR results in a more accurate derivative. Moreover, the derivative performed in Equation 3.5 can be thought of as a 1 tap finite impulse response (FIR) filter; as such, this filter introduces a half clock cycle delay. Errors are created when the half clock cycle delayed data from the derivative block is combined together with undelayed data through additions and multiplications later in the signal path. This non-integer delay, however, can be compensated through linear all pass filters as suggested in [9], at the cost of increased complexity. Second, a zero order hold (ZOH) operation that inherently occurs in the conversion from a discrete time to a continuous time signal introduces an magnitude droop on the control signal, similar to the effects of a low pass filter. As noted in [9], the droop can be modeled with the following equation:

$$H_{ZOH}(f) = \frac{1 - e^{-j2\pi fT}}{j2\pi fT} = e^{-j\pi fT}\,\mathrm{sinc}(fT) \tag{3.6}$$

where f represents the frequency of interest and T represents the sampling period. As noted in Equation 3.6, at 20MHz, with a sampling frequency of 100MHz, the inversion signal will experience an attenuation of 0.6dB. From the above information, it can be implied that a simple solution is to increase the sampling frequency. By doubling the sampling rate, the derivatives will become more accurate, the absolute value of the half cycle delay will decrease by a factor of 2, and the ZOH attenuation will decrease to only 0.14dB. With a sampling rate of 200MHz, as shown in Figure 3.5, the error drops to less than 40mV. However, increasing the sampling frequency is not always desirable as it increases the power consumption of the DSP.

Another solution involves the use of a more advanced derivative such as an FIR derivative described in [9]. In simulation, a simple 6 tap FIR filter with the following

coefficients: [-0.1142 0.1528 0.2985 0 -0.2985 -0.1528 0.1142] (generated through MatLab's FDAtoolbox) is used to perform the derivative. A multitap FIR filter allows for a more accurate derivative even when operating at a lower OSR. In addition, with an even number FIR filter, this derivative does not suffer from a half cycle delay exhibited in the previous design. As shown in Figure 3.6, the inversion performance now approaches that of a 200MHz DSP, with an error of approximately 50-60mV. Performance can likely be further improved with a higher order FIR derivative at the cost of increased complexity.

Although errors in the power supply envelope do not affect the transmitted signal, these errors are still an important consideration because they will degrade the efficiency of the PA. Supply undershoots require that the bias level of the supply be increased to maintain the PA in saturation and will result in a constant penalty in efficiency. Supply overshoots, on the other hand, results in excess power dissipation during the overshoot itself. Errors in the supply consists of two components: errors resulting from an imperfect inversion of the signal envelope in the DSP block, mentioned above, and errors resulting from the switching ripple introduced in the powertrain. Switching ripple is commonly reduced by lowering the corner frequency of the LC filter; however, an excessive reduction will increase the filter dynamics that the DSP must compensate and will invariably lead to increased DSP error. DSP inversion errors can be reduced by increasing the algorithm complexity through the techniques discussed above; however, DSP complexity cannot be increased arbitrarily such that the efficiency gained in the powertrain is offset by the increased power consumption of the DSP.

3.4 Optimization

Because the goal of the envelope tracking power supply is to minimize PA power usage, it is important that the DSP portion consumes as little power as possible. When designing a DSP, the bit widths of the data path are an important consideration because this parameter affects the performance and impacts the power consumption of the logic. Ideally, a DSP would sample, process, and output the envelope data using a constant number of bits predetermined by the desired output signal accuracy. In reality, however, the actual

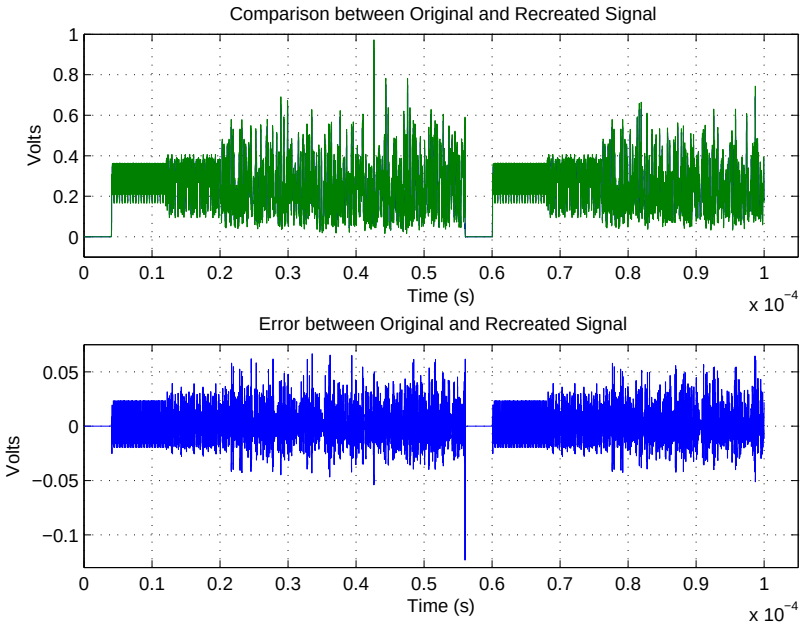

(a) 100us Sample

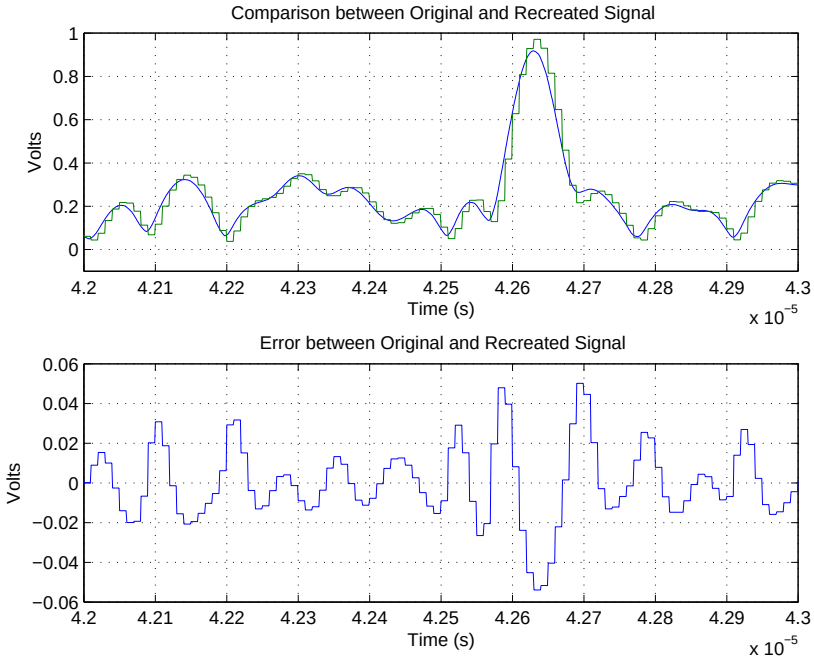

(b) 1us Sampler

Figure 3.4: 100MHz LCR Inversion with Simple Derivative

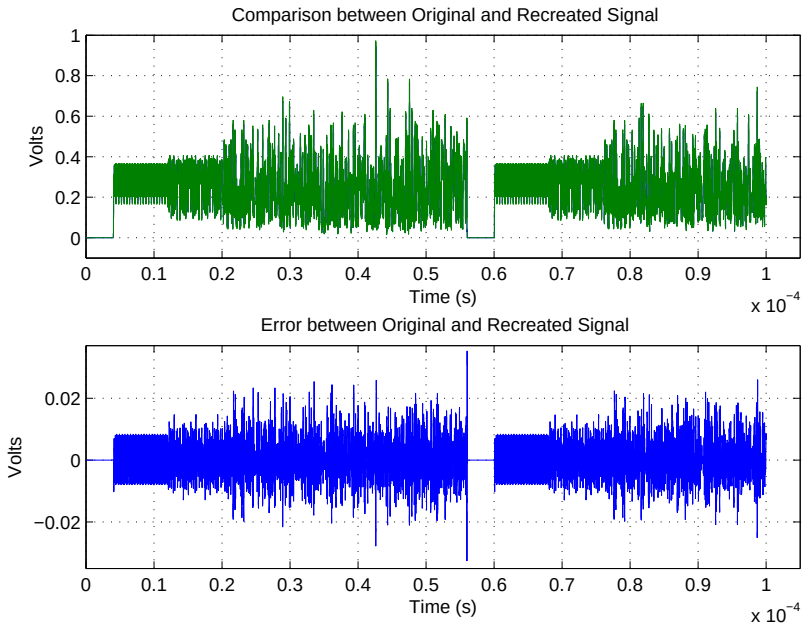

(a) 100us Sample

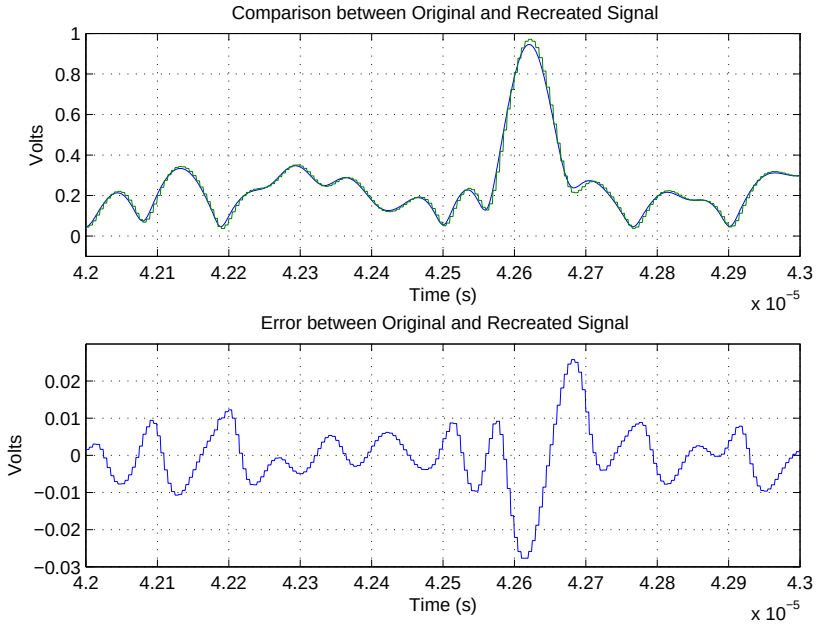

(b) 1us Sampler

Figure 3.5: 200MHz LCR Inversion with Simple Derivative

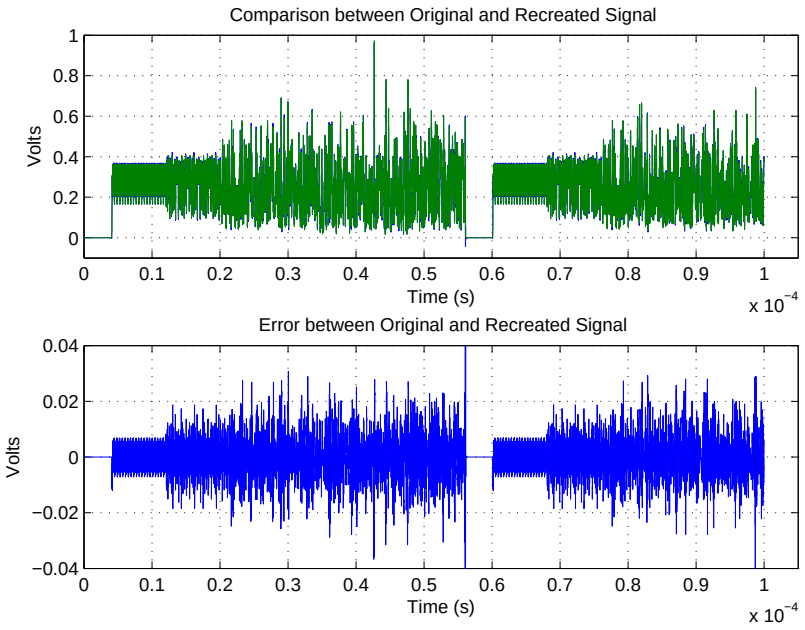

(a) 100us Sample

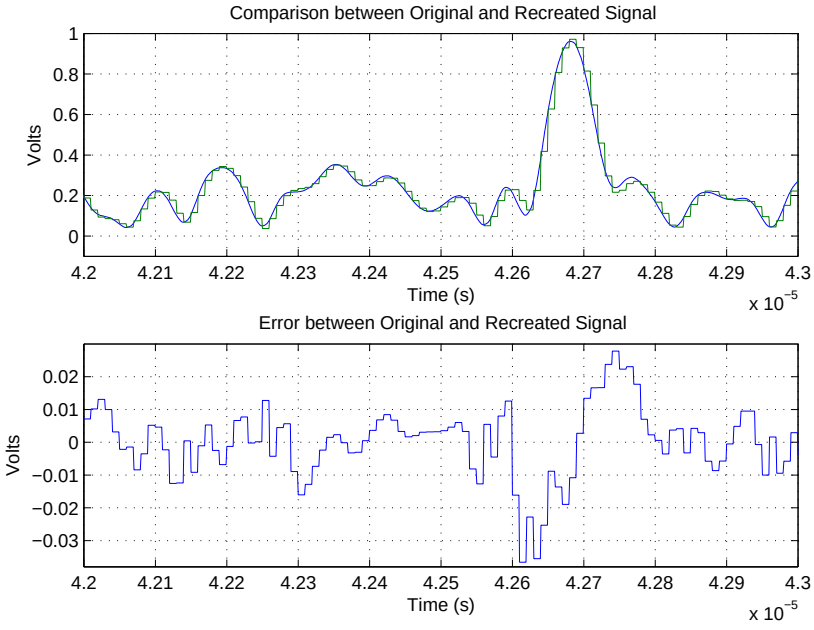

(b) 1us Sample

Figure 3.6: 100MHz LCR Inversion with FIR Derivative

implementation may require a much wider range depending on the arithmetic operations. Arithmetic operations such as constant multipliers may scale the signal to such a degree that it can no longer be represented with the available bits.

Therefore, the first method used to minimize the bit widths is to limit the range of the constants used so that the resulting output will not clip or fall below the quantization noise floor. One method is to combine adjacent constant multiplications together. The algorithm contains the equations $I = C\frac{dv}{dt}$ and $V = L\frac{di}{dt}$. To perform these derivatives, the differences between two consecutive data points are calculated and scaled by the sampling frequency of 100 MHz. This value is then multiplied by the load capacitance, which is a small value on the order of nano-Farads. Because these two values vary greatly from 100e6 to 1e-9, an enormous number of bits would be needed to carry out the two operations separately. By combining these two constants together, the scaling constant becomes 0.1 and the operation requires significantly less bits. Because the operations in the algorithm are for the most part linear, the combination of constants need not only occur when they are located sequentially one after each other; they can occur anywhere along the data path without changing the transfer function. Moreover, constants can also be combined with an artificially added scaling constant in order to limit the output signal within a certain range, as long as the scaling operation is reversed further down the data path, as shown in Figure 3.7.

The second method is to optimize the allocation of bits used to represent the data after each arithmetic block. In the Xilinx DSP blocks, data is represented in fixed point notation (FPN). A fixed point represented value can be converted to a rational decimal point value as follows:

$$\begin{aligned} X_{decimal} =& A_N * 2^N + A_{N-1} * 2^{N-1} + \cdots + A_1 * 2^1 + A_0 * 2^0 + \\ & B_1 * 2^{-1} + B_2 * 2^{-2} + \cdots + B_{M-1} * 2^{-(M-1)} + B_M * 2^{-M} \end{aligned} \tag{3.7}$$

N + M represents the total number of bits, where A represents the integer bits coefficients and B represents the fractional bits. As seen in Equation 3.7, the allocation of bits is a tradeoff between accuracy and range: adding an extra bit to the fractional portion allows for

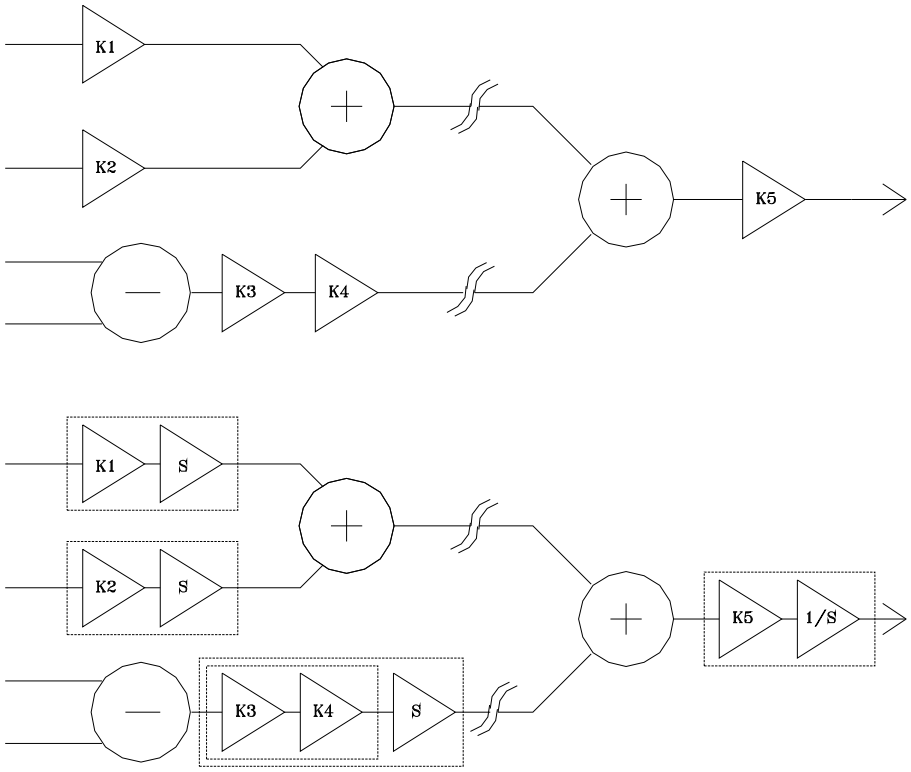

Figure 3.7: Illustration of Constant Scaling

a smaller LSB, but limits the maximum size of the data. Therefore, the goal of optimization is to maximize the allocation of bits to the fractional bin without clipping the signal. To perform this, the maximum signal range for the output of every block is calculated, starting with the input of the DSP and progressing along the data path until the output is reached. Once the maximum range is determined for each stage, the minimum number of integer bits that can represent the signal without clipping can then be found and assigned. To further improve signal representation, the signal may also be checked to determine if it can be represented as a signed or unsigned integer. If the signal is guaranteed to always remain positive, then the sign bit can be reallocated to the fractional bits as well.

The results of the DSP optimization are shown below in Figures 3.8, 3.9, and 3.10. The following three figures display comparisons between the output of a DSP using floating point resolution and a DSP using fixed point notation. A sample 802.11a waveform is processed in its respective DSP block and is then filtered by a continuous time LCR filter implemented via state space equations. In Figure 3.8 and 3.9, the DSP was allocated 8 bits. Note that the derivative used in these DSP simulations consists of a simple one tap FIR filter, chosen for symplicity, shown in Equation 3.5. In the unoptimized 8 bit DSP, shown in Figure 3.8, the algorithm did not undergo any constant rescaling optimization; however, the bit allocation after each mathematical operation was adjusted to ensure that signal clipping did not occur. An unoptimized 8 bit DSP displayed an error of approximately 130mV. Figure 3.9 displays an 8 bit DSP output with constant rescaling optimization. This DSP exhibits an error of 60mV, an improvement of one extra bit. If a more accurate DSP is required, the bitwidth of the DSP can be increased to 9 bits; as shown in Figure 3.10, the majority of the error is contained within 30mV. Although the error exhibited by the DSP is more than an LSB (an 8 bit data path could ideally yield an error of 4mV), it has been observed in these simulations that the error still follows the expected behavior of being halved with each additional bit used.

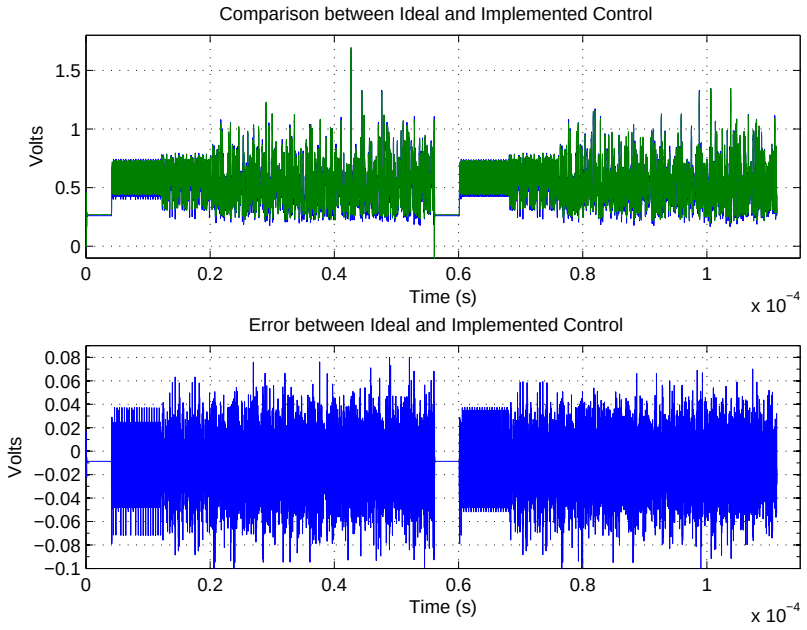

(a) 100us Sample

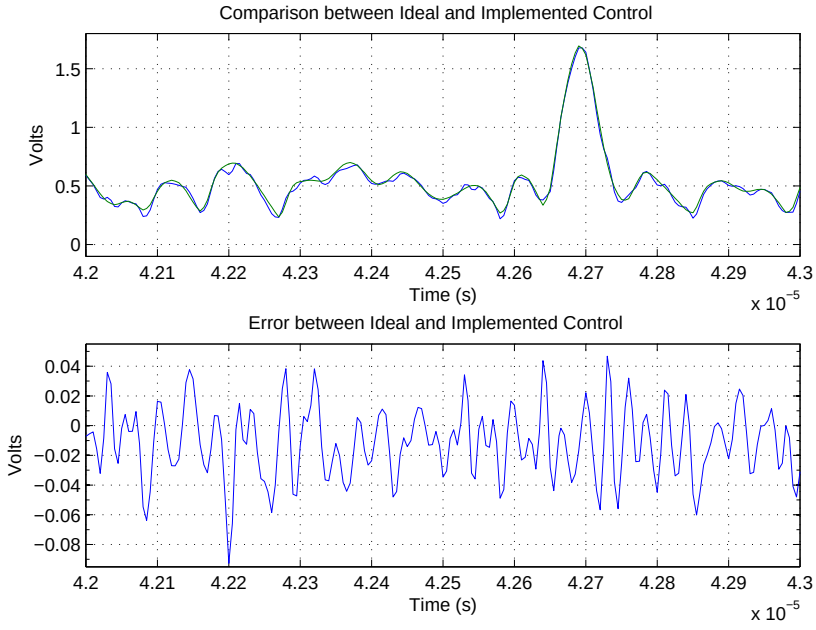

(b) 1us Sample

Figure 3.8: 8 bit DSP without optimization

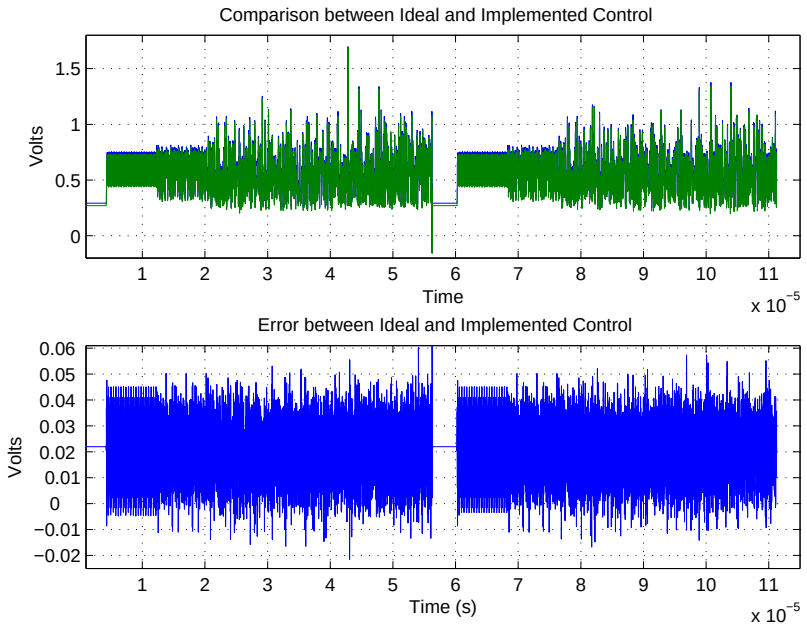

(a) 100us Sample

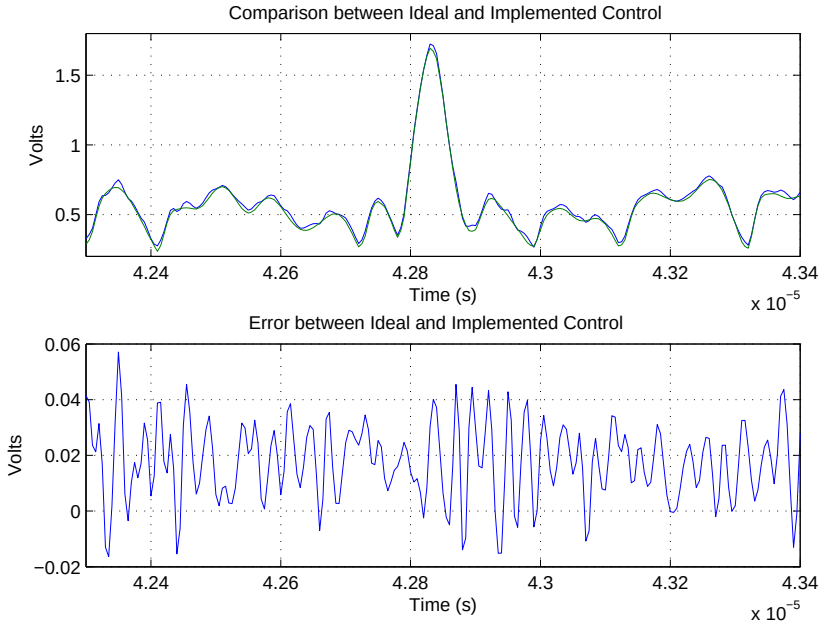

(b) 1us Sample

Figure 3.9: 8 bit DSP with optimization

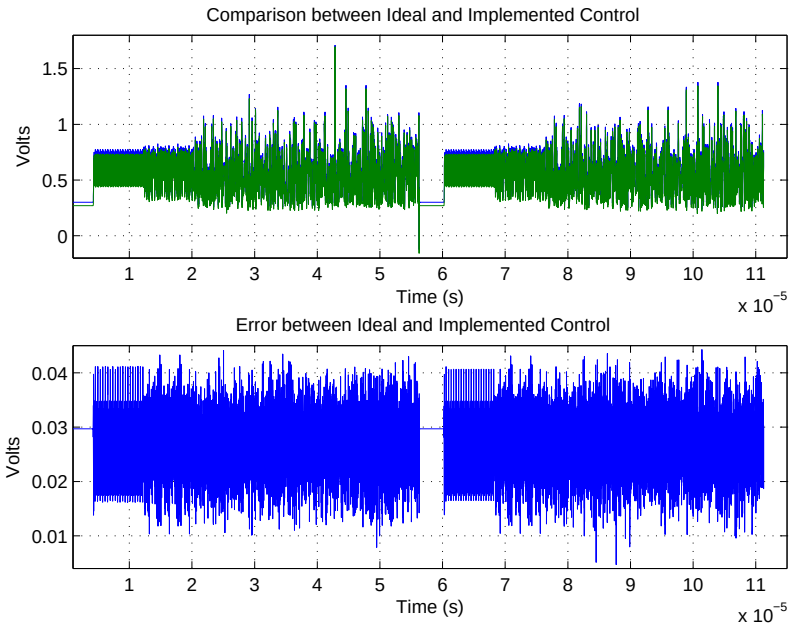

(a) 100us Sample

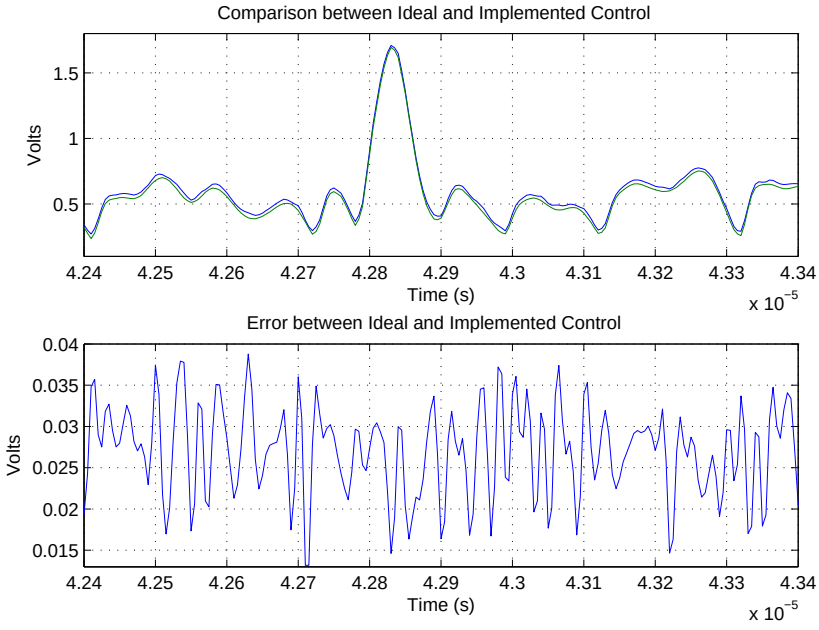

(b) 1us Sample

Figure 3.10: 9 bit DSP with optimization

3.5 Sigma Delta Design

The final stage in the DSP processing is the sigma delta processor. As noted in [6], the sigma delta allows the bit width of the output signal to the DPWM to be reduced to a realizable level without a significant loss in the signal to noise ratio (SNR). The sigma delta shapes the quantization noise such that the in band noise is pushed out of band where it can be subsequently filtered out. With the use of the feedforward correction algorithm, the effects of the sigma delta can be further improved. As mentioned above, the cutoff frequency of the buck converter's LRC filter can be set below the desired signal's bandwidth. By lowering the corner, additional quantization noise near the signal's bandwidth can be filtered out without degrading the desired signal due to compensation by the feedforward processing.

In this design, the sigma delta function is implemented using a multi-bit 2nd order error feedback design described in [10] and [6] and illustrated in Figure 3.11. A second order modulator is chosen for its additional noise shaping capability compared to a first order noise shaper and its resistance to limit cycles without introducing too much complexity with respect to stability and overflow handling. The actual implementation of the sigma delta is shown in Figure 3.12. Due to the dual edged PWM design, the sigma delta must process two sets of data: the rising edge and falling edge duty cycle command. Ideally, the sigma delta would operate at twice the DSP frequency, either through dual clocks or a dual-edged clocking scheme: the input to the sigma delta would alternate between the rising and falling command at every clock cycle. However, because this work is limited to one singled-edge triggered clock due to hardware and toolflow limitations, both command data must be processed together within one cycle. To accomplish this, two identical sigma delta modulators are connected together so that they operate in parallel. The rising command is processed by the first modulator and the resulting error signal, which ideally would have been fed back to itself to process the falling data in the next clock cycle, is now passed to the second modulator to be processed with the falling data. Like the first modulator, the error signal of the second modulator is fed back to the first modulator to be processed with the rising command data in the next clock cycle and the process is repeated.

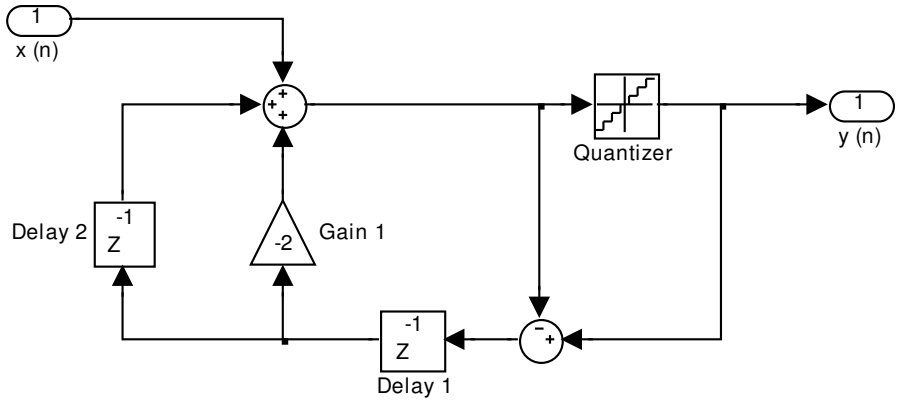

Figure 3.11: Second Order Error Feedback Sigma Delta Modulator

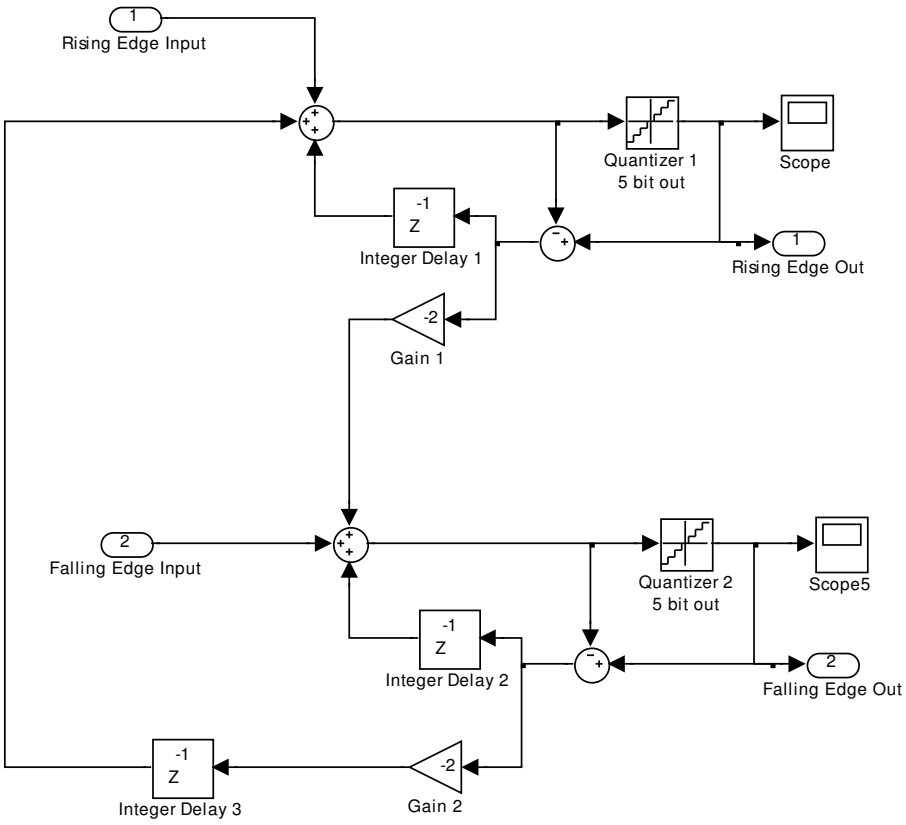

Figure 3.12: Dual Edged Sigma Delta Modulator

Simulation results of this block are shown in Figure 3.13 and in Figure 3.14. In both figures, a 500KHz sine wave is fed through an ideal feedforward control block. The resulting control signal is converted from an amplitude to a pulse width representation via a dual edged natural sampler. The pulse width data is then fed to a 5 bit dual edged DPWM which generates pulse widths for the powertrain. The entire simulation occurs within the MatLab Simulink environment. The first plot in Figure 3.13 illustrates the powertrain output; note that the limited quantization levels are quite noticeable. The second plot in Figure 3.13 illustrates the output with a sigma delta modulator; the sigma delta modulator increases the resolution such that the limited DPWM resolution is no longer noticeable. In both plots, the outputs are superimposed onto an ideal 500KHz sine wave for reference. Note that the ripples in both plots in Figure 3.14 are due to the switching harmonic added by the powertrain. These results suggests that in designs with relatively low quantization levels, even a modest increase in SNR provided by the sigma delta may yield a noticeable boost in tracking accuracy when viewed through the perspective of a transient simulation.

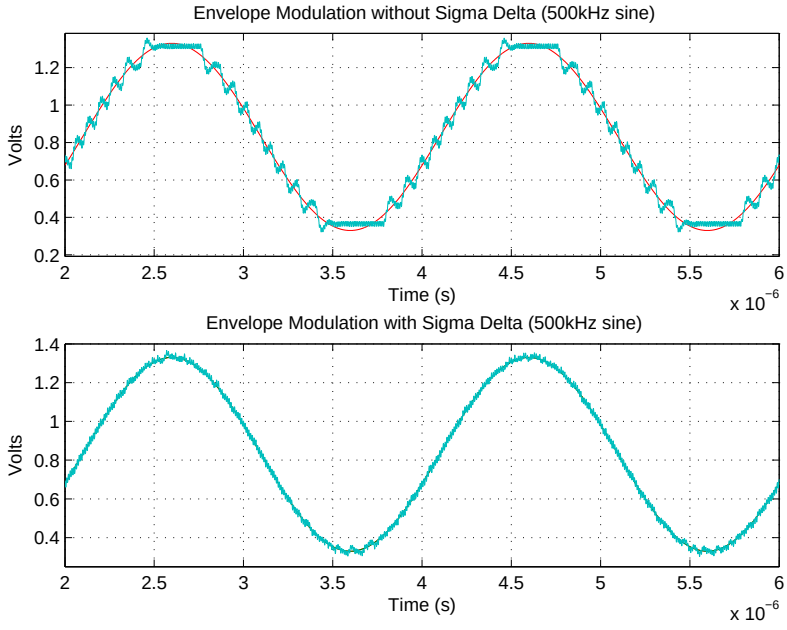

Figure 3.13: Modulator output with and without a Sigma Delta Modulator

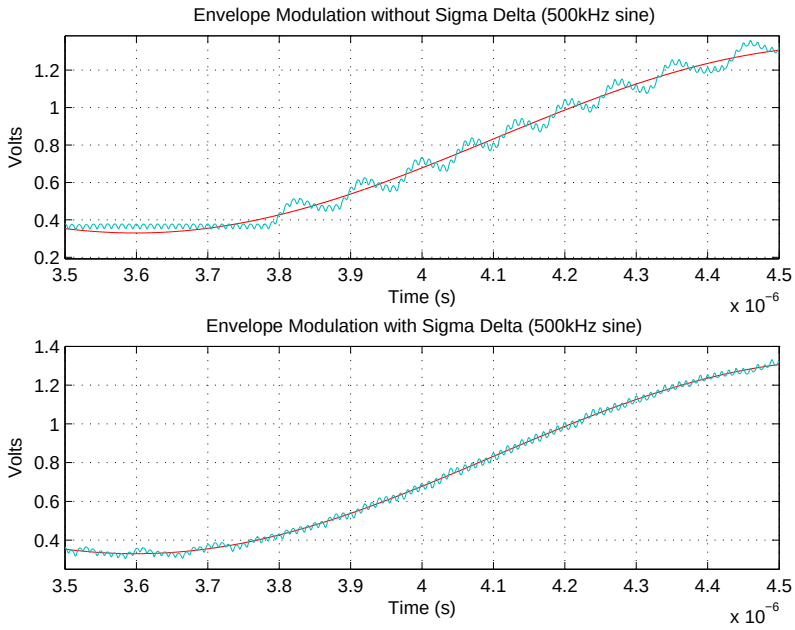

Figure 3.14: Zoomed in Comparison of the Sigma Delta Modulator

Chapter 4

On Chip Circuit Design

4.1 Introduction

The on chip portion of the envelope tracking system, shown in Figure 4.1, takes in the 5 bit digital control signal and converts it to the appropriate voltage to supply the radio frequency power amplifier (RFPA). This circuit consists of two main components: the digital pulse width modulator (DPWM) and the powertrain. The DPWM converts the digital control words into the respective pulse widths. The powertrain then takes the pulse widths and converts them back to the envelope tracking supply voltage. The chip has been fabricated in a 90nm CMOS technology.

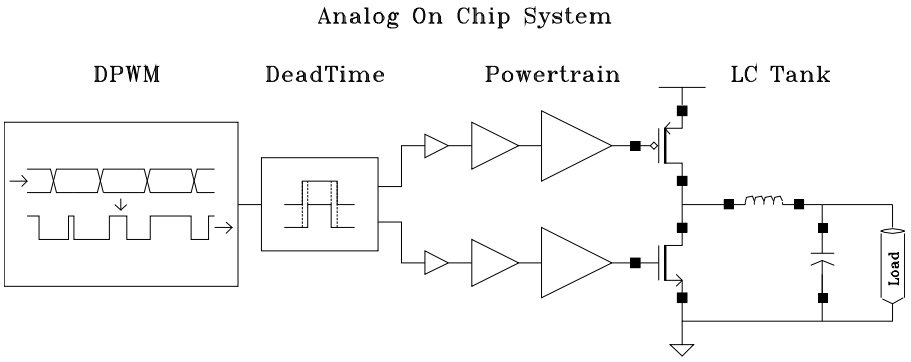

Figure 4.1: System Diagram of the on Chip Datapath

4.2 Ring MUX DPWM

Several different types of architectures can be used to implement the DPWM. Three such topologies have been previously investigated in J. Xiao's PhD Thesis [11]: the counter comparator, the tapped-delay, and the ring oscillator multiplexer (MUX) DPWM. As noted in [11], the counter comparator is unsuitable due to its high power consumption and high clock frequency requirement. The tapped delay line, consisting of an external clock driving a chain of delay cells, consumes less power; however, this implementation requires additional hardware to generate a clock and to control the delay cells with a feedback loop. Moreover, [11] mentions that the tapped delay line is also unsuitable due to the strict symmetry requirements necessary for precise delay matching between phases in multiphase applications [11]. The last topology, chosen for this design, is the ring oscillator MUX. J. Xiao's investigation found that this topology has both the benefit of low power consumption and an inherently symmetric structure beneficial for multiphase operation.

The ring oscillator MUX DPWM consists of three major components: the ring oscillator, the multiplexer bank, and the pulse generator. As illustrated in Figure 4.2, the ring oscillator generates a sequence of evenly spaced rising edges. For an N bit DPWM running at a $1/T_s$ clock, each edge is spaced approximately $T_s/2^{N+1}$ apart. The mux bank selects two command edges that indicates when the pulse should rise and fall within a clock cycle. The two command edges are then used by the pulse generator to create the desired pulse.

4.2.1 Ring Oscillator

As mentioned above, the design calls for a 5 bit double edged DPWM running at 100MHz. In a double edged pulse, both the rising and falling edge of the pulse can be controlled independently, with each edge having its own 5 bits of resolution. To generate this pulse, 64 taps are needed: the first 32 supply the rising edge of the pulse; the latter 32 supply the falling edge. One important implication in a double edged design is that an even number of taps is needed if each edge were to have an equal amount of resolution. However, this requirement precludes the use of single-ended ring oscillators which require

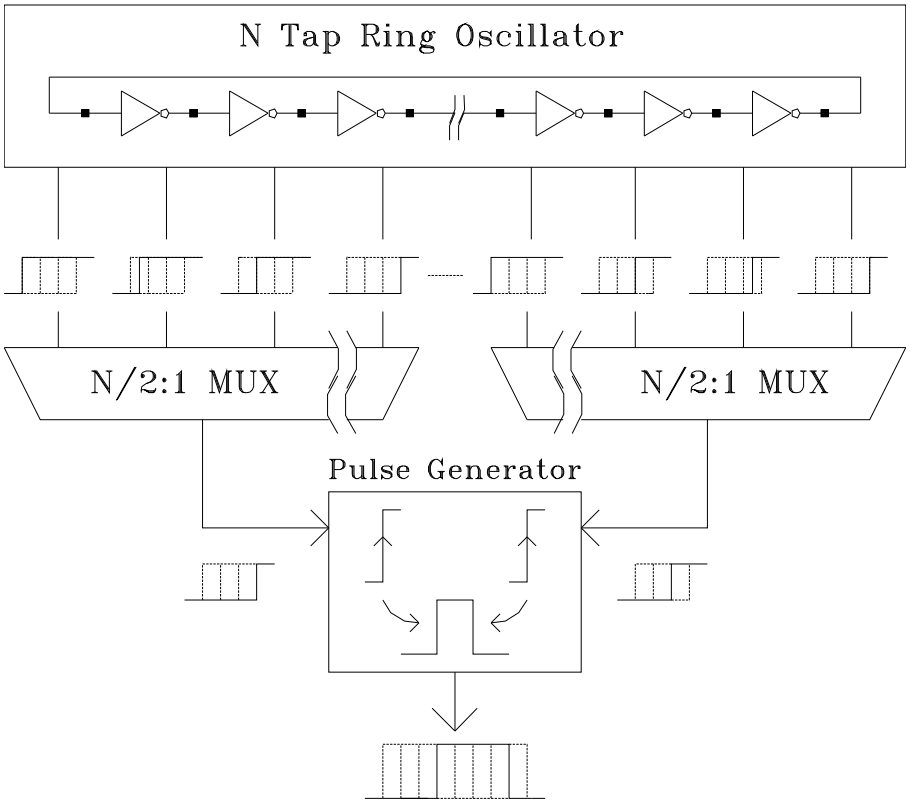

Figure 4.2: Block Diagram of the DPWM

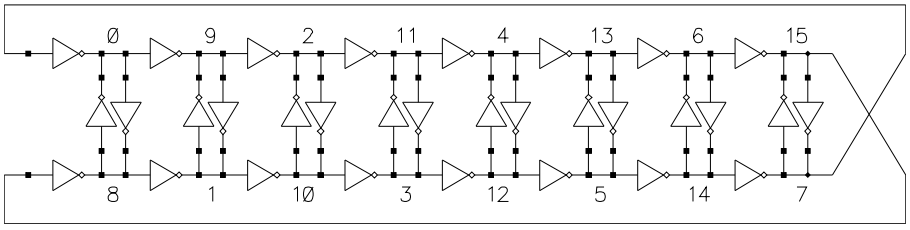

Figure 4.3: System Diagram of the Differential Ring Oscillator

an odd number of stages to oscillate.

Thus a differential ring oscillator configuration is used instead, as shown in Figure 4.3 and explained in [11]. A differential ring can be thought of as two separate rings, each representing a polarity, running in parallel with one another. Each polarity contains half the total number of stages required and the tap at the end of each ring is fed into the beginning of the opposite ring. In this way, oscillation can be maintained with an even number of total taps. The associated taps in the two polarities are then synchronized together with keeper inverters to enforce their differential nature. This scheme yields two benefits. First, a differential ring supports an even number of stages, which is necessary for doubled edged operation. In addition, an even number of taps allows for the full use of a binary encoded data bus and lends itself to multiphase operation. Each phase of an N bit M phase DPWM can be created by merely offsetting the taps $\frac{2^N}{M}$ stages away from each other. Another benefit is that a differential ring yields twice the number of taps for a given inverter propagation delay and clock frequency. Because each subring only has half the number of stages and the clock frequency of a ring is defined as $T_{RING} = 2 * T_{INVDELAY} * 2^N$, [12], the required inverter delay for a differential ring is twice that of a singled ended ring for a given clock frequency and number of taps. The extra margin in delay time allows for the use of larger sized inverters for matching and linearity purposes and for the use of lower supply voltages to reduce power consumption. One should note however that these benefits are partially offset by the additional logic and power consumption drawn by the keeper inverters. To minimize these drawbacks, the keeper inverters are commonly kept to a minimum size.

The oscillation frequency of the ring is controlled by varying the supply voltage to the inverters. The relationship between clock frequency and supply voltage is shown below:

$$F_{RING} = \frac{1}{2 * T_{INVDELAY} * 2^N} \tag{4.1}$$

$$T_{INVDELAY} = 0.69 * R_{EQ} * (C_{INTRINSIC} + C_{LOAD}) \tag{4.2}$$

$$R_{EQ} = \frac{3}{4}\frac{V_{DD}}{I_{DSAT}}(1 - \frac{7}{9}\lambda VDD) \tag{4.3}$$

$$I_{DSAT} = k * V_{DSAT} * (V_{DD} - V_{TH} - 0.5 * V_{DSAT}) \tag{4.4}$$

$$F_{RING} = \frac{k * V_{DSAT} * (V_{DD} - V_{TH} - 0.5 * V_{DSAT})}{2 * 2^N * 0.52 * C_{LOAD} * V_{DD} * (1 - \frac{7}{9}\lambda V_{VDD})} \tag{4.5}$$

The supply voltage versus ring oscillator frequency plot is shown in Figure 4.4. Because the available tuning range is not very wide, careful simulation and the adjustment of the ring inverter sizing should be used to target the desired frequencies of operation.

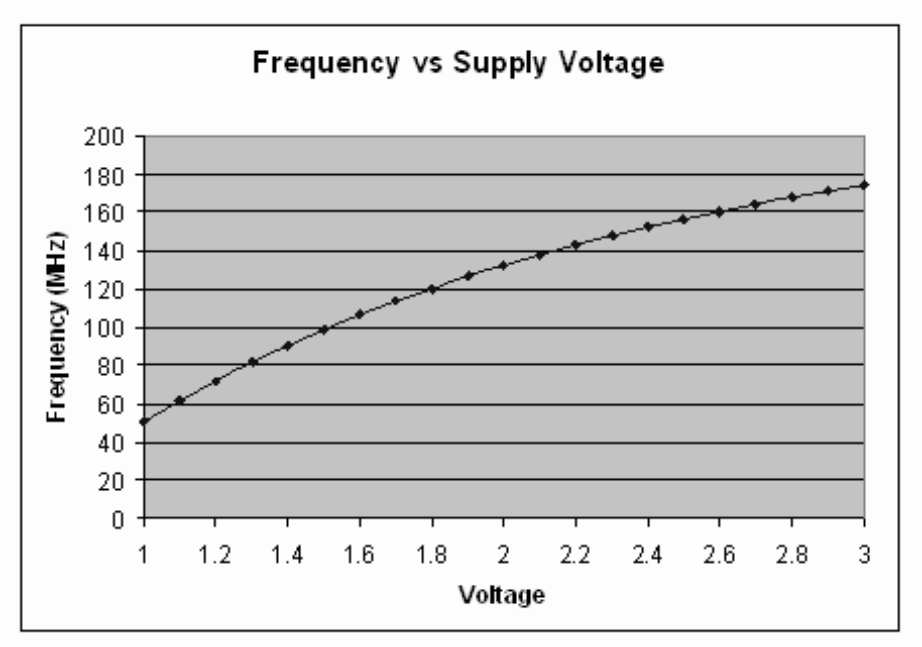

Figure 4.4: Voltage Supply vs Ring Frequency Relationship

One issue related with supply regulation is the interface between the varying and fixed supply logic. As noted above, the ring frequency is tuned by varying the supply to the ring inverters. However, the rest of the logic is supplied with a single voltage supply. Without any interfaces such as level shifters, the logic immediately following the supply-varying ring logic would experience power and performance penalties; this occurs because the supply-varying logic may not be able to drive the input gate of the fixed supply logic up to the constant supply voltage. Although level shifters could solve this problem with only a slight increase in design complexity, the supply voltage limits the tuning range to less than 1V before the transistors enter subthreshold. Another solution is to shift the range of the tuning voltage of the ring above the 1V standard supply. In this way, the ring inverters will never under drive the standard logic and the tuning range is now extended to over 1V without the risk of subthreshold operation. To handle the increased stress,

the ring inverters are implemented with thick oxide gates designed to operate at around 2.5V. To interface between the two supply domains, a thick oxide inverter biased with a standard 1V supply is used. As long as the supply tuning range is kept between 1V and 2.5V, the interface inverter will not experience any of the static power consumption problems mentioned above. Moreover, because the interface device is supplied with 1V, the interface output will not damage the standard logic. One drawback to this approach are the slower inverter delay times of the thick oxide inverters compared to the standard inverters. However, this degradation should not affect the functionality of the DPWM as long as the delay times are matched across all interface inverters. Moreover, the asymmetry between the rise and fall times at the output of the thick oxide interface inverter does not present an issue with the duty cycles because only the rising transitions are used to control the duty cycle of the output pulses. This configuration is illustrated in Figure 4.5.

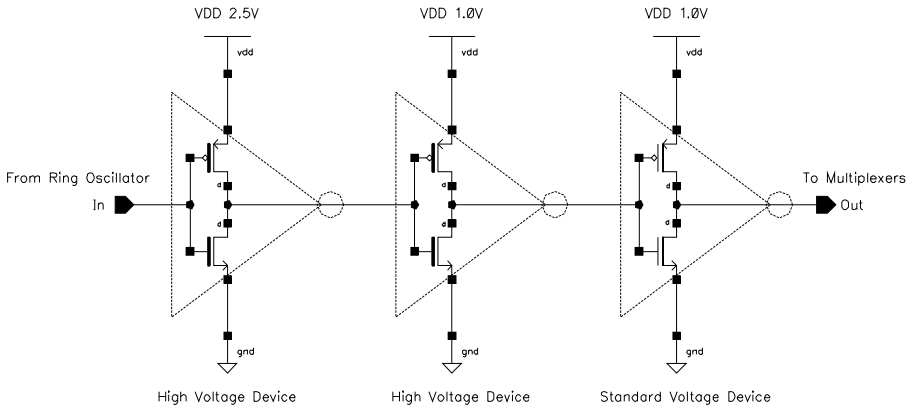

Figure 4.5: Implementation of a Supply Level Shifting Scheme

4.2.2 Multiplexer Bank

The multiplexer (MUX) banks are responsible for routing the correct tap to the pulse generator so that the rising and falling edge of the pulse is created at the desired time. To properly satisfy this responsibility, the MUX has two main objectives. First, the MUX must switch fast enough so that a rising command edge is not lost during a transition. Second, the delays from the inputs to the output must be uniform to ensure that the edges

are still uniformly spaced away from each other; deviations in delays will inevitably lead to a nonlinear DPWM output. This design contains four 5 bit multiplexer banks, one for each edge in a pulse for each phase.

Each five bit MUX bank consists of transmission gates arranged in a binary tree fashion as shown in Figure 4.6. Regeneration inverters are placed at each branch to minimize propagation delay due to switch resistance and capacitance along the signal path. The binary tree format has two main benefits. First, a binary encoded data bus can be fed directly to the transmission gates in the MUX bank, no decoding or other logic processing is necessary. Each level in the binary tree corresponds to one bit in the 5 bit select bus. Second, the regular nature of a binary tree assists in the matching of delays between all inputs of the MUX bank.

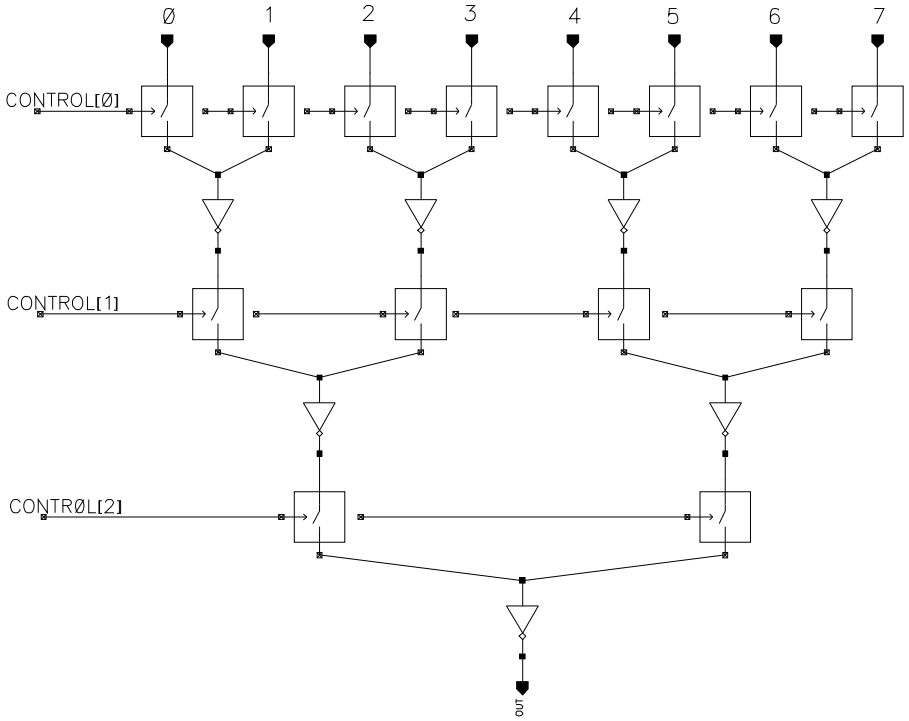

Figure 4.6: System Diagram of a 3-bit Multiplexer Bank

In this design, the timing relationship between data arrival and MUX switching is critical. If the MUX switches too early or late, glitches may be introduced which will cause multiple rising or falling command edges to be sent to the pulse generator. To ensure proper

synchronization, one tap in the ring is selected as the reference tap for each phase, as noted as tap 0 in Figure 4.3. This tap serves to mark the beginning of each new switching cycle and is used to clock the registers that lock the input data for the multiplexers. Because the MUX switches at tap 0, this implies that the MUX must be able to switch taps within one inverter delay; otherwise, a rising edge from tap 1 may be delayed. A timing diagram illustrating the tap selection scheme is shown in Figure 4.7. In Figure 4.7, rising edges are sequentially delayed from one another by one inverter delay. The rising edge in tap 0 indicates the beginning of a new PWM cycle. The dotted box indicates the maximum time period in which the MUX can switch tap selection without delaying an edge or generating a glitch. Note that the actual available time may be less due to the non-zero rise times of the rising edges. The vertical dotted line denotes the center of the PWM pulse: edges rising to the left of the line are used to create the rising edge, while edges to the right form the falling edge of the PWM pulse.

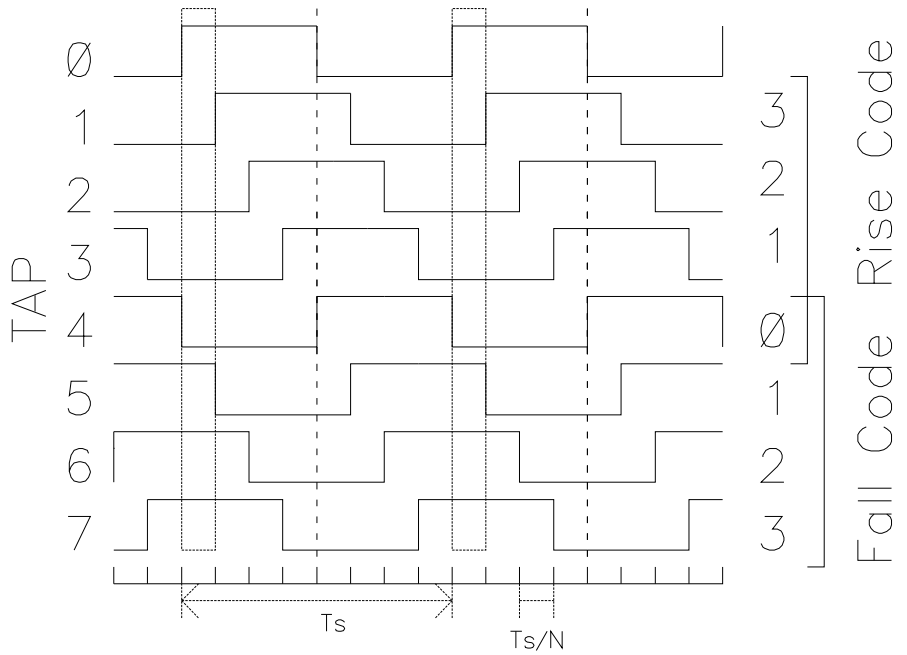

Figure 4.7: Timing diagram illustrating the Outputs of an 8 Tap Ring Oscillator.

An alternative to this design which is less timing sensitive is to use multiple multiplexers to select a rising edge as shown in [11]. In this method, one MUX bank is used

to select the current rising edge tap while the other MUX bank selects the tap for the next cycle. A third MUX can then be used to alternate between the two MUX outputs. In this way, the MUX banks can have up to one clock cycle for the transmission gates to switch and for the internal nodes to settle to the proper voltage.

4.2.3 Pulse Generation

The pulse generation block creates the desired output pulses based on when the rising and falling command edges are received. The main components of this block consist of two edge-to-pulse converters connected to the inputs of an SR latch, as shown in Figure 4.8. When a command edge is received, an edge to pulse converter creates a brief pulse at the set or reset inputs of the SR latch, depending on the type of edge received. Based on the input levels, the latch then generates a high or low signal at its output. The edge to pulse converter, in Figure 4.9a, consisting of an inverter based delay line and an AND gate, is necessary to prevent the set and reset inputs to the latch from being asserted at the same time. The timing of the delay line is set such that the pulses from the converter last no longer than the width of the minimum duty cycle, which corresponds to minimum time between the arrival of two command edges. An SR latch was chosen to generate the pulses due to its robust nature. Glitches such as multiple or missing set and reset signals will only render the pulse at the current and possibly next clock cycle incorrect; when normal input are resumed, the SR latch will operate properly as well without the use of an external reset signal. Although the SR latch is capable of generating complementary outputs, only a single output is used by the powertrain. To ensure that the rise and fall times are symmetrical to minimuze duty cycle errors, the unused complementary output is loaded with an dummy inverter. Simulations indicate that rise and fall times differ by only 1 ps or less than 1% of the minimum pulse width at the output of the SR latch.

In the event that no pulse is required for one clock cycle, the MUX bank and taps are arranged such that the rising and falling edges arrive at the same time to create a zero width pulse. However, this is an unacceptable input to the SR latch as this will inherently cause the set and reset inputs to the latch to be asserted simultaneously. To prevent this from happening, a suppression circuit, shown in Figure 4.9b, is added to detect

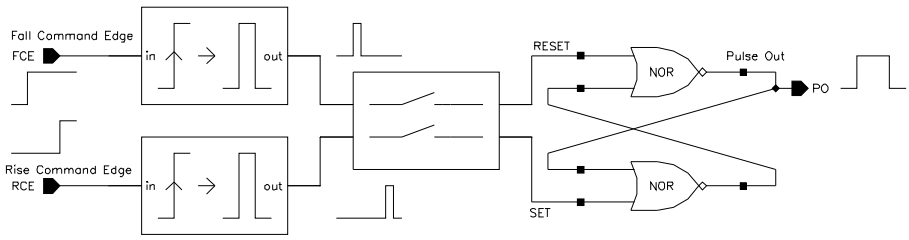

Figure 4.8: System Diagram of the Pulse Generator

and suppress any simultaneously generated set and reset pulses to the SR latch to prevent latch metastability.

4.3 Powertrain

The powertrain is comprised of a half-bridge network and an LC filter. A half-bridge network, shown in Figure 4.1, consists of a high and low side switch connected to positive and ground rail, respectively. These switches turn on and off to create a square wave of a given duty cycle. The output square wave is then fed into an LC tank that filters out the high frequency components, leaving only the desired DC term behind. A careful analysis of class D supply operation is performed in [1]. In an ideal switched mode power supply, the efficiency can theoretically approach 100% and the output voltage follows the relationship: $V_o = V_{IN} * D$, where D denotes the duty cycle of the square wave. In a dynamic class D PA design, the duty cycle is continuously varied so that the output voltage tracks a desired signal. In this situation, the LC tanks must be carefully designed to suppress the switching ripple without excessively attenuating the desired frequency content.

Because the design aims for high efficiency operation, the powertrain is optimized to minimize loss. The source of loss in this powertrain is categorized into two main components: conduction and dynamic losses. In this application, where the load is modeled as a resistor, conduction losses scale proportionally to duty cycle and consist primarily of losses due to switch resistance and inductor resistance. Dynamic losses scale proportionally to the switching frequency, which is usually fixed, and are due to energy wasted driving gate capacitance, hard switching losses when charging the output capacitance, and power dissi-

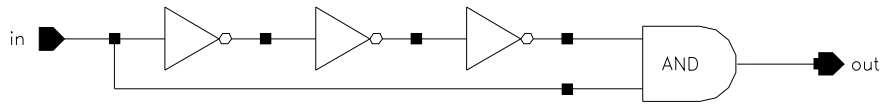

(a) Edge to Pulse Generator

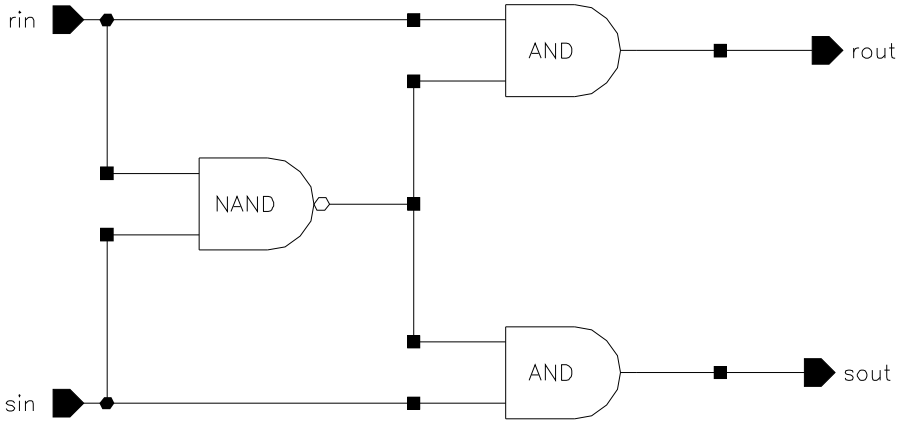

(b) Pulse Suppression Circuit

Figure 4.9: Implementation of Pulse Generator Sub-blocks

pated across the body diode during dead time switching. To minimize loss, the high side and low side switches are scaled such that the sum of losses due to conduction and dynamic sources is minimized. Because the optimum sizing varies with duty cycle, the powertrain must be optimized over the entire expected range of operation. As explained in [1], this process is accomplished by maximizing the following average efficiency equation:

$$\bar{\eta}_{supply} = \frac{E_{PA}}{E_{battery}} = \frac{\int_{-\infty}^{\infty} g_s(P_{PA}) \cdot P_{PA}\, dP_{PA}}{\int_{-\infty}^{\infty} g_s(P_{PA}) \cdot P_{battery}(P_{PA})\, dP_{PA}} \tag{4.6}$$

where P_{PA}, $g_s(P_{PA})$, and $P_{battery}(P_{PA})$, represents a given power output, the probability density function of P_{PA}, and the power consumption of the powertrain operating at P_{PA}, respectively. In this design the high and low switches are implemented on chip. They are designed and optimized by Stauth according to the process outlined in [1].

In addition to the main switches, the power train also includes dead time controls and buffer inverters. The adjustable dead time controller modifies the incoming control pulses to prevent the high and low side switches from activating simultaneously to eliminate

short circuit current which would lead to excessive power loss. In certain conditions, careful tuning of the dead time controller can also be used to accomplish zero voltage switching operation to further improve power efficiency. Finally, the buffer inverters are used to increase the driving strength of the signal so that the powertrain switches are switched with sufficiently fast rise and fall times.

4.4 Layout Considerations

In traditional switched mode power supplies, the switching rate is set significantly higher than the control bandwidth to allow for the use of feedback regulation to ensure that the output voltage level is correct. As a result, nonlinearities in the ring and any resulting distortion in the output will be attenuated. However, this design uses a feedforward control approach and, as a result, all nonlinearities not anticipated in the control block will be present at the output. Although ring nonlinearities can be measured and corrected to a degree in the digital feedforward control, such corrections may not perfect. As noted above, quantization levels and DSP complexities all limit the predistortion capabilities of digital control system; therefore, care should be taken to minimize the amount of correction required.

In the case of the DPWM circuit, layout considerations began at the lowest level, that of the inverter. The individual ring inverters were sized approximately four times the minimum size to improve matching against process variations that would affect the propagation delay between taps. Furthermore, dummy inverters were placed at the ends of the ring to minimize edge variation artifacts. The next modification rearranges the placement of the inverters as shown in Figure 4.10. Although connectivity remains the same, the new format has several advantages. First, this arrangement allows for a compact and regular layout. With all the inverters oriented in the same direction, the ring can be realized with two rows of sixty-four standard inverter cells. Second, the cross connections allow every tap to be available at both edges of the ring. Previously, extra wiring would be needed if a tap was needed at the opposite edge of the ring, resulting in varying wire capacitance at every tap. In this design, almost every inverter will see an identical load.

The shape of the ring can also be manipulated to improve symmetry. The ring is bent into a U shape to minimize the delay associated with returning the signal back to the beginning of the ring. The rearrangement mentioned above works hand in hand with the U shape by making the ring longer but thinner; although length does not affect symmetry in a U shape, the thickness affects the amount of extra metal needed at the bends of the U. Delay variations can be further minimized by increasing the number of bends in the ring. However, this greatly increases the complexity involved to wire out the ring to the multiplexers.

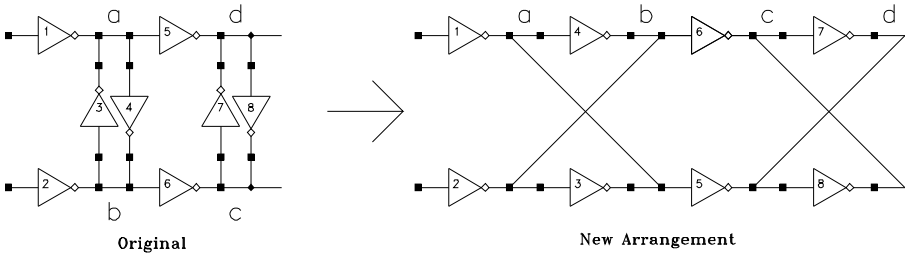

Figure 4.10: The rearranged inverter configuration in the differential ring

The last consideration to ensure symmetry is to ensure that all the traces leading from the ring taps to the MUX banks to the pulse generators are of equal lengths. Equal trace lengths will, to the first order, equate to identical delays. Ideally, the traces used to generate both rising and falling edges of both phases would be the same. However, this is difficult to do due to limitations in space and metal layers. Consequently, only traces used to create a command edge (either rising or falling) in a single phase were equalized. As shown in Figure 4.11, each edge of the ring contains half the total number of taps necessary to generate an edge of a pulse. Each 32-1 MUX was broken in half at the MSB tap and each half was placed on either edge of the ring. This enables each tap to reach its respective MUX without having to cross across the ring. Moreover, the inputs to the MUX were made such that they lined up exactly with the taps on the ring to minimize the possibility that the outer taps experienced more delay than the inner ones. Within each MUX, the delay of every input can be easily kept constant due to the inherent nature of a binary tree MUX. The output of each half of the MUX are then joined together and fed into the pulse generator.

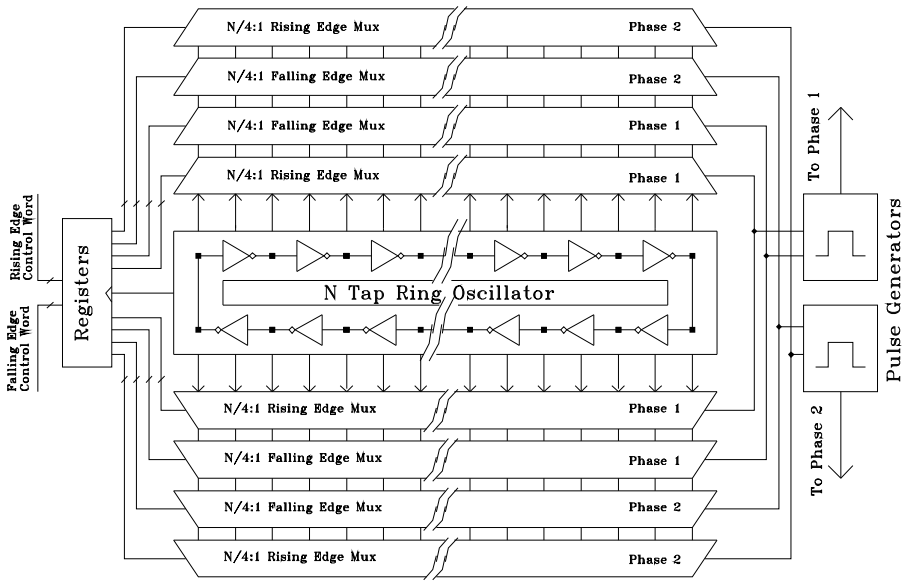

Figure 4.11: Layout diagram optimized for equal delays between taps

4.5 Experimental Results

The following section describes the preliminary experimental results of the chip. The first test measures the linearity of an unloaded chip. Input codes are loaded into the chip and the resulting output pulse widths are measured at steady state. Figure 4.12 and Figure 4.13 displays the output of the test for Phase 1 and Phase 2 outputs, respectively. From the figures, both Phase 1 and Phase 2 output are rather linear in the lower codes; however, Phase 1 displays several non-linear deviations at the higher codes.

In the second test, input codes are again loaded into the test chip at steady state; however, in this case, the chip is loaded with an LCR tank and the output DC voltage is measured. The LCR load consists of a 82nH inductor, 6.8nF capacitor, and 5 ohm resistive load. With a 1V supply, the maximum DC draw is 200mA. Again, as shown in Figure 4.14 and Figure 4.15, the DC voltage increases linearly with increasing input codes in the lower codes. However, at the higher codes, the voltage displays a kink, with Phase 1 kinking upwards and Phase 2 kinking downwards. One possible cause is that the analog and digital VDD are shared. At low loads, the internal supply is stable and the digital logic performs as expected. However, as the load increases, the voltage ripples created by the powertrain

switches may adversely affect the DPWM. For example, as can be seen in Figure 4.11, in Phase 1 the rising multiplexers are closer to the ring than the falling multiplexers; however, this is reversed for Phase 2. At low codes, the ring buffers can sufficiently drive the data line connecting the multiplexers together. However, with an unstable supply, the buffer drive may be insufficient and there may be appreciable delays as the signal travels down the data line. In this case, because the data arrives earlier at the rising MUX for Phase 1 and later at the rising MUX for Phase 2, the resulting output pulses and voltages will be larger and smaller respectively. In dual phase mode, this behavior does not appear and the output increases quite linearly with an increasing input. One explanation is that the two outputs cancel out each other's non-linearities. Another explanation is that dual phase operation imposes a more constant current draw on the supply, leading to a more stable internal supply voltage, which leads to increased linearity in the digital DPWM logic. During experimentation, it was observed that voltage ripple on the supply pins decreased dramatically when the chip operated in a dual phase configuration.

In the third test, the code switching capabilities of the test chip are measured. The input codes are swept from 0 to 31 to 0 at a rate of 100 MHz and the resulting output pulse codes are measured. Ideally, a monotonic triangle waveform would be generated; bumps in the plot indicates that either the registers failed to latch onto the correct code or the MUX banks failed to switch on time. As shown in Figure 4.17, the output does generally track the input, but small variations occur, suggesting that although the DPWM is able to track the input codes, it may be having trouble reading the fastest changing input LSB bit.

Figure 4.18 and Figure 4.19 plots out several representative output pulses for small and large input codes, respectively. In the plots, the pulses are rather stable and do not display any appreciable droops, suggesting that the power switches are sufficiently sized to handle the current ripple. However, a large amount of ringing occurs at the transitions of each pulse indicating that there may still be an excessive amount of inductive parasitics in the connection between the die and the board. Further tests could not be carried out due to ESD issues leading to chip damage before more thorough tests could be performed.

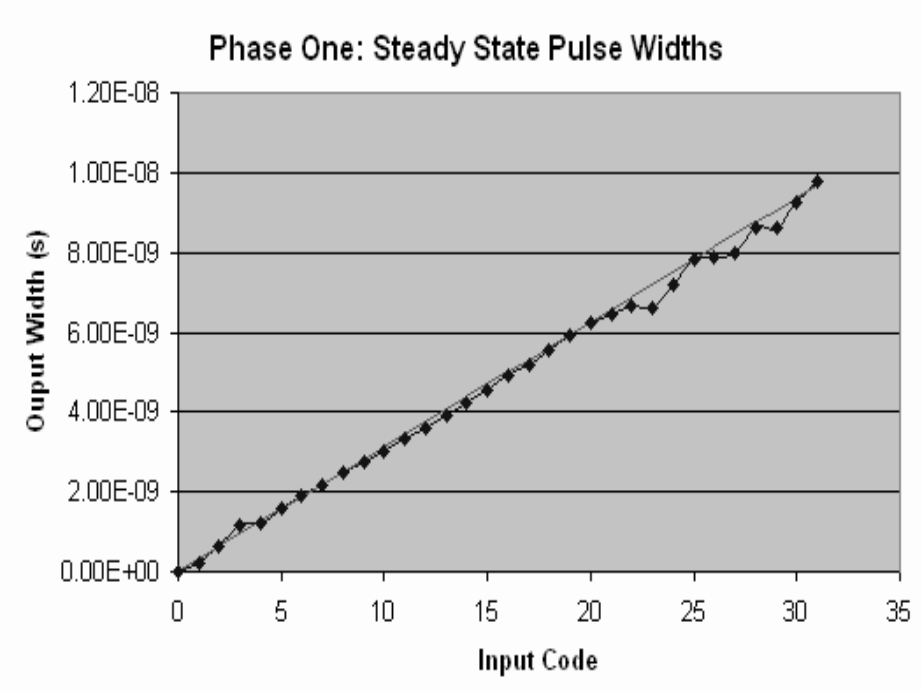

Figure 4.12: Phase One: Steady State Pulse Widths

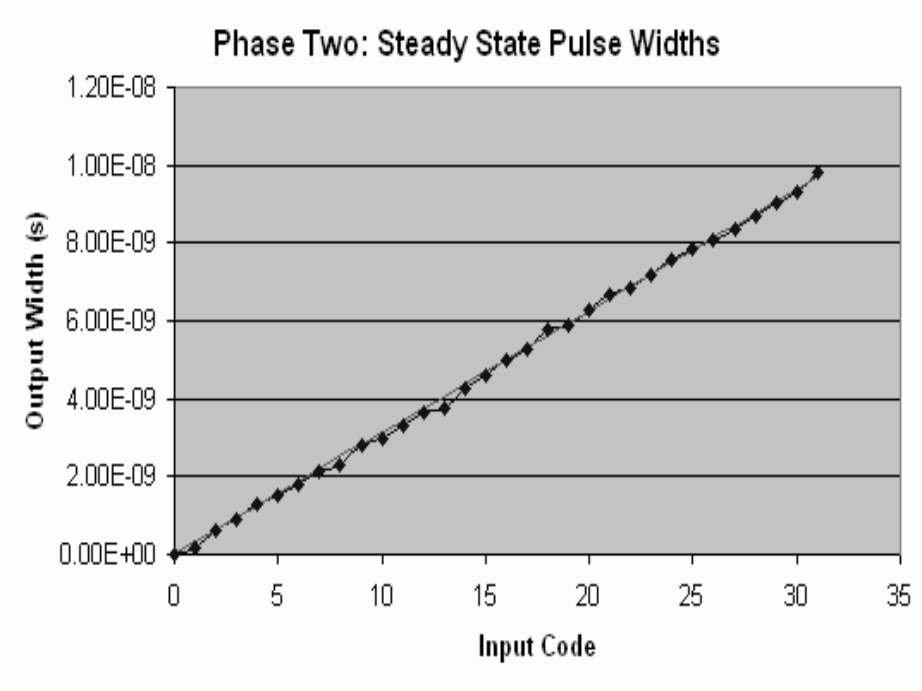

Figure 4.13: Phase Two: Steady State Pulse Widths

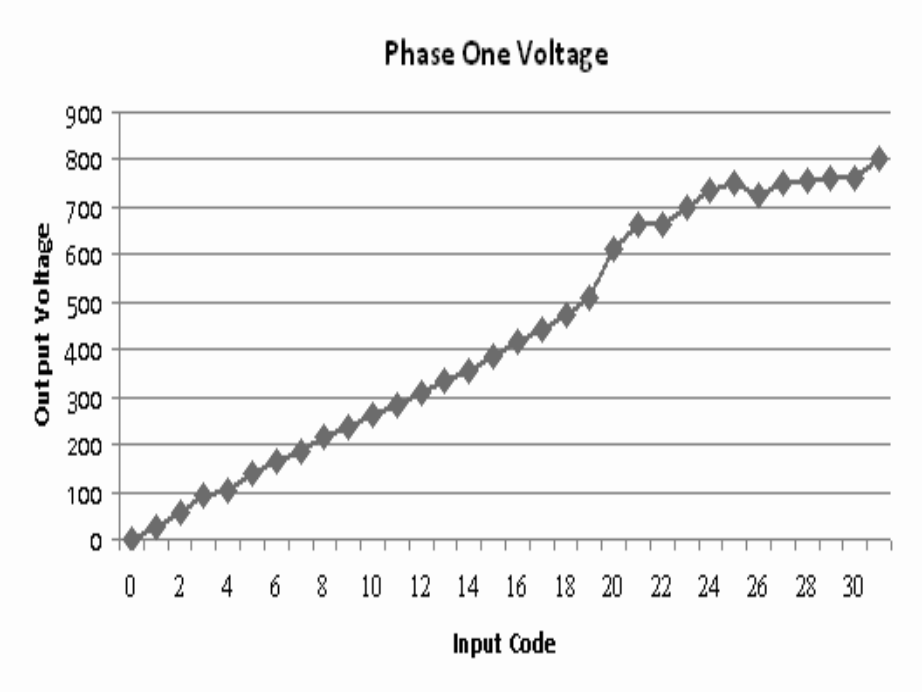

Figure 4.14: Phase One: Steady State Output Voltage

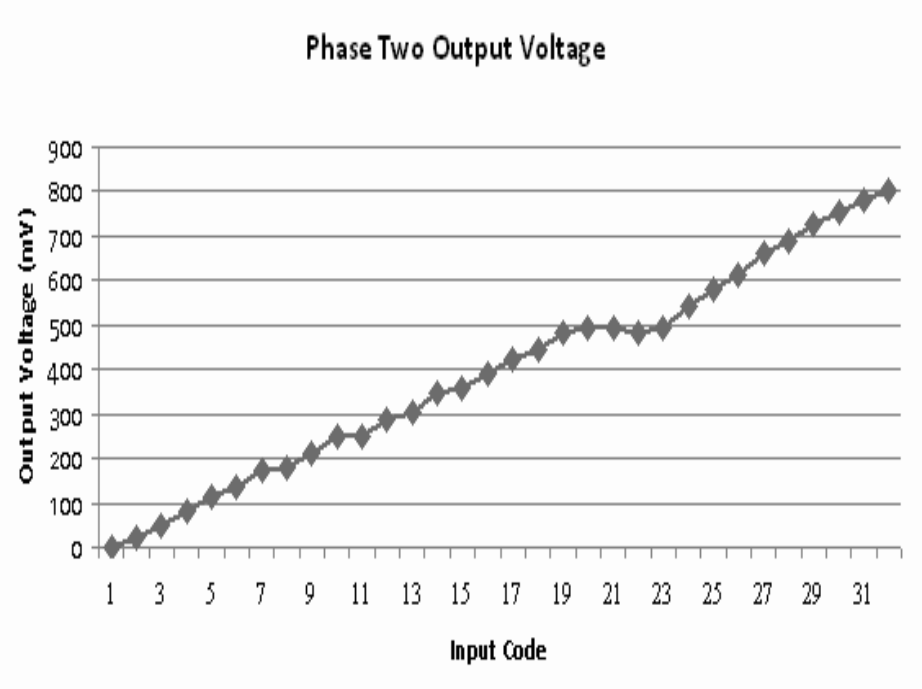

Figure 4.15: Phase Two: Steady State Output Voltage

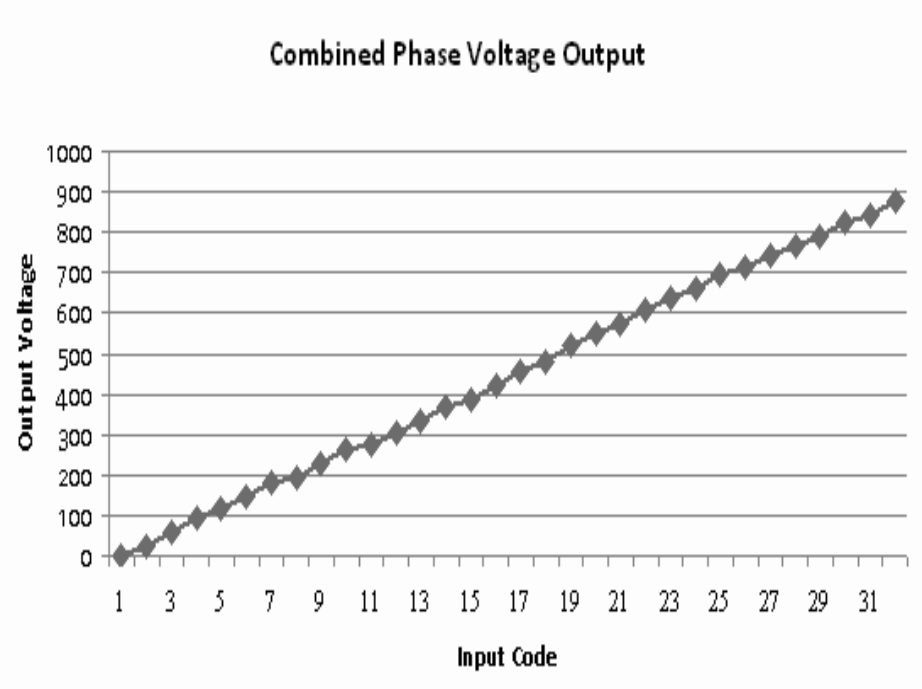

Figure 4.16: Dual Phase Steady State Output Voltage

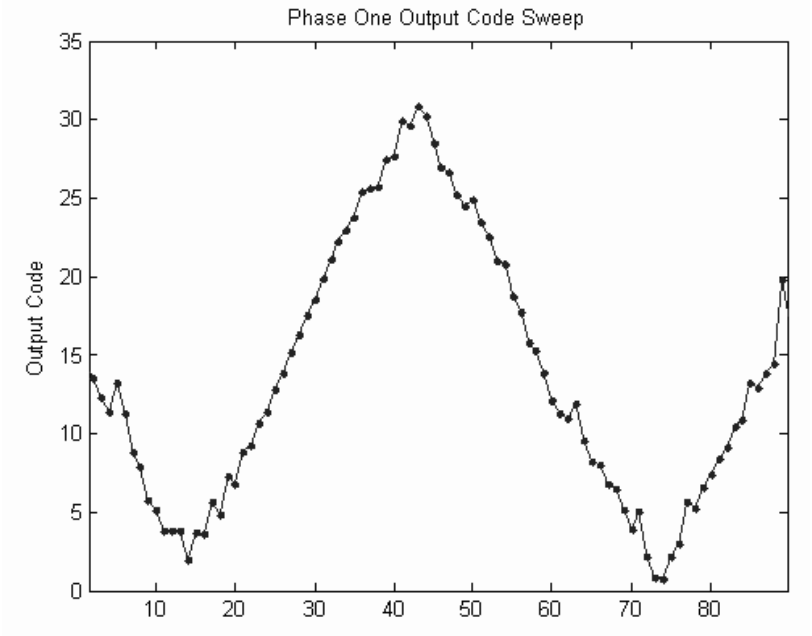

Figure 4.17: Phase One Output Pulse Width Sweep

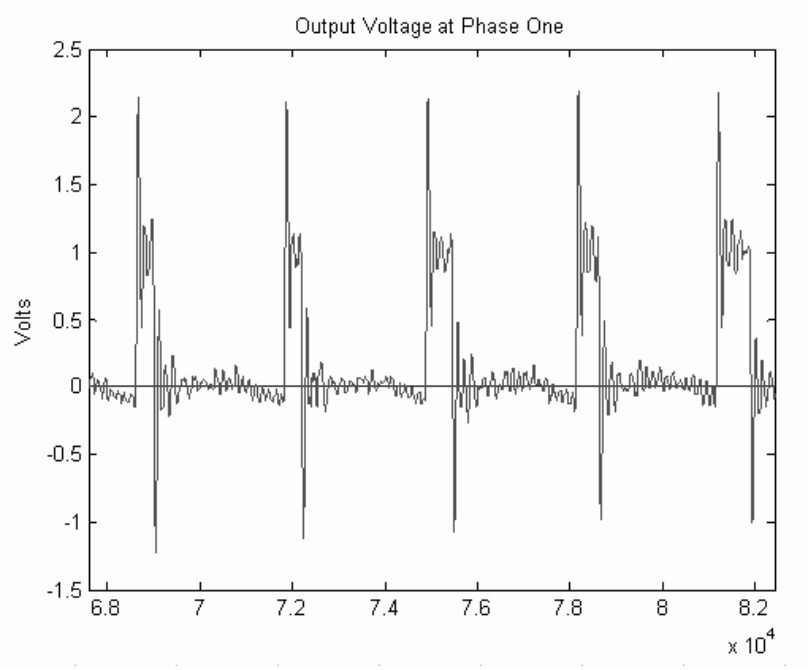

Figure 4.18: Phase One: Small Code Output Pulses

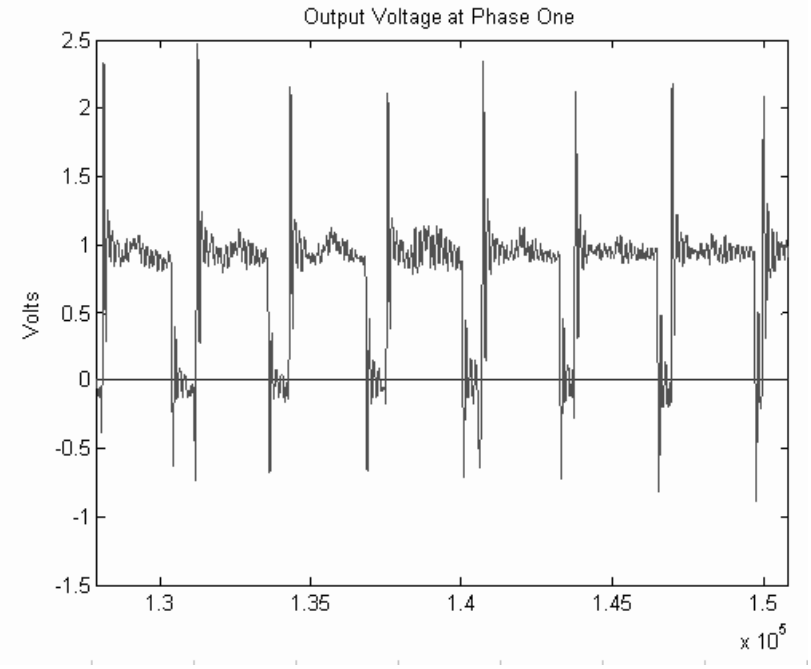

Figure 4.19: Phase One: Large Code Output Pulses

4.6 Analysis and Improvements

One observation in the experimental results was that the output pulse widths did not always match the desired input code. One possible reason is because the DPWM clock and the FPGA clock are not synchronized together. Because the phase between the data and the sampling clock are not locked, there will be no guarantee that the input data will have settled to the correct value. However, this situation will normally not be a problem because the majority of the 802.11a signal power is located at relatively low frequencies, as shown in the power spectral density (PSD) of the envelope signal in Figure 4.20. At these low frequencies, the input codes will appear as a DC signal relative to the clock and data settling will not be an issue. However, to ensure accuracy at all times especially during high frequency transitions, data locking should be used to ensure that the data is ready when sampled by the DPWM registers.

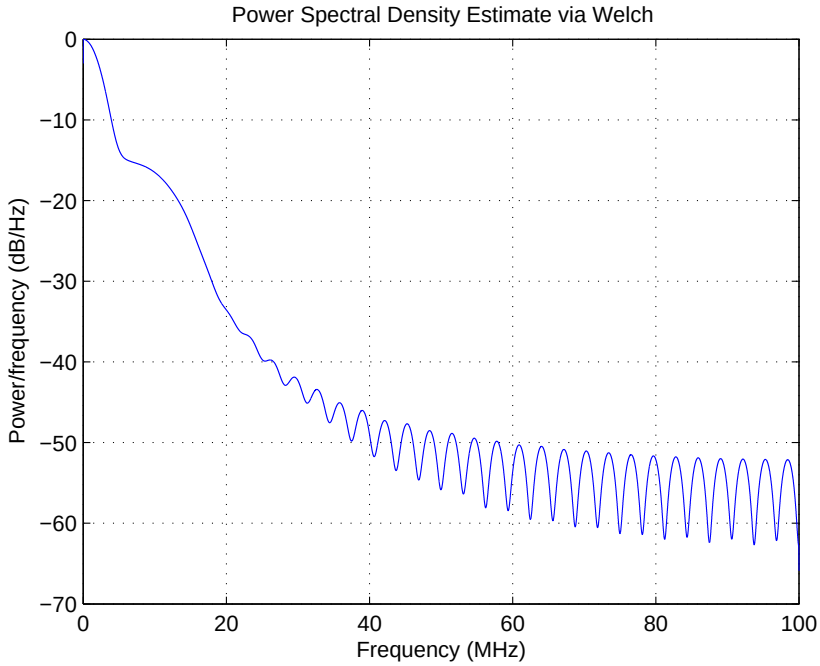

Figure 4.20: Power Spectral Density of a sample 802.11a waveform

Several asynchronous clocking methods are available to accomplish this task. The most common approach is to use a single clock for both the FPGA and DPWM: the data

sampling clock in the DPWM can be forwarded to the FPGA to synchronize the two components together. Any clock skew due to the large routing distance necessary between the DPWM chip and the FPGA board may be corrected via clock retiming schemes. Another possible approach is to synchronize the DPWM to the FPGA by forwarding the FPGA clock along with the control data. A bank of registers located on the DPWM chip then latches the control data on the negative clock edge of the FPGA. In this manner, the control data will always be sampled after a sufficient amount of settling time has passed. This latched data is then resampled at the DPWM clock rate by another bank of registers. The latter solution may be useful in situations where the FPGA clock cannot be externally controlled or synchronized.

Another area for improvement is the layout of the on chip metal traces. Because of the large current draws created by the switching regulator, parasitic inductances and resistances on the test board and on the package will invariably create voltage swings in the on chip supply. These current draws occur primarily when the switches initially close and charge the output capacitance as well as when the current ripples across the inductor. The resulting voltage swings are highly undesirable because they will degrade chip performance. First, because the power and digital supplies are combined together, voltage swing may lead to unpredictable digital behavior. Voltage swings in the digital supply may lead to varying propagation delays through the logic gates which may adversely affect the linearity of the circuit. The solution to this issue is simply resolved by adding an isolated supply rail and ground rail for the digital logic.

Voltage swings due to parasitic inductances will also reduce the efficiency of the powertrain. During high current swings, the inductors will absorb energy that would ideally be delivered to the load; hence reducing power efficiency. To counteract these effects, bypass capacitances are used recover energy from the inductors and release it back into the load when required. To minimize the effects of board level parasitic inductances due to traces, external bypass capacitors are used. Ideally, the external bypass capacitor should be placed next to the package pins so that any voltage ripple generated by parasitic trace inductances may be shunted to ground. In the existing solution, however, the VDD and ground pins are situated away from each other on opposite sides of a chip. This configuration makes it

difficult to place a bypass capacitor close to the VDD and ground pins to provide adequate ripple suppression. In future designs, the metal traces inside the chip should the rearranged so that the VDD and ground pins may be placed next to each other to eliminate this problem.

Within the chip itself, bypass capacitors are used to suppress the inductance due to the bond wires and the package. The existing design contains approximately 350pF of on chip N-well capacitors, used primarily for its high capacitance density. However, N-well capacitors contain non trivial amounts of ESR which reduces its bypassing capabilities at high frequencies. To compensate for this, metal oxide metal capacitors (MOM), which have minimal ESR, can be added in addition to the N-well capacitors; however, they were not initially used due to their substantially lower capacitance density. In the current design, a large portion of the die is unused and is occupied with a metal grid used to transfer current from the supply pads to the powertrain switches. To improve ripple rejection, the metal grid can be replaced with a large array of multilevel MOM caps. In this way, the low ESR MOM capacitors will bypass the high frequency transients, while the N well capacitors will handle the larger slower transients. Moreover, because the MOM array consists of metal layers, it can still perform the function of current transfer without substantial resistive losses.

Chapter 5

Conclusion

In conclusion, this report discussed the implementation of a feedforward envelope tracking power supply. First, an overview of the feedforward approach was discussed, including the various sources that require compensation in the supply path, such as the LC tank and the varying PA load. Equations to model and invert these sources were discussed and implemented using Simulink and Xilinx digital blocks. Considerable effort was also made to optimize these algorithms to operate efficiently. Second, digital processes to improve the output quality, through the use of a natural sampling and a sigma delta converter, were also discussed.

Third, the implementation of the 5 bit 100MHz DPWM and buck converter was discussed. Because the control system does not contain feedback, a considerable effort was made to design and layout for linearity. Although the chip could not be fully tested due to ESD and data synchronization issues, initial results were positive: running in a two phase mode, the supply output was relatively linear and was able to source approximately 200mA of current to a 5 ohm load. However, further improvements can still be made in respect to the data interface between the digital control blocks and the DPWM. Moreover, powertrain performance can be increased through improvements in the supply rail design.

Bibliography

[1] J. Stauth, "Dynamic power supply design for high-efficiency wireless transmitters," Master's thesis, University of California at Berkeley, 2006.

[2] A. A. Saleh, D.C. Cox, "Improving the power added efficiency of fet amplifiers operating with varying-envelope signals," *IEEE Transactions on Microwave Theory and Techniques*, vol. 51, pp. 51–56, 1983.

[3] F. Wang, A. H. Yang, D. Kimball, L. E. Larson, P. Asbeck, "Design of wideband-width envelope tracking power amplifiers for ofdm applications," *IEEE Transactions on Microwave Theory and Techniques*, vol. 53, pp. 1244–1255, 2005.

[4] B. Sahu,G. A. Rincon-Mora, "A high efficiency linear rf power amplifier with a power tracking dynamically adaptive buck-boost supply," *IEEE Transactions on Microwave Theory and Techniques*, vol. 52, 2004.

[5] N. Wang, V. Yousefzadeh, D. Maksimovic, S. Pajic, and Z. B. Popovic, "60% efficient 10-ghz power amplifier with dynamic drain bias control,," *IEEE Transactions on Microwave Theory and Techniques*, vol. 52, 2004.

[6] C. Pascual, Z. Song, P. T. Krein, et. al, "High-fidelity pwm inverter for digital audio amplification: Spectral analysis, real-time dsp implementatoin, and results," *IEEE Transactions on Power Electronics*, pp. 473–485.

[7] Z. Lukic, K. Wang, A. Prodic, "High-frequency digital controller for dc-dc converters based on multi-bit l a pulse-width modulation," *Applied Power Electronics Conference and Exposition*, 2005.

[8] J. Stauth, "Personal communications," 2006-2007.

[9] A. V. Oppenheim, R. W. Schafer, *Discrete-Time Signal Processing.* Prentice Hall, 2 ed., 1999.

[10] D. Johns, K. Martin, *Analog Integrated Circuit Design*, pp. 546–547. Wiley, 1 ed., 1996.

[11] J. Xiao, *An Ultra-Low-Quiescent-Current Dual-Mode Digitally-Controlled Buck Converter IC for Cellular Phone Applications.* PhD thesis, University of California at Berkeley, 2003.

[12] J. Rabaey, A. Chandrakasan, B. Nikolic, *Digital Integrated Circuits*, pp. 200–202. Prentice Hall, 2 ed., 2002.