Title
Digitally Enhanced Wireless Transceivers for Multi-mode Reconfigurable Radios

Author
Nanda, Rashmi

Publication Date
2012

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

Digitally Enhanced Wireless Transceivers for Multi-mode Reconfigurable Radios

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Electrical Engineering

by

Rashmi Nanda

2012

ABSTRACT OF THE DISSERTATION

Digitally Enhanced Wireless Transceivers for Multi-mode Reconfigurable Radios

by

Rashmi Nanda

Doctor of Philosophy in Electrical Engineering

University of California, Los Angeles, 2012

Professor Dejan Marković, Chair

Digitally enhanced wireless transceivers are gaining prominence due to their promise of greater integration, flexibility to adapt to varying SNR conditions, performance, and area benefit that comes with CMOS feature size reduction. This thesis discusses the benefits of going digital in a system that has until now been dominated by analog or discrete-time signal processing. Digitally assisted radio transceiver architectures that incorporate a wide range of tuning parameters to control the RF carrier, signal bandwidth, and baseband modulation schemes are presented. These architectures rely on high-speed A/D and D/A conversion close to the antenna in order to maximize the extent of digital signal processing in the radio chain. We examine specific challenges that such a system faces in the context of minimizing noise, maximizing signal bandwidths, and enabling efficient filtering and sample rate conversion. The thesis also presents prototype architectures for a radio receiver and transmitter that employ a range of signal processing techniques to enable high-throughput, power-efficient systems.

The dissertation of Rashmi Nanda is approved.

Alan N. Willson

Sudhakar Pamarti

Milos D. Ercegovac

Dejan Marković, Committee Chair

University of California, Los Angeles

2012

I dedicate this dissertation to my husband, Rohit, for his incredible patience and unwavering love and support over the course of my research.

TABLE OF CONTENTS

List of Figures

LIST OF TABLES

ACKNOWLEDGMENTS

This thesis would not have been possible without my advisor, Dejan Marković's, constant support and mentoring. He helped me throughout my time in graduate school and gave me sufficient freedom to work on things that would hold my interest. I am especially grateful for his patience during times when the experimental part of the research was not going ahead as planned. I am also thankful for his guidance and timely discussions, academic and otherwise, which helped me achieve my goals at UCLA. I would also like to thank my committee members, Professor Willson, Professor Ercegovac, and Professor Pamarti for their useful comments during the qualifying examination, which helped define the scope of this work. I also appreciate the funding support provided by the Center for Circuits and Systems Solutions (C2S2) during the course of my graduate study. Also, I would like to thank ST Microelectronics and Samsung Corporation for the infrastructural support they provided.

I want to thank my friends at UCLA for the help they provided during the 6 years of my stay. I appreciate Nitesh Singhal's help during research discussions and his guidance when I was setting up my experiments in the circuits lab. I would also like to thank Abhishek Ghosh who was ready to engage in research discussions at all times. In addition, all the group members in DMgroup were very supportive when I needed their assistance. Cheng-Cheng Wang was extremely helpful during the final phase of chip testing and I will always be grateful for his objective approach when it came to debugging problems in the test setups. Henry Chen taught me the use of all the relevant equipments in the lab and most of the experimental results in this thesis were obtained with his help and knowledge. I am also grateful to Sarah Gibson, Vaibhav Karkare, Chia-Hsiang Yang, Victoria Wang, Richard Dorrance, Vahagn Hokhikyan, Hariprasad Chandrakumar and other members of DMgroup for providing useful insights during the group meetings.

And finally, I have to thank my husband Rohit who was unfailing in his support and encouragement to keep me motivated so I could get closer to completing my degree. I would also like to thank my parents for the support and encouragement they provided, especially during times when I needed it most.

Portions of Chapters III and VI have been reprinted, with permission, from: R. Nanda, H. Chen,

and D. Marković, “A Low-Power Digital Front-end Direct-sampling Receiver for Flexible Radios,” *IEEE Asian Solid State Circuits Conference*, November. 2011. Copyright © 2011, IEEE. D. Marković was the PI.

Portions of Chapter III have been reprinted, with permission, from: R. Nanda and D. Marković, "Digitally Intensive Receiver Design: Opportunities and Challenges", *IEEE Design & Test of Computers*, November 2012. Copyright © 2012, IEEE. D. Marković was the PI.

VITA

2002–2006	Bachelor of Technology (B.Tech.), Electronics and Electrical Communication Engineering, Indian Institute of Technology, Kharagpur, India.
2006–2008	Master of Science (M.S.), Electrical Engineering, University of California Los Angeles (UCLA), California.
09/08–12/08	Teaching Assistant, Electrical Engineering Department, UCLA, Taught sections of EE 10 (circuit analysis class) under direction of Professor Babak Daneshrad.
04/12-06/12	Teaching Assistant, Electrical Engineering Department, UCLA. MS Online EE 216B (advanced graduate class on DSP architecture design and optimization) under direction of Professor Dejan Marković.
09/06–09/12	Research Assistant, Electrical Engineering Department, UCLA.

PUBLICATIONS

R. Nanda and D. Marković, "ΣΔ Modulators for Low Power Digitally Intensive Radio Transmitters", to appear in the Proceedings of 2012 Asilomar Conference on Signals, Systems, and Computers.

R. Nanda and D. Marković, "Digitally Intensive Receiver Design: Opportunities and Challenges", IEEE Design & Test of Computers, (Special issue on digitally-assisted wireless systems) November 2012.

R. Nanda, H. Chen, and D. Marković, "A Low-Power Digital Front-end Direct-sampling Receiver for Flexible Radios," IEEE Asian Solid State Circuits Conference (ASSCC), November 2011.

R. Nanda and D. Marković, "Energy-efficient Retiming and Scheduling of Datapath Dominant Digital Systems," ASP J. Low Power Electronics, vol. 7, no. 3, pp. 341-349, Aug. 2011.

Rashmi Nanda, Chia-Hsiang Yang, Dejan Marković, "DSP Architecture Optimization in Matlab/Simulink Environment," in Proc. Int. Symposium on VLSI Circuits (VLSI'08), June 2008, pp. 192-193.

CHAPTER 1

Introduction

The evolution of CMOS integration is erasing the rigid boundary that used to divide a communication system into two distinct parts, the digital processing core and the RF front-end. While the digital core was tracking Moore's law and benefiting from the speed and performance enhancement provided by CMOS scaling, the RF front-end struggled with issues of low-voltage headroom, poor linearity, and small dynamic range. Adding to these problems was the need to support a wide variety of standards in a single device. Emerging standards such as LTE and WiMAX support data bandwidth up to 20MHz across multiple frequency bands that span several hundred MHz of the RF spectrum. A summary of the RF carriers and signal bandwidths for radio receivers in various wireless standards is shown in Figure 1.1. The multi-mode and multi-band specifications call for RF carriers anywhere between 300MHz to 5.8GHz and signal bandwidths between 200kHz to 20MHz.

The obvious solution was to replicate a conventional front-end multiple times, each realization

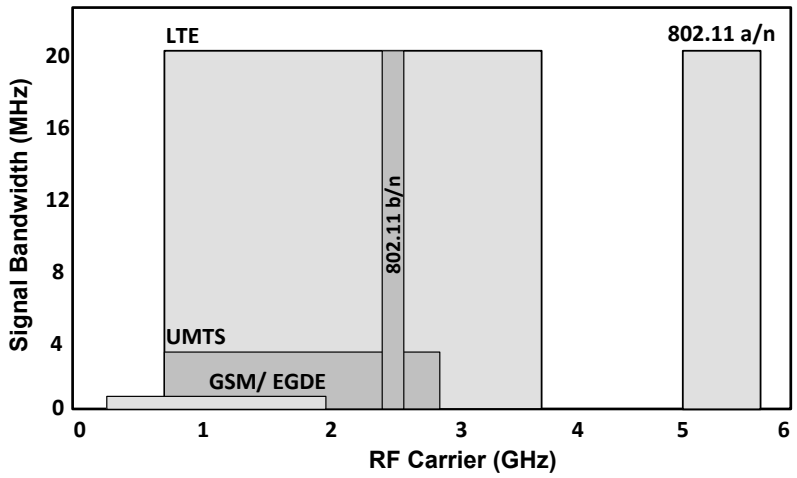

Figure 1.1: Multi-mode, multi-band specifications for wireless standards.

optimized to function in a narrow frequency band. For example, the conventional radio receiver, shown in Figure 1.2(a), will need several realizations of components like LNAs, analog amplifiers, mixers, and lowpass filters when the system has to operate over a wide range of carriers and signal bandwidths. But this solution proves to be costly not only because it occupies a large area, but it also costs significant design effort to build multiple front-ends and optimize them for a given set of specifications.

Eventually the RF community began to envision the holy grail of all RF front-ends, a "software-defined radio terminal" capable of being programmed to any radio standard, at any data rate, and any RF carrier [Mit95], [Abi07]. The ideal software-defined radio would be an FPGA-like device programmed externally through a computer interface. On the radio receiver (Rx) side this would mean a highly linear, high-speed ADC digitizing the received RF signal that is processed by the software programmable core. On the transmitter (Tx) side, a high-speed, highly linear DAC interface converts the digital RF data received from the programmable core into analog form suitable for transmission. Components like low-noise amplifiers (LNAs), power amplifiers (PAs), and bandpass filters will remain in the chain, but we get rid of analog components like mixers, amplifiers, and baseband filters that are necessary for signal conditioning in a traditional front-end.

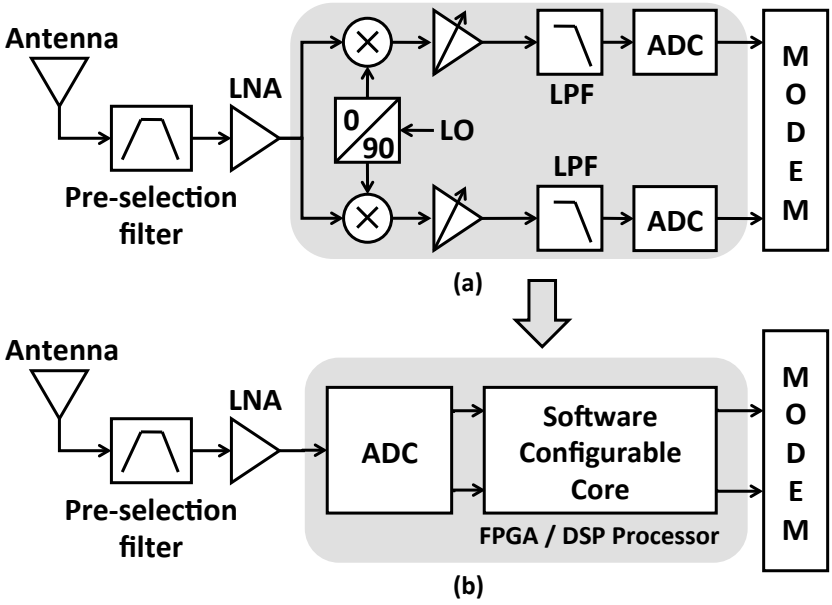

Figure 1.2: (a) Traditional radio receiver front-end (b) The ideal software-defined radio receiver.

The distinction between the two approaches is shown for the radio receiver in the block diagram of Figure 1.2. In Figure 1.2(b), the software radio replaces all the signal conditioning components of the radio receiver in Figure 1.2(a). The RF signal is first digitized by a high-speed ADC, following which it is amplified and filtered by the software-defined core. Parameters like carrier and bandwidth are programmed into the core in order to support multi-mode, multi-band operation.

In addition to solving the problem of multi-mode, multi-standard radio design, software-defined radios become drivers of cognitive radio systems. Cognitive radio refers to the concept of intelligent radios that sense a wide spectrum (e.g. 500MHz) and locate available bandwidth not in use by the primary users of that spectrum [Hay05]. This available bandwidth is then open for transmission or reception by the secondary user employing the cognitive radio. A cognitive radio system has two main components. The first is a spectrum sensing processor that identifies the location and span of the available bandwidth. The second component is a flexible radio transceiver that adjusts its RF carrier and signal bandwidth so as to tune itself to the position of the available band. The software-defined radio terminal is ideally suited to act as this flexible radio transceiver.

Both the receiver and transmitter have distinct challenges associated with their design when it comes to incorporating flexibility and processing signals over a wide bandwidth. First, on the receiver side, a highly linear ADC is needed to support wideband digitization across a spectrum spanning a few GHz. Not only must the ADC digitize the signal of interest, but it must also ensure that any blockers and interference present in the received wideband signal do not degrade the quality of the signal of interest. To achieve this without any narrowband filtering operations before the ADC is tough and demands dynamic range requirements that supersede the maximum dynamic range mandated by the target radio standards. Second, adding to the tough ADC specifications is the challenge of realizing a software-defined processor core that can effectively condition the multi Gs/s ADC output. If the software core is to ideally replicate the traditional radio receiver tasks of mixing, lowpass filtering, and decimation, then we are looking at 10-14 bits of digital signal processing (DSP) starting at a few Gs/s. To put this requirement into perspective, a typical DSP processor from TI can only handle a maximum of 350MHz clock rate [Ins08]. A third challenge is to move towards greater CMOS integration and implement all RF components in scaled technology. This would enable a system-on-chip realization of the entire transceiver chain that includes

the front-end and the baseband modem. In the next section we look at various solutions proposed to enable a fully-flexible integrated radio receiver.

1.1 Flexible Radio Receivers

Although the dream of a truly software-defined radio receiver has not yet become a reality, several versions of flexible radio receivers, demonstrating varying degrees of success, have been implemented. The original effort [GNS09] was to incorporate flexibility in the RF and analog filters that are part of a traditional analog front-end. The flexible characteristics were realized using a matrix of active and passive components inside the filters. The overall baseband selectivity could be switched between 3^{rd}- and 5^{th}-order by reconfiguring the filters. This approach was limited due to the fixed set of filter configurations that could be realized from the component matrix. Also, incorporating flexibility in analog components is difficult and relies mainly on altering fixed configurations of active and passive circuit components. In addition, the bulky nature of analog filters, the intrinsic non-linearity of the components, low dynamic range, and poor scaling of power with new technology generations make the design process extremely challenging.

A second approach implemented the filtering operation using discrete-time switched-capacitor circuits [MHM05], [SLE08] to alleviate the problems posed by analog components. This class of receivers uses discrete-time analog signal processing to downconvert, downsample, and filter the received signal. The processing occurs on RF data sampled at the carrier Nyquist rate, without analog-to-digital conversion [MS04]. The objective is to sufficiently suppress the out-of-band interferers in the received signal before the signal can be passed onto the ADC for digitization [BMC06], [MCB09]. This approach significantly lowers the dynamic range required by the ADC since the ADC only has to process the signal of interest in a small bandwidth. The frequency response of the switched capacitor filters is varied by tuning their capacitor sizes and sample rates. Overall, the flexibility achieved with discrete-time processing is much wider compared to using fixed set of analog components.

Another interesting use of discrete-time signal processing was made in [RDB08], where the authors used a pre-processing unit that uses discrete-time Fast Fourier Transform (FFT) before

digitization. The idea was to implement filtering of the received signal in the frequency domain by first taking the FFT of the signal, and then selecting the frequency bands of interest from the FFT output. A 4096-point analog FFT with sample rate at 5 Gs/s was used to accommodate RF carriers up to 5GHz. Once the FFT output samples are digitized, they must pass through an inverse FFT operation in the digital domain that brings the samples back to the time domain.

But RF receivers still lacked the capability to scale in performance and area with advanced CMOS technologies. The concept of digitally assisted front-ends has made this scaling possible with the advantage of fine-grain tunability of the components in the receiver chain [HMS05], [HZS08]. Digitally assisted front-ends use DSP components to implement signal-conditioning circuits in the RF receiver chain. In addition to being low cost and low power, digital components are more flexible. Reprogramming the frequency response of digital filters can be done on-the-fly by changing tap coefficients or bypassing cascaded stages.

In this work, we adopt an approach similar to the software-defined radio terminal of Figure 1.2, where a high-speed ADC is followed by a processor core, except that instead of using a generic FPGA/DSP processor we use a specialized digital front-end (DFE) expressly suited for performing RF signal-conditioning operations. A conceptual view of this front-end is shown in Figure 1.3. The RF signal received from the antenna is pre-filtered by a wide bandwidth RF filter, amplified, and digitized using a high-speed, high dynamic range ADC. The ADC is optimized to maximize its performance in the bandwidth of interest around the target RF carrier. The digital front-end unit first downconverts the digitized signal by mixing it with the RF carrier. The next task is to downsample the signal from the ADC sample rate of f_{s1} to the modem sample rate of f_{s2}. In the downsampling process, the DFE must lowpass filter the signal to minimize noise injection in the baseband, which occurs due to aliasing. Lowpass filtering is also necessary for channel selection, when the unwanted channels in the received signal are attenuated. The DFE must ensure that the combined effect of finite wordlength processing and sample-rate conversion does not degrade the noise figure of the receiver chain to levels below the target specifications of the receiver.

The block diagram in Figure 1.3 shows the ADC directly digitizing the received signal without any significant pre-conditioning. There have been other versions of digitally assisted front-ends where mixing and lowpass filtering is done before digitization. For example, the work in [HMS05]

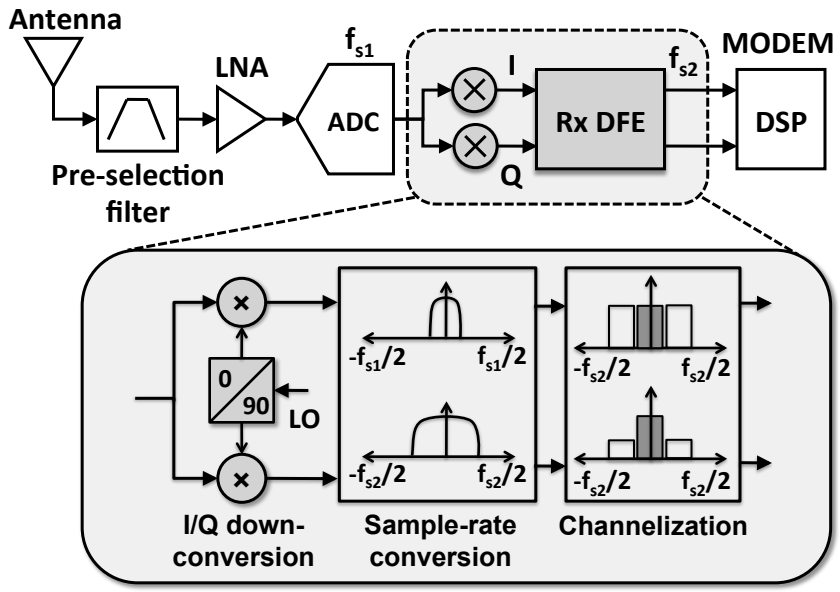

Figure 1.3: Conceptual digitally intensive radio receiver terminal.

proposed mixing of the RF signal in the analog domain, followed by analog lowpass filters and an ADC operating at a fixed rate of 104 Ms/s. Subsequent signal conditioning like channel selection, additional filtering, and downsampling is done digitally using DSP components that constitute the DFE. This approach lowers the burden on the ADC, which can now operate at lower dynamic range since a significant part of the received signal has been filtered in the analog domain. In essence this is similar to the discrete-time signal processing approach where the switched capacitor filters attenuated the out-of-band content in the received signal before digitization. However, we still face the problem of implementing analog mixers and lowpass filters that do not scale well in power or area. Hence our objective is to move the ADC as close to the antenna as possible, as is shown in Figure 1.3, and enable a direct-sampling receiver.

A direct sampling receiver samples and digitizes the bandpass RF signal at GHz rates and subsequent DSP blocks in the digital front-end, operating at high input sampling frequency, realize the downconversion and filtering. The main advantage of this approach lies in avoiding the use of bulky analog filters and mixers and having a fully reconfigurable front-end that scales in power and area with advancement of technology. The ADC, however, must operate at GHz sample rates with high dynamic range, making it very power hungry. Energy efficiency is therefore critically

important for the ADC. Implementing energy-efficient DFEs is also very challenging given the high throughput requirements of the digital mixer and filters that process the GHz-rate ADC output.

An example of such a conceptual multi-mode radio terminal, fully configurable through parameters in a digital front-end was presented in [HZS08]. This work proposed the use of a 17-bit ADC with 100dB dynamic range to cover standards like EDGE, IS-95, and UMTS between 0-5GHz. Such an ADC, however, is infeasible in present-generation CMOS at power levels ranging around a 100mW that is necessary for portable handset type of applications. With future scaling of CMOS, ADCs are expected to improve in sample rate and operate at reduced power levels and it is likely that such ADCs will become feasible with advancing technology. In present technology, however, a fair compromise is to reduce the range of RF carriers and radio standards such that an ADC with reasonable power consumption can handle the required wideband digitization. The work in [BAM09] presented an example of such an ADC operating up to 2.4GHz RF carrier and 20MHz signal bandwidth with an SNDR (signal to noise plus distortion ratio) of 37dB and a maximum power consumption of 26mW. Although the SNDR and dynamic range for this work was lower than the required 10-12 bits we need for standards like WCDMA and LTE, this work was a good indicator of the kind of ADC architectures needed to digitize narrowband radio signals at high sample rates.

In this thesis we analyze a high-speed direct-sampling receiver (Figure 1.3) and demonstrate the physical implementation of an energy-efficient digital front-end processor. The focus is on the system-level realization of a direct-sampling receiver that includes a pre-selection filter, a high-speed ADC, and a low-power DFE. We model a continuous-time $\Sigma\Delta$ ADC architecture and show that its SNR performance is ideally suited for implementation in a direct-sampling receiver. We also derive the specifications of the passive bandpass LC filter that will be required before the ADC in order to pre-filter the RF signal before digitization. And finally, we look at the micro-architecture and system design of a digital front-end that can be coupled with the modeled ADC to perform all the signal conditioning operations digitally and handoff data to the baseband MODEM. In the next section we switch gears and discuss the second problem of flexible radio transmitter design.

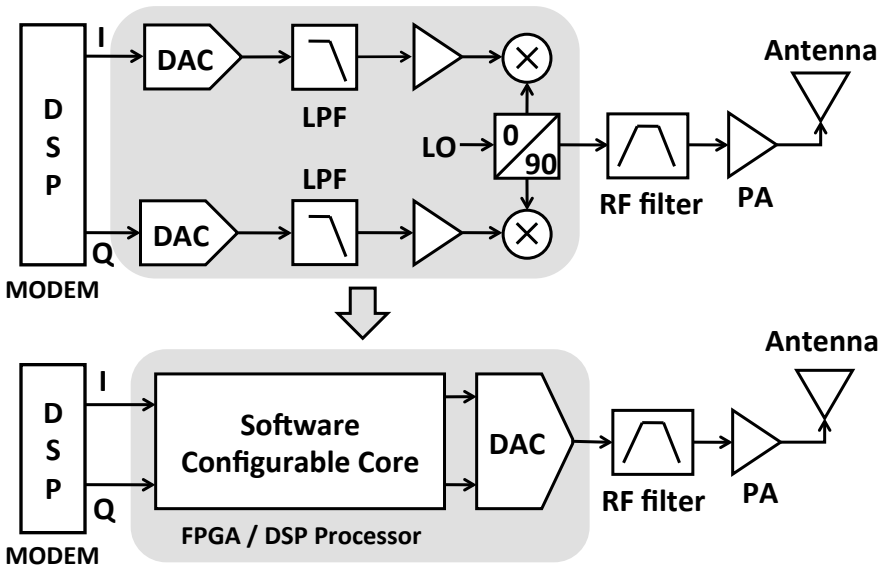

Figure 1.4: Traditional radio transmitter front-end and the ideal software-defined radio Tx.

1.2 Flexible Radio Transmitters

On the transmitter side, the ideal software-defined radio terminal can be visualized as shown in Figure 1.4. The challenge here is to get the baseband signal modulated at the RF carrier without leaking energy beyond the mandated threshold (defined by the strictest radio standard) in the adjacent radio bands. Traditionally, the radio transmitter uses a low-sample rate D/A converter to transform the baseband digital signal, sampled at rate f_{s2}, to the analog domain. Images of the baseband signal present at multiples of f_{s2} need to be suppressed after D/A conversion. Analog lowpass filters perform this suppression, following which the signal is mixed with the RF carrier. Additional RF filtering after mixing ensures that the transmitted signal adheres to the spectral emission mask and receive band noise specifications for frequency division duplex (FDD) operations. Finally, the transmitted signal is processed by a power amplifier that increases the output power level (~20dBm), before being passed to the antenna for transmission. Non-linear components in the Tx chain inject noise in the baseband signal that degrades the error vector magnitude (EVM) [MRM04] of the transmitted signal, ultimately leading to a degraded bit error rate [Bre03] on reception. The transmitter must restrict the extent of this noise injection to ensure that the EVM remains below levels required by the target radio standard.

A fully flexible, multi-mode, software-defined radio terminal replaces the analog lowpass filters and mixers, relying on a high-speed, linear DAC to transform the RF modulated waveform to the analog domain. The software-defined core will have to function at high sample rates in order to attenuate the baseband images to power levels lower than the emission mask and upconvert the baseband data to the RF carrier frequency. For a purely digital mixing process, the software core must upsample the baseband data to at least twice the RF carrier frequency, which translates to greater than 4Gs/s for RF carriers higher than 2GHz. FPGA or DSP processors, however, cannot perform this throughput-intensive task in present technology. Hence the best trade-off is to follow the approach we took on the receiver side and design a digital front-end tailored to perform the tasks of a flexible transmit modulator.

In this context, three classes of RF modulators have emerged:

- Outphasing modulators
- Polar modulators
- Linear modulators

The polar and outphasing modulators promise greater efficiency in the power amplifier (PA) stage since they resort to the use of nonlinear PAs that process constant-envelope phase-modulated signals. This is made possible by transforming the original I and Q baseband data to an alternate dimension.

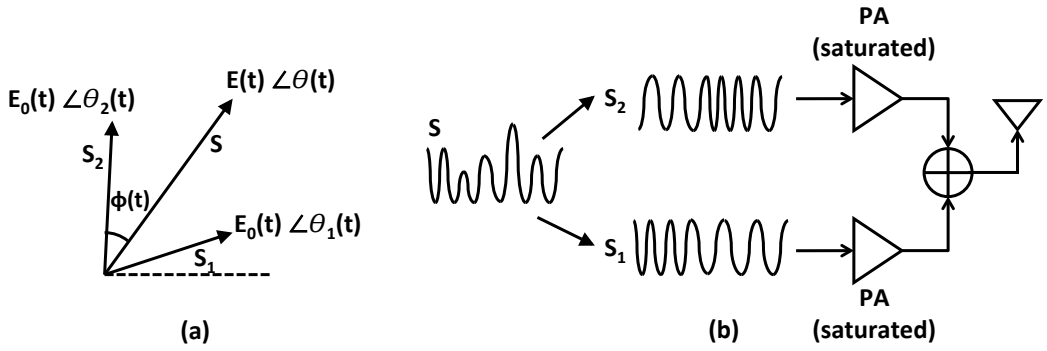

Figure 1.5: Outphasing transmitter (a) Defining the signal vectors (b) Power amplifiers.

The modulator for an outphasing PA [HLA09a] converts any arbitrary modulation, which has amplitude and phase information, into two phase-modulated waveforms. The phase-modulated waveforms, which have constant and equal amplitude but varying phase, when combined together regenerate the amplitude and phase information present in the original signal. As shown in Figure 1.5, the phase modulation is performed on θ, and the outphasing angle, ϕ, is responsible for amplitude modulation. The outphasing signals $\mathbf{S}_1$ and $\mathbf{S}_2$ are created by transforming the original signal $\mathbf{S}$ into two constant-envelope signals with phase $(\theta+\phi)$ and $(\theta-\phi)$. In the first step, a VCO generates the carrier, whose phase varies with θ. This phase-modulated carrier then branches into two paths that delay and advance the phase of the input by ϕ. The two signals are amplified by two independent PAs and combined at the PA outputs. This approach requires precise matching of the signals during their recombination in order to successfully generate the original waveform. A slight gain or phase mismatch between the two paths can result in significant distortion and spectral re-growth in the adjacent bands.

In a polar modulator [SRM03] the I, Q components are transformed to the polar (R, θ) domain and combined in a single output PA, as shown in Figure 1.6. The phase signal, θ, is passed to an all-digital PLL (ADPLL) that synthesizes the RF carrier frequency along with the phase modulation, given by the variation of θ with time. A constant envelope signal, phase modulated by θ is input to a saturated PA. The component R drives a DC-to-DC switching converter, the output of which modulates the PA supply voltage.

In the polar modulator the amplitude signal, R, and the phase signal, θ, undergo completely different sets of signal processing steps before being combined in the PA. The recombination of the

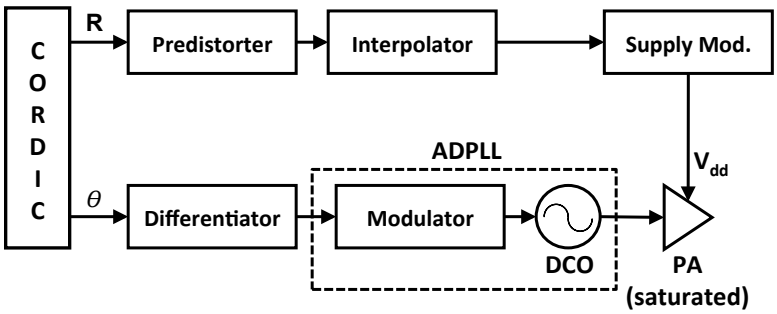

Figure 1.6: Polar modulator combines the amplitude, R, and phase, θ, signals in the PA.

amplitude and phase signals is compromised due to the inherent filtering and different signal processing on each path, resulting in distortion of the final output. The results presented in [KRS08] demonstrate the sensitivity of the polar architecture to the timing alignment of the phase and amplitude paths. It is also challenging to implement a linear, highly efficient, and wideband supply modulator, which is needed for this approach unless we choose to use a pre-driver circuit before the PA to combine the amplitude and phase signal, as was done in [SWR05], [BMK11]. The pre-driver approach, however, will require a linear power amplifier (PA) that results in loss of output efficiency in the PA stage.

The linear modulator, on the other hand, does not rely on any transformation of the (I, Q) data to an alternate dimension. Hence it does not suffer from the problem of mismatches during the direct combination of the (I, Q) components. The price to pay in this case is the use of a linear PA in the output stage that results in an overall lower efficiency of the transmitter [HLA09a]. But the linear modulator has several benefits to compensate for this loss. In addition to being versatile, linear modulators can support complex modulations like 64-QAM as well as OFDM type of data multiplexing [PZH08]. They are not constrained by the bandwidth limitations of the polar modulator and are therefore an attractive solution for high data rate communication systems.

As mentioned before, the signal processing in a linear transmitter occurs in the Cartesian coordinate system on the I and Q components of the symbols. The attractive trends of speed and power scaling of digital CMOS, coupled with the relative ease of its system integration, has encouraged the digitization of the linear transmitter. Digital signals, however, are rife with quantization noise and also have images of the baseband signal present at multiples of the modem sample rate f_{s2}. The transmitter must suppress these images and quantization noise in the RF spectrum to ensure that the output signal satisfies emission mask requirements. This requirement sets a maximum limit on the transmitted power levels in frequency bands away from the signal of interest.

The basic signal processing flow of a digitally intensive linear modulator is shown in Figure 1.7. The (I,Q) components are upsampled, lowpass filtered, and combined in an RFDAC, which is responsible for D/A conversion as well as upconversion of the baseband data to the RF carrier frequency. The baseband signal from the modem at sample rate f_{s2} is upsampled by the Tx digital front-end (DFE) to a higher sample rate f_{s1}. In the process it filters the images of the baseband

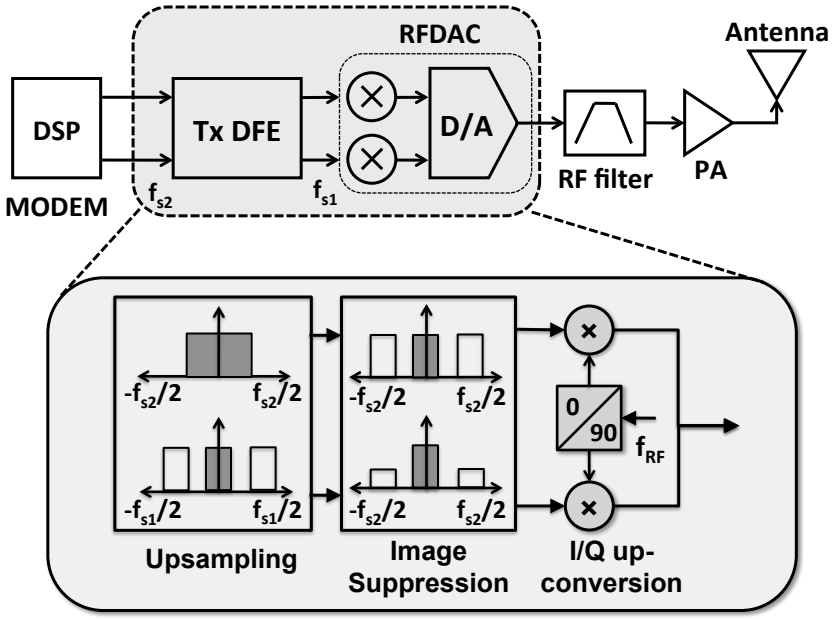

Figure 1.7: Basic signal flow in a linear modulator.

signal present at multiples of f_{s2} to ensure that the images do not violate the emission mask. The DFE also performs fractional sample-rate conversion if the frequency f_{s1} and f_{s2} are not related by integer factors. The DFE architecture can be designed to include tuning parameters that vary attenuation characteristics, upsampling ratios, signal bandwidths, and extent of quantization noise suppression. The RFDAC following the DFE is a current-steering D/A converter as well as a mixer. The I and Q current signals centered on the RF carrier are combined at the output of the RFDAC.

The wordlength of the digital samples input to the RFDAC and the sample rate f_{s1} are important factors that determine whether the quantization noise and amplitude level of the images in the RF spectrum are sufficiently low. Noise suppression can be guaranteed if the digital word in the data stream is sufficiently wide, since quantization noise power is inversely proportional to the number of bits. The work in [BMK11] uses a 17-bit RFDAC to ensure that noise suppression is achieved for a 5MHz bandwidth signal without the aid of an on-chip bandpass filter. The work in [ESK07] uses digital signals as wide as 10-bits and sampled at 307.2Ms/s to achieve an acceptable noise level of 56dB in the adjacent band, and image suppression of 44dB at 5MHz signal bandwidth. The low sample rate of 307.2Ms/s is sufficient to ensure 44dB suppression for a 5MHz signal, but the sample rate will have to increase significantly (more than linear) for higher bandwidth

signals. Also, longer wordlengths lead to higher power consumption in the RFDAC. Increasing wordlength also increases the probability of mismatches in the unit current cells used in the D/A converter [AGA04] that ultimately degrades the linearity of the DAC.

A proposed solution to the problem of reducing wordlength in the RFDAC is the use of ΣΔ modulation [PZH08]. This technique reduces the DAC resolution while shaping the added quantization noise away from the signal of interest. An RF bandpass filter following the RFDAC can filter out the additional out-of-band quantization noise introduced by the ΣΔ modulator (SDM). ΣΔ modulation is a promising step towards enabling low-power digital-to-RF signal generation, but conventional modulators [PZH08], [HLA09b] require a high degree of signal oversampling. Oversampling is needed to improve the in-band SNR of the transmitted signal and also to increase the 3 dB bandwidth of the RF bandpass filter, shown after the D/A convertor in Figure 1.7, and enable its on-chip integration. Higher sampling frequency, however, increases the power consumed in the digital blocks, since power consumption is directly proportional to the sample rate.

In [HLA09b] the authors use a sample rate of 4 times the RF carrier (f_{s1} = 4Gs/s @ f_{RF} = 1GHz) with a 1-bit ΣΔ output for a 5MHz QPSK modulated signal. Although this approach results in a 1-bit DAC, it limits the range of RF carriers due to the high oversampling ratio of $4f_{RF}$. The work in [PZH08] uses a DFE that upsamples the baseband data to $2f_{RF}$ (f_{s2} = 5.4Gs/s @ f_{RF} = 2.7GHz) before being processed by an SDM and RFDAC. The high DFE output sample rate enables the use of a low 3^{rd}-order ΣΔ modulator and a low-Q bandpass filter. This approach, however, suffers from high power consumption of close to 100mW in the digital core (DFE + SDM) that is a result of processing digital data at rates of 5.4Gs/s.

In this work we aim to improve the power efficiency of the digital core while also ensuring that a passive integrated BPF can eliminate any spectral noise that violates the emission mask after D/A conversion. This was achieved through co-optimization of the noise transfer function in the ΣΔ modulator and the bandpass filter response such that we minimize the sample rate, f_{s1}, as well as the DAC resolution. The ΣΔ modulator lowers the sampling frequency for a given RF filter bandwidth through the use of an FIR noise shaping filter. This allows the DFE and SDM to operate at a relatively low frequency of $f_{RF}/3$ (f_{s1} = 900Ms/s @ f_{RF} = 2.7GHz). We also use a time-interleaved RFDAC architecture that aids in filtering the out of band quantization noise. The

quantization noise introduced by the SDM is suppressed to some extent (~20dB) by the inherent filtering operation of the 2-channel time-interleaved DAC. An integrated RF bandpass filter can then filter out the remaining noise. Consequently, the digital chain functions at a reduced sample rate and switching power without compromising the extent of on-chip integration of the transmitter.

1.3 Thesis Outline

This thesis continues with a discussion on the choice of ADC architectures and optimization techniques in Chapter II and is followed by the design of a high-speed energy-efficient DFE processor in Chapter III. We take a detailed look at the digital components in the transmitter chain in Chapter IV. This will include a discussion on the FIR noise shaping employed in the ΣΔ modulator. Chapter V will talk about the architecture and filtering effects of the time-interleaved RFDAC. This is followed by a description of the prototype test chips for the Rx DFE and the RF transmitter fabricated in 65nm CMOS and their measured results in Chapter VI.

CHAPTER 2

A/D Converters for Flexible Receivers

A fully reconfigurable direct-sampling radio receiver covering standards like GSM, EDGE, UMTS, WLAN, and LTE will need to cover all the specifications listed in Table 2.1. The specifications demand support for RF carriers between 450MHz-5GHz, channel bandwidths of upto 40MHz, and dynamic range of 90dB, if we look at the strictest requirement of the individual standards. In scenarios where the wideband system has to process several strong radio signals, the dynamic range requirements get worse. According to the work in [HMS05], the digitization process for such a multi-standard radio receiver terminal will require an ADC operating at 10GHz with 17 bits of resolution.

Such an implementation is infeasible in present technologies with acceptable power levels for portable handset type of applications. So we restrict the scope of our work to RF carriers up to 2.7GHz, signal bandwidths of 20MHz, and standards like WiMAX/LTE that require 10-12 bits at

Standard	Carrier Frequency (MHz)	Channel Bandwidth	ADC Dynamic Range (dB)
GSM	380-1900	200kHz	90
EDGE	850-1900	200kHz	87
GPRS	850-1900	200kHz	84
CDMA2000	450-2100	1.228MHz	80
W-CDMA / UMTS	728-2690MHz	3.84MHz	70
IEEE802.11a	5150-5825	20MHz	55
IEEE802.11b	2400-2495	20MHz	55
IEEE802.11n	2400/5000	20/40MHz	55
LTE	728-3800MHz	1.4/3/5/10/15/20MHz	76/72/70/70/70/70

Figure 2.1: RF carriers, channel bandwidths, and dynamic range requirements for various radio standards.

the MODEM input over a 20MHz wide signal band. For wideband digitization, the actual dynamic range will have to be higher than 12 bits, since we not only have to digitize the signal of interest but also ensure that any strong signals/blockers present in the received signal, which contains several different RF channels, do not distort the signal of interest. Hence the target dynamic range has to be greater than 12 bits with a margin of at least 6-10dB to accommodate such wideband sampling. A $\Sigma\Delta$ ADC was the preferred choice due to the oversampled nature of the digitization process, and also because of its expected lower power consumption.

System-level decisions regarding the ADC sample rates, number of bits, and order of $\Sigma\Delta$ modulation go a long way in determining whether the SNR specifications of various radio standards can be achieved at power levels of 50 to 100mW. The ADC sample rate is the most important variable, since it not only determines the operating frequency of the ADC, but also that of the subsequent blocks in the DFE chain. The ADC sample rate has a close relation with the RF mixing process, which is done immediately after digitization.

Mixing is the process of translating the RF modulated data to the baseband. When we sample a continuous-time signal at a finite rate of f_{s1}, blocks in the analog spectrum spanning a frequency of f_{s1} fold back in the digital domain to be centered between $-f_{s1}/2$ and $+f_{s1}/2$. An RF signal modulated at carrier f_{RF} has an RF bandwidth of $2f_{RF}$. According to Nyquist criterion, the signal must be sampled at rates above $2f_{RF}$. But sampling at twice the Nyquist rate is not mandatory for radio applications, since the signal bandwidth is restricted to only several MHz (20MHz). We can undersample at a rate $f_{s1} < 2f_{RF}$ and still obtain a replica of the signal.

This process is shown for two combinations of f_{RF} and f_{s1} in Figure 2.2. In Figure 2.2(a), the ADC sample rate is greater than $2f_{RF}$ resulting in the RF data positioned at f_{RF} before and after digitization. In Figure 2.2(b) we undersample the signal at $f_{s1} = 2f_{RF} - W$, which is lower than $2f_{RF}$. This results in an image of the original RF data at $(f_{s1} - W)/2$ after digitization. The noise level, shown in dark grey, rises after undersampling in the second case, since the fixed quantization noise power spreads in a smaller frequency range. In the first case, the digitized RF samples have to be multiplied by the sine and cosine of f_{RF} during mixing, while in the second case the multiplication is done with the sine and cosine of $(f_{s1} - W)/2$. But implementing a digital mixer at the GHz sample rate of f_{s1} becomes either timing infeasible or very power inefficient for a multi-bit ADC output.

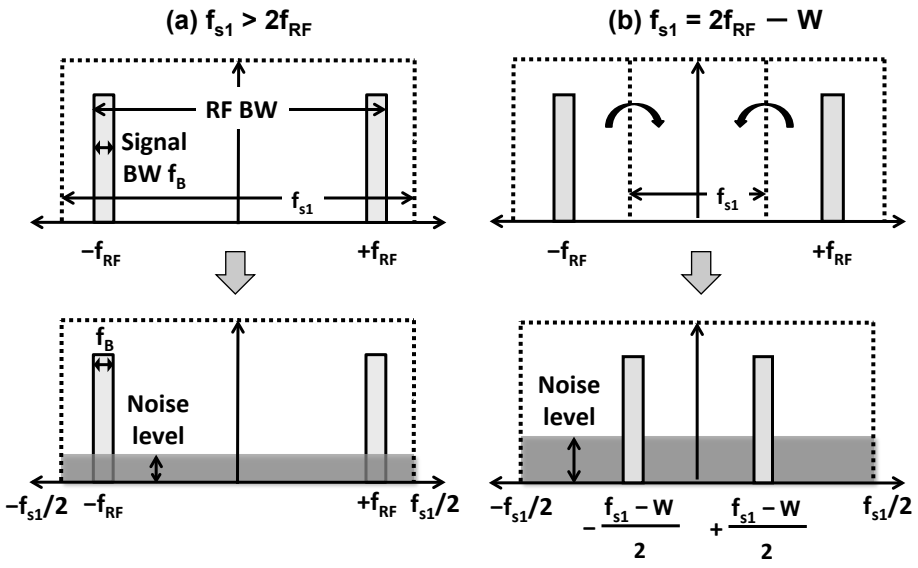

Figure 2.2: ADC sampling (a) Nyquist criterion ($f_{s1} > 2f_{RF}$) (b) Undersampling ($f_{s1} = 2f_{RF} - W$).

In this scenario, the second case of undersampling can be exploited to make the mixing operation trivial after digitization.

A special case is when $f_{s1} = W = f_{RF}$ and the signal can be translated to the baseband directly after sampling. This seems like an attractive solution, but since there is no distinct mixer following the ADC, we will now need two ADCs, one for the I path and the other for the Q path. Both paths are sampled by clocks that are $\pi/2$ out of phase to generate the I and Q samples at the baseband. Having two ADCs, however, will result in mismatches between the characteristics of the two converters, leading to I/Q imbalance in the baseband data.

A second solution is when $f_{s1} = 4/3f_{RF}$ and $W = 2/3f_{RF}$, resulting in an image of the signal at $1/3f_{RF}$, shown in Figure 2.3. For a sample rate of $f_{s1} = 4/3f_{RF}$, the frequency of $1/3f_{RF}$ corresponds to the $\pi/2$ position in the digital domain. This implies that the ADC output has to be multiplied with samples of $\sin(\pi/2 \cdot n)$ and $\cos(\pi/2 \cdot n)$, $n \in I$, during mixing. These functions are repeated sequences of $\{1, 0, -1, 0\}$ that make the mixing operation trivial. The mixing can then be implemented using just a multiplexer that either selects the input, its compliment, or a zero depending upon a control signal from a counter. This solution only needs a single ADC and hence avoids the problem of I/Q imbalance. The penalty is an increased sample rate of $4/3f_{RF}$ which is 33% larger than the first case of $f_{s1} = f_{RF}$, but the advantage lies in using a single ADC instead of two ADCs that were needed

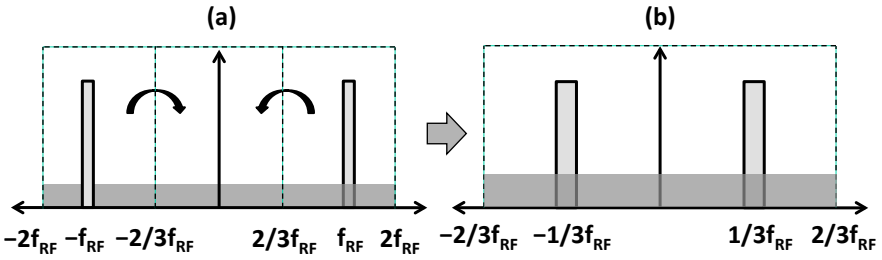

Figure 2.3: (a) Undersampling at $f_{s1} = 4/3f_{RF}$. (b) RF data centered at $\pi/2$.

in the first case. Hence we chose a sample rate of $f_{s1} = 4/3f_{RF}$ in this work. Before we move on to selecting the ADC topology, it is worthwhile to review the basics of ΣΔ modulation that will be used extensively throughout this chapter.

2.1 ΣΔ Modulation

Since the ADC sample rate is $4/3f_{RF}$ and the RF data is positioned at $1/3f_{RF}$ (not DC) after digitization, we need a bandpass ΣΔ modulator. An example of a 2nd-order bandpass ΣΔ modulator is shown in Figure 2.4. The modulator uses a bandpass filter B(z) to shape the quantization noise e(n) introduced by the 1-bit quantizer QN. The quantization occurs on samples of $-$v(n$-$2), which is a result of bandpass filtering x(n), the difference between u(n) and y(n). The overall effect is to shape the added quantization noise, e(n), away from the signal of interest, which is the bandpass input signal u(n). The modulator transfer function is derived in equations 2.1-2.6.

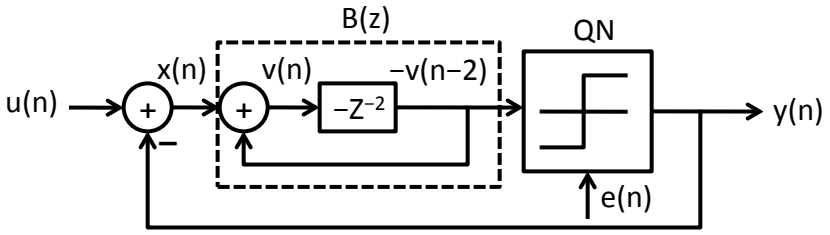

Figure 2.4: A 2nd-order bandpass ΣΔ modulator with 1-bit quantizer.

$$x(n) = u(n) - y(n) \tag{2.1}$$

$$v(n) = x(n) - v(n-2) \tag{2.2}$$

$$y(n) = -v(n-2) + e(n) \tag{2.3}$$

$$v(n) = u(n) - e(n) \tag{2.4}$$

$$y(n) = -u(n-2) + e(n-2) + e(n) \tag{2.5}$$

$$Y(z) = -z^{-2}U(z) + E(z)(1+z^{-2}) \tag{2.6}$$

From the z-transform of the output we can see that $Y(z)$ is a combination of the input $U(z)$ and the filtered noise, $E(z)(1+z^{-2})$. The function $(1+z^{-2})$ has notches at $\pi/2$, a desirable property since the input signal $U(z)$ is centered at $\pi/2$. The power spectral density of the ΣΔ modulator output for a sinusoidal input is shown in Figure 2.5. We can see that the notches at the $\pm\pi/2$ position shape the quantization noise away from the sinusoid, thus maximizing the SNR in a limited bandwidth around the signal of interest. The excess noise can be filtered away and after decimation we get the spectral content in the bandwidth marked f_B.

The primary variables that affect the noise performance of the ΣΔ modulator are the oversampling ratio ($f_{s1}/2f_B$), the number of output bits, M, in the quantizer, and the order of modulation, N.

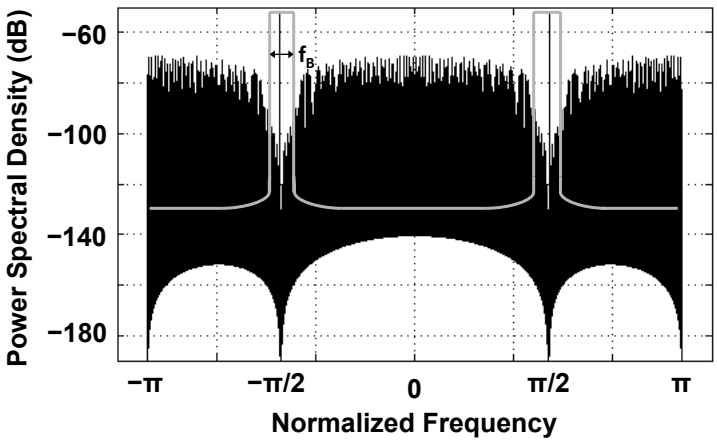

Figure 2.5: Power spectral density of a 2nd-order bandpass ΣΔ modulator with 1-bit output.

The order of modulation is determined by the order of the noise shaping function. For example, the noise shaping function $(1+z^{-2})$ implies second-order modulation. Since we fixed the sample rate of the modulator at 4/3f_{RF}, the design variables are now restricted to the order, N, and the number of output bits, M. A third variable available to us is the choice between a discrete-time modulator and a continuous-time modulator. In a discrete time-modulator the sample and hold operation is performed outside the ΣΔ loop. The modulator in Figure 2.4 is an example of a discrete-time topology. On the other hand, in a continuous-time modulator the sampling operation is done at the input of the quantizer inside the ΣΔ loop. We take a look at the trade-offs associated with these design variables in the next section.

2.2 ADC Architectures

Various ADC architectures with varying sample rate, number of output bits, and order of ΣΔ modulation can be implemented and a summary of published results is presented in the graph in Figure 2.6. We show a comparison between 2 classes of ADC architectures in Figure 2.6. The first, shown by circles, are lowpass ADCs that process baseband signals, and are preceded by a mixer. These represent the A/D converters used in conventional radio chains. The power numbers for references [SP12], [DSM04] are doubled from that of the published results to account for their usage in both I and Q channels. The second class, shown by diamonds, are ADCs that directly digitize

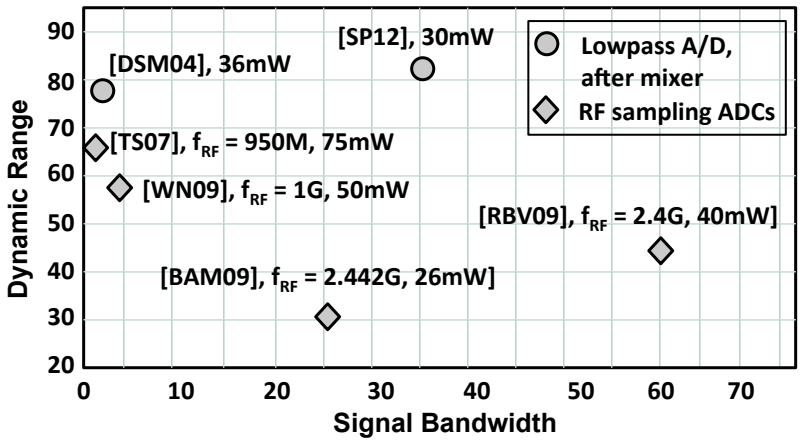

Figure 2.6: Comparison of various ADC architectures.

the RF signal before down conversion and are suitable for use in direct sampling receivers, and all of these use ΣΔ modulation in one form or another [TS07], [WN09]. It is apparent that the RF sampling ADCs must improve in dynamic range in order to meet the specifications of prevalent radio standards. Realizing a power-efficient, high dynamic range, and high-speed ADC within a limited power budget is one of the biggest bottlenecks a direct sampling receiver must overcome.

Let us now narrow down on potential ΣΔ ADC architectures for the direct-sampling receiver. A natural choice is the 2^{nd}-order 1-bit ΣΔ modulator, where the feedback loop is stable and the design complexity is low. This topology when implemented for continuous-time ADCs at f_{RF} = 2.4GHz, bandwidth of 20MHz, and a sample rate of 3.2GHz only obtains a maximum SNR of 37dB [BAM09]. In order to improve SNR, we can increase the modulation order, N and the number of output bits, M in the quantizer. One approach is to keep a 1-bit quantizer and only increase the order of modulation. In [RBV09], the authors use a 6^{th}-order bandpass 1-bit ΣΔ modulator, achieving an SNDR of 40dB for a 2.4GHz RF signal and a bandwidth of 60MHz. If the modulation order is increased further for greater SNR performance, we face problems of higher power consumption, large area, and mismatch between the 2^{nd}-order resonators that are used in the ΣΔ loop. Hence, a better choice is to vary both the modulation order and the number of quantizer bits to meet the required resolution of greater than 12 bits.

In order to avoid nonlinearity and mismatch problems in the multi-bit feedback DAC, it is necessary to minimize the number of output bits in the quantizer. With this in mind, we propose a continuous-time 4-bit 4^{th}-order bandpass ΣΔ ADC. This choice of order and resolution can enable the target dynamic range of approximately 80dB in a 20MHz bandwidth. A continuous-time architecture is the preferred choice since it makes high-speed operation of the ΣΔ loop feasible. The addition of an anti-aliasing filter before the quantizer is also an attractive feature of this ADC. More details on this filtering process will be presented in Section 2.4. In the next section we will look at the ADC modeling, signal and noise transfer functions, SNR variation, and the effect of circuit-level non-idealities on the ADC performance.

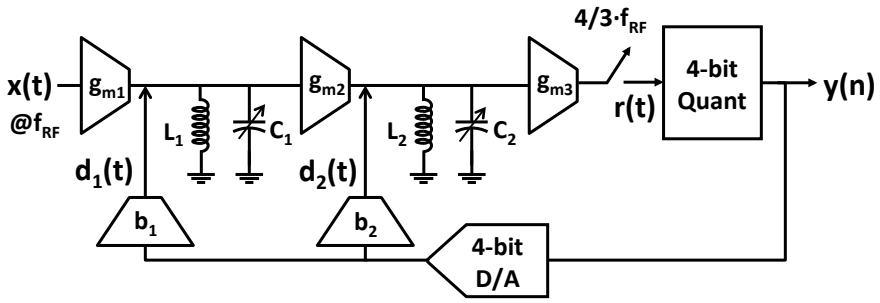

Figure 2.7: 4^{th}-order bandpass ΣΔ ADC with 4-bit quantizer. LC resonators are tuned at f_{RF}.

2.3 ADC Modeling

The proposed ΣΔ ADC model is shown in Figure 2.7. The input signal x(t) centered at carrier f_{RF} is applied to a transconductance amplifier (g_{m1}) that converts the voltage waveform to a current signal. The signal is then added to the first feedback current signal, $d_1(t)$, from the DAC, after which it passes through a 2^{nd}-order LC resonator centered at f_{RF}. A second transconductance amplifier (g_{m2}) is next in the chain, followed by addition of feedback signal, $d_2(t)$, from the DAC and another resonator, also centered at f_{RF}. The continuous-time signal r(t) at the output of the second LC resonator is sampled at $4/3f_{RF}$ and passed to the 4-bit quantizer for digitization. The feedback DAC takes the quantized signal as input and generates analog feedback signals $d_1(t)$ and $d_2(t)$. The capacitive component of the LC resonator is tunable to adjust the resonance frequency to variable carrier, f_{RF}.

The ADC model can be simplified to the equivalent form shown in Figure 2.8, where the continuous-time input signal x(t) is pre-filtered by $G(s)$ before being sampled at $4/3f_{RF}$. The feedback signal from the D/A (modeled as a zero-order-hold / ZOH) is processed by the continuous-time loop filter $H(s)$, following which it is sampled and added to the pre-filtered samples from $G(s)$. The output y(n) is obtained by adding the quantizer error, e(n), to these discrete samples. $G(s)$ is the open-loop transfer function from the input x(t) to the quantizer input r(t) (Figure 2.7). From Figure 2.7, we can see that $G(s)$ is a 4^{th}-order continuous-time function given by:

$$G(s) = \frac{k(s+\frac{w_R}{Q})^2}{(s^2+\frac{w_R}{Q}s+w_R^2)^2} \tag{2.7}$$

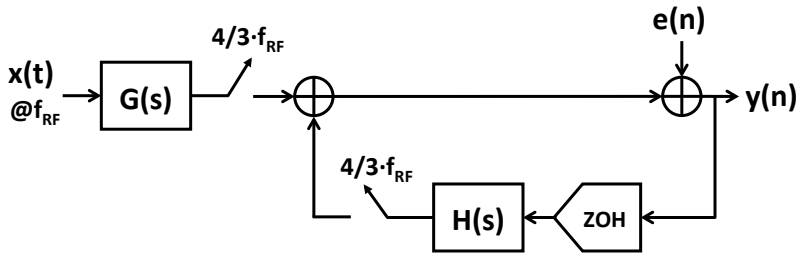

Figure 2.8: Simplified model of the bandpass ΣΔ ADC. ZOH is the zero-order-hold function.

Here w_R is the resonance frequency of the LC tank, and Q is the selectivity of the LC resonator. Ideally, the value of Q should be infinite, but non-idealities in the resonator result in finite selectivity. The non-ideality appears in the form of an intrinsic wire resistance, R_W, which is in series with the inductor. When R_W is taken into account, the selectivity, Q, of the resonator is given by:

$$Q = \frac{w_R}{w_B} = \frac{w_R}{R_W/L} = \frac{2\pi\sqrt{L}}{R_W\sqrt{C}} \tag{2.8}$$

Here w_B is the signal bandwidth of the resonator and is expressed as:

$$w_B = \frac{R_W}{L} \tag{2.9}$$

The resonance frequency, w_R, in Eq. 2.8 approximately equals $\frac{1}{\sqrt{LC}}$ under the assumption $w_B << \frac{1}{\sqrt{LC}}$. For a fixed implementation of the inductor, the values of R_W and L remain fixed. In Figure 2.8, $H(s)$ is the open-loop transfer function from the D/A output to the quantizer input. The z-domain equivalent, $H(z)$, of the loop filter transfer function can be computed by discretizing the cascaded functions $H(s)$ and ZOH. The closed-form expression for $H(z)$ is obtained using Eq. 2.10.

$$H(z) = Z\left\{L^{-1}\left\{\frac{1-e^{-sT}}{s}H(s)\right\}\right\} \tag{2.10}$$

The function $Z\{\}$ represents the z-transform and the function $L^{-1}\{\}$ is the inverse Laplace

transform. The sample period T equals 3/(4f$_{RF}$). The z-domain noise transfer function, $NTF(z)$, of the ADC can be computed from the loop transfer function $H(z)$ using Eq. 2.11.

$$NTF(z) = \frac{1}{1 - H(z)} \tag{2.11}$$

The signal transfer function, $STF(w)$, is a combination of the pre-filter $G(jw)$ and the noise transfer function, $NTF(e^{jwT})$. It is given by Eq. 2.12, which follows the derivation in [ST04].

$$STF(w) = G(jw)NTF(e^{jwT}) \tag{2.12}$$

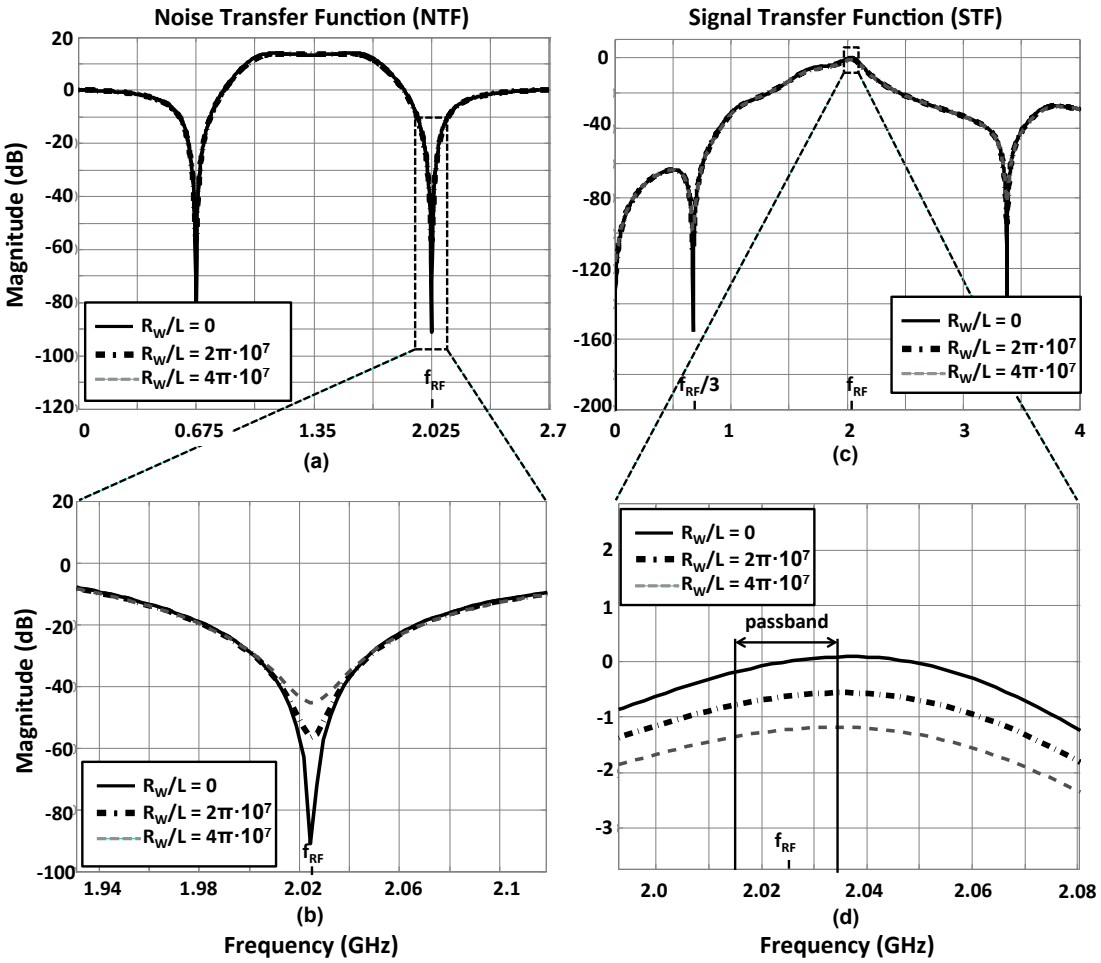

Figure 2.9: (a) Noise transfer function at $R_W/L = 0$, $2\pi10^7$, and $4\pi10^7$. (b) Zoomed-in NTF shows notch at f_{RF}. (c) Signal transfer Function at $R_W/L = 0$, $2\pi10^7$, and $4\pi10^7$.
(d) Zoomed-In STF shows passband around f_{RF}.

The ratio R_W/L strongly affects the noise and bandwidth performance of the ADC. This can be observed from the variation of the STF and NTF as a function of R_W/L. Figure 2.9 plots the NTF at three distinct values of R_W/L. The chosen values of R_W/L corresponds to bandwidth, w_B, increasing in steps of 10MHz, starting from the ideal value of 0 corresponding to infinite Q, and then going to $2\pi\times\ 10\times10^6$ and $2\pi\times\ 20\times10^6$. The transfer function, evaluated at f_{RF} = 2.025GHz and f_{s1} = 2.7GHz, has notches at frequencies $f_{RF}/3$ and f_{RF}. The solid line shows the ideal STF when $R_W = 0$. When R_W/L is increased, the notch is less steep, indicating an increase in noise power. This can be observed from the zoomed-in NTF shown in Figure 2.9(b), where the depth of the notch reduces from 90dB to 45dB when R_W/L is raised from 0 to $2\pi\times\ 20\times10^6$.

The signal-transfer-function for the continuous time bandpass ADC is shown in Figure 2.9(c). The STF has 0dB magnitude around the resonance frequency f_{RF} = 2.025GHz, and notches around DC, $f_{RF}/3$, and $5/3f_{RF}$. The notch at DC occurs because the function $G(s)$, in Eq. 2.7, has a zero close to DC. The notches at $f_{RF}/3$ and $5/3f_{RF}$ appear because of the NTF nulls in these bands, as can be predicted from Eq. 2.12. The variation of the STF with increasing R_W/L is shown in the zoomed-in plot of Figure 2.9(d). It can be seen that the peak of the response reduces with the increase in R_W/L. The passband range of 20MHz around the RF carrier is nearly flat for all 3 cases. The attenuation at the passband edge w.r.t. f_{RF} is 0.16dB for $R_W/L = 0$, and reduces to 0.1dB for $R_W/L = 4\pi 10^7$. The small droop in the passband can be corrected in the digital domain by varying the passband characteristics of the digital FIR filters in the DFE.

Hence, from the perspective of optimizing the NTF and minimizing the noise power, the target is to lower the coil resistance, R_W as much as possible. In general, it is difficult to minimize the resistance, and in practical realizations of bandpass filters, active Q-enhancement techniques are used. A common technique is to use a negative resistance [BAM09] in parallel with the inductor that lowers the intrinsic wire resistance of the inductor. The negative resistance is realized by using a pair of cross-coupled transistors in parallel with the LC tank as shown in Figure 2.10. The cross-coupled pair realizes a resistance of $-2/g_m$ ($g_{m1} = g_m = g_m$) parallel to the LC circuit,

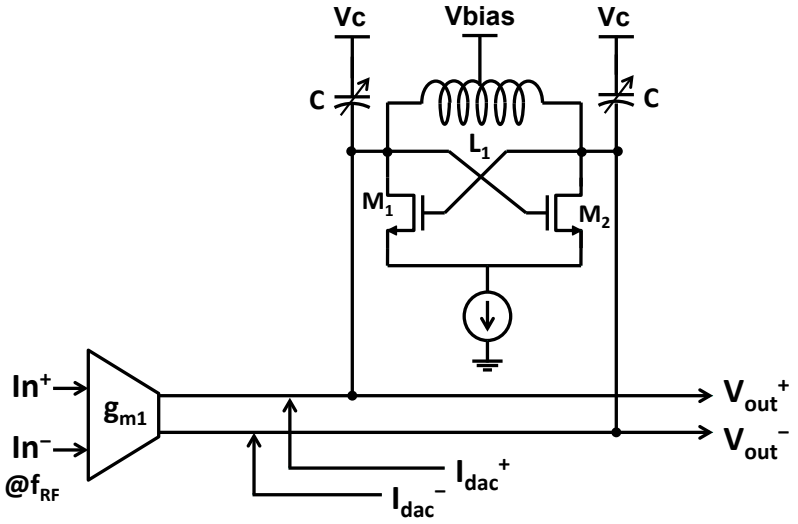

Figure 2.10: Q enhancement of inductor with negative resistance realized using cross-coupled transistors M_1 and M_2.

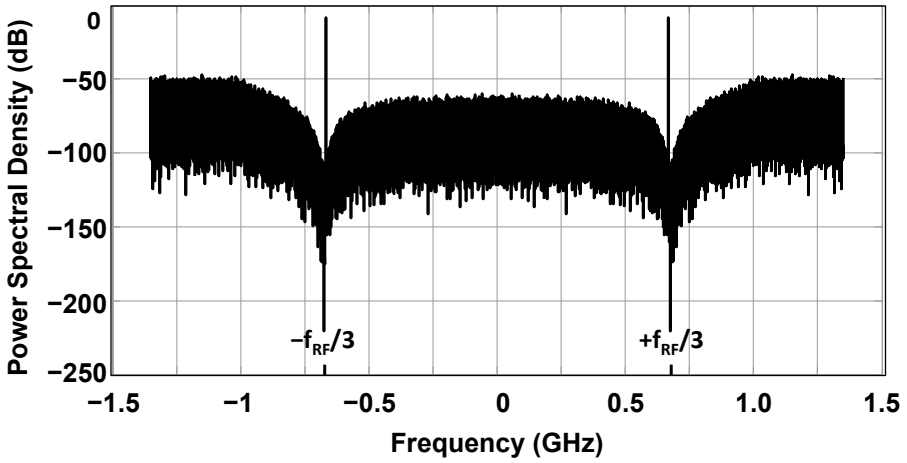

Figure 2.11: Spectrum of a 2.025GHz RF modulated sinusoid undersampled at 2.7GHz ($4/3f_{RF}$). The signal replica is centered at 0.675GHz.

which reduces the effect of the coil resistance R_W. Here g_m refers to the transconductance of the transistors.

Figure 2.11 shows the power spectral density of an 8MHz sinusoid modulated at f_{RF} = 2.025GHz and digitized at f_{s1} = 2.7GHz. The noise spectrum has a notch at $f_{RF}/3$ = 675 MHz. The R_W/L ratio for the inductors in the LC resonator sections was chosen to be $(2\pi)\times 20 \times 10^6$, since this value results in the ADC meeting the target SNR performance. After undersampling, the signal replica is formed at $f_{RF}/3$ = 675MHz. The oversampling ratio for a bandpass ΣΔ ADC is given by the ratio of half the sampling frequency and the signal bandwidth, $f_s/(2f_B)$ [NST96]. For a sampling frequency of 2.7GHz and a 20MHz signal bandwidth, the oversampling ratio of the ADC is 67.5. The maximum SNDR was 83.5dB for this example and the dynamic range was 85dB which makes it close to 14-bit resolution, providing sufficient margin for wideband digitization above the mandatory 12-bit resolution.

The variation of SNR with input signal amplitude at RF carriers of 2.0GHz and 2.7GHz is shown in Figure 2.12. The LC resonators were tuned to f_{RF} in each case. As expected, the SNR increases linearly with the input signal amplitude, and when operating close to maximum signal amplitude, the SNR drops due to overloading of the quantizer. For a fixed signal bandwidth, the SNR improves marginally with higher RF carrier frequency due to increase in oversampling ratio. The maximum SNR for the 2.7GHz RF signal (f_{s1} = 3.6GHz) is around 90dB, which is 6.5dB

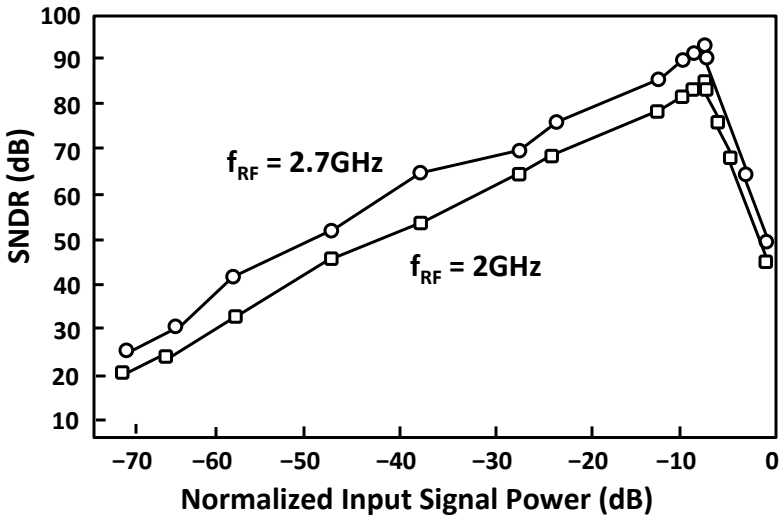

Figure 2.12: SNR variation with input signal amplitude and RF carrier frequency.

higher than the 2.0GHz signal. A major factor that limits the achievable SNR in continuous-time ADCs is the jitter in the sampling clock. We analyze the effect of clock jitter in the next section.

2.4 Clock Jitter Analysis

Continuous-time ADC architectures have increased sensitivity to clock jitter. We verified this by adding jitter to our simulation models and computed the SNR under the influence of a jittered clock. The clock jitter was derived from a zero-mean Gaussian distribution with a standard deviation equal to the RMS value of the clock jitter. The variation of SNR with increasing jitter is presented in Figure 2.13 for a 2GHz RF carrier and a 2.66GHz clock ($4/3 \cdot f_{RF}$). The jitter was varied to a maximum of 500fs, which is 1.33% of the clock period, when the SNDR degrades to around 45dB. The results in Figure 2.13 indicate that the system can tolerate about 250-300 fs of jitter. The SNR degradation is sharp if the jitter standard deviation exceeds 300 fs.

This jitter specification can be met by integrating a low-jitter PLL with the receiver [HSP08]. Alternatively, jitter sensitivity of the continuous-time ADC can be reduced by employing a switched capacitor resistive (SCR) feedback DAC [OGM05]. The SCR DAC, however, has an exponentially decreasing feedback pulse that requires a higher peak current compared to conventional switched-

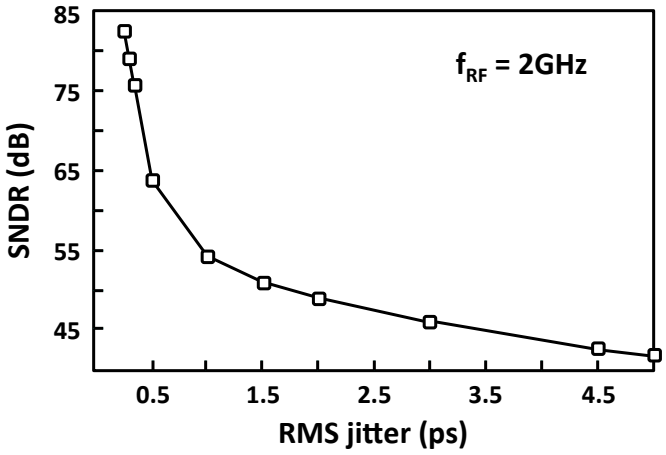

Figure 2.13: SNR degradation with increasing jitter in the clock.

current DACs, resulting in higher power consumption.

Since the undersampling ADC digitizes a limited bandwidth of $4/3f_{RF}$, we need a pre-selection filter before the ADC to prevent aliasing of any signal content that aligns with the desired signal after undersampling. In the next section we study the requirements of this pre-filtering operation.

2.5 Pre-selection Filter

A pre-selection filter is needed before the ADC to prevent aliasing of signals outside the range $\{2/3f_{RF}, 2f_{RF}\}$. Undersampling of the RF signal at $4/3f_{RF}$ will cause all spectral blocks spanning a frequency range of $4/3f_{RF}$ to fold into the baseband. We are interested in obtaining a replica of the spectral content in the region $\{2/3f_{RF}, 2f_{RF}\}$. Any content outside this region needs to be attenuated so that the replica remains uncorrupted after sampling.

The expected attenuation characteristics of the pre-selection filter are shown in Figure 2.14. The plot shows the interference locations that alias at $\pm f_{RF}/3$, where the signal replica of f_{RF} eventually forms after undersampling. To ensure that the replica is formed without any corruption, the pre-selection filter must attenuate these interferers. The attenuation in the 20MHz band around the interference locations must exceed the linearity requirement at the MODEM input. Therefore the interferers must be suppressed by > 72dB (12-bit linearity) with respect to the signal. The

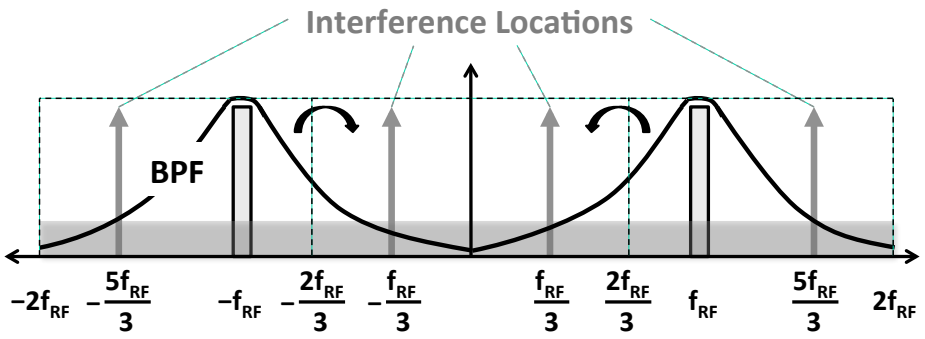

Figure 2.14: Expected attenuation characteristics of the pre-selection filter.

anti-aliasing property of the continuous-time ΣΔ ADC can be used to suppress the interferers at the aliasing locations of $f_{RF}/3$ and $5/3f_{RF}$. The STF, shown in Figure 2.9(c), has notches at these frequencies, and will suppress any interference that can alias with the desired signal centered at f_{RF} after undersampling. Figure 2.15 shows the zoomed-in plots of the notches at $f_{RF}/3$ and $5/3f_{RF}$ for f_{RF} = 2.025GHz. As expected, the STF notches become less sharp with increasing values of R_W/L. At $R_W/L = (4\pi)\times10^7$, we get approximately 100dB of attenuation at $f_{RF}/3$ = 675MHz, and around 80dB at $5/3f_{RF}$ = 3.375GHz. At higher RF carriers, more attenuation is expected since the notches become steeper due to higher oversampling ratios. Therefore, a pre-selection filter for anti-aliasing is not necessary in this system.

The second task of the pre-selection filter is to lower the dynamic range of the received signal in the ADC band of $\{2/3f_{RF}, 2f_{RF}\}$ by suppressing signal content that lies beyond the bandwidth of 20MHz centered around f_{RF}. For this purpose, a pre-selection filter can be implemented using a passive bandpass LC filter. Passive filters are attractive due to their minimal power consumption, but their 3dB cut-off is limited by the Q factor of the inductor, since no Q-enhancement techniques are used. For an on-chip implementation, the Q needs to be minimized. For a 2.025GHz RF modulated signal, an attenuation of 2.2dB can be obtained at a frequency of $f_{RF}-100$MHz = 1.925GHz from a 2^{nd}-order passive LC filter with Q = 10. When combined with the STF of the continuous time ADC, the total attenuation after digitization, is 5.4dB. This attenuation limits the in-band SNR degradation to 3.7dB, when a 0dB blocker is present at $f_{RF}-100$MHz. The value of Q = 10 for the inductor can be realized on-chip so that the system can be fully integrated. The filter can be

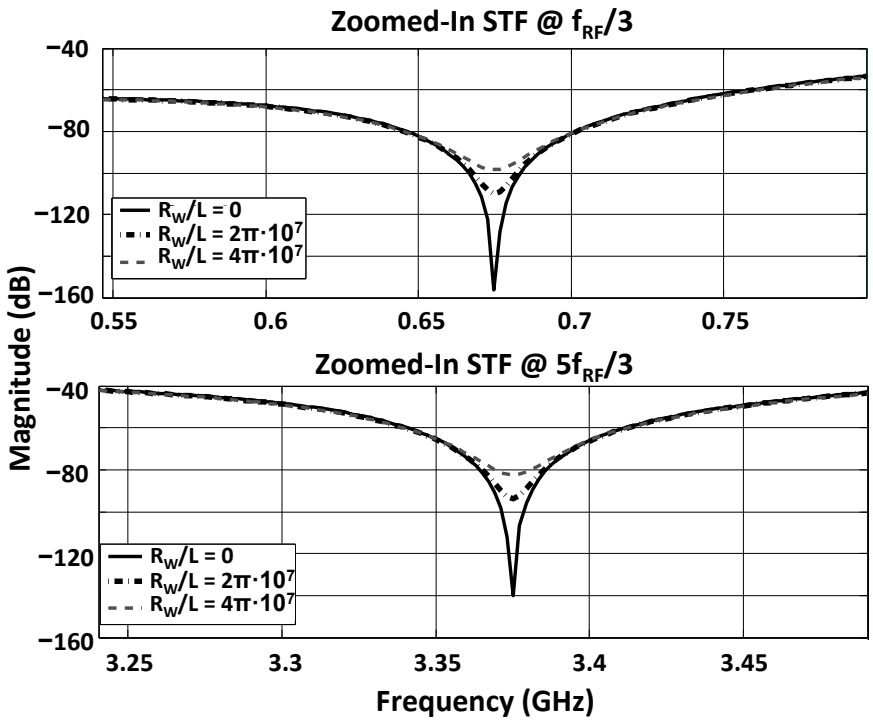

Figure 2.15: Zoomed-in view of notches in the STF at $f_{RF}/3$ and $5f_{RF}/3$.

centered in different frequency bands by the use of a variable capacitor in the LC topology, as was done for the ADC.

After the RF signal is pre-filtered and digitized, it has to be down-converted, filtered, and decimated to the MODEM sample rate. These tasks are performed by a high-speed digital front-end (DFE). Details of the DFE architecture are presented in the next chapter and the results from the chip implementation of the DFE is discussed are Chapter VII.

CHAPTER 3

Digital Front-end for Flexible Receivers

Once the RF signal has been digitized by a high-speed ADC, the digital signal has to be down-converted, decimated, and lowpass filtered before it can be handed off to the baseband modem. This task is performed by the digital front-end. As discussed in Chapter II, the choice of the ADC sample rate of $f_{s1} = 4/3f_{RF}$ will ensure that the mixing operation becomes a trivial multiplication by $\{1, 0, -1, 0\}$. Due to speed and dynamic range limitations in the A/D converter, the maximum RF carrier was restricted to 2.7GHz in Chapter II. Hence, the DFE is designed to process signals modulated at a maximum carrier of 2.7GHz, which sets a maximum sample rate f_{s1} of 3.6GHz. Once the digital data has been downconverted, it has to be decimated from sample rate f_{s1} to the modem sample rate f_{s2}. This decimation has to be preceded by lowpass filtering, without which any noise or interference content in the wide spectrum of f_{s1} will alias and corrupt the signal of interest centered at DC [PM96]. The lowpass filtering and sample rate conversion can be done by a combination of cascade integrated comb filters, polynomial interpolation filters, and finite impulse response (FIR) filters. We look at how each of these filters can be utilized in a flexible manner along the DFE chain. We begin with a description of cascade integrated comb (CIC) filters.

3.1 CIC Decimation Filters

Cascade integrated comb or CIC filters are ideal for decimation by large integer factors and therefore suitable candidates for filtering at the start of the DFE chain where the oversampling ratio is highest. The structure of a typical CIC decimation filter is shown in Figure 3.1(a). The structure is simple to implement requiring only adders, delays, and downsamplers to decimate the input signal at rate f_s to an output sample rate of f_s/D. The CIC filter characteristic is controlled by three tuning

parameters

- N - the number of integrator and differentiator sections, also known as the order of the filter.
- M - the differential delay in the integrator and differentiator loops.
- D - the decimation factor.

The variation of the frequency response with these parameters is shown in Figures 3.1(b)-(d). The response is a sinc function with notches at frequency $(n)\cdot f_s/D$ when D is even and at $(n+1)\cdot f_s/D$ when D is odd and $n \in 1, 2, 3$. Figure 3.1(b) shows the response of a 1^{st}-order CIC (N=1) when the decimation factor D = 4 and the differential delay M = 1. The frequency response has notches at $f_s/4$ and $f_s/2$, which are the locations that fold back and alias with the signal at baseband after decimation. The notches ensure that any signal content in these locations will cause minimum corruption of the signal of interest. The 1^{st}-order response in Figure 3.1(b) provides approximately 25dB of attenuation with respect to DC in the frequency region around the notches. If a strong interference present at locations $f_s/4$ and $f_s/2$ requires more than 25dB of attenuation, then the order of the CIC can be increased to alter its response. Figure 3.1(c) shows the response of a 2^{nd}-order (N = 2) decimation filter. The attenuation in the frequency region around $f_s/4$ and $f_s/2$ is approximately 35dB for this filter, an increase of 10dB compared to the previous case. The order of any higher order CIC can be reduced by bypassing one or more of the cascaded stages in the chain. For example, a 5^{th}-order filter can be easily reconfigured to function as first order if all the integrator and differentiator stages except the first are bypassed.

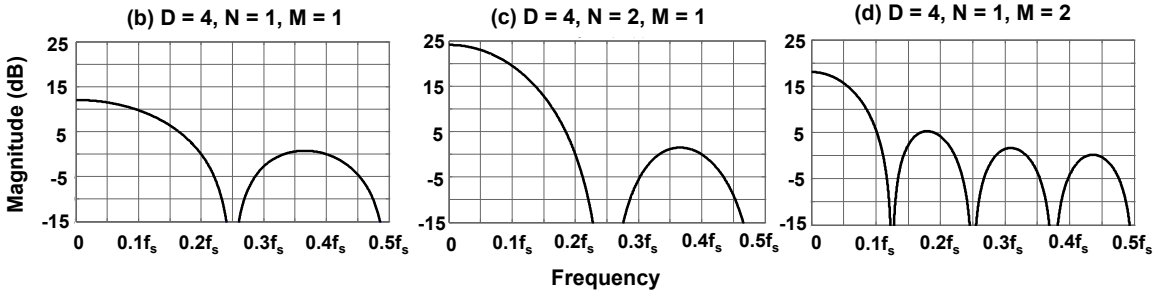

Figure 3.1: (a) CIC architecture. (b)-(d)Response of the CIC at different values of D, N, and M.

The final parameter in the CIC filter response is the number of differential delays, M, in the feedforward and feedback loops of the differentiators and integrators. Figure 3.1(d) shows the frequency response of the 1st-order CIC when the differential delay is increased to 2. The additional delay introduces extra notches in the response at locations $f_s/8$ and $3f_s/8$ without altering the extent of attenuation at the notches (25dB). The differential delay can hence be used to introduce additional nulls if strong interference appears in these strategic locations.

The first CIC stage in the DFE has to support the highest sample rate, a maximum of 3.6Gs/s. This filter becomes a throughput bottleneck if conventional CIC structures are used. Traditionally, CIC filters are implemented using the recursive structure, shown in Figure 3.1(a). The main problem with this architecture is the throughput restriction on f_s imposed by the integrators to the left. Owing to their recursive nature, these integrators cannot be pipelined or parallelized. Moreover, the adders in the integrator loops have long wordlengths (10-14 bits) [KYW98] that make the operation at GHz rates impossible without architectural changes. Several signal-processing techniques were used to optimize the CIC architecture in our implementation. First, although the CIC integrator loops indicate an apparent recursive nature, the transfer function of the filter can be represented as a feed-forward function. An example the transfer function, $H_{CIC}(z)$, for a 2nd-order filter, is given by Eq. 3.1.

$$H_{CIC}(z) = \frac{(1-z^{-D})^2}{(1-z^{-1})^2} = (1+z^{-1})^2(1+z^{-2})^2\ldots(1+z^{-2^{P-1}})^2 \tag{3.1}$$

The modified transfer function can be realized by the filter structure shown in Figure 3.2. We use the relation that z^{-M} followed by decimation by k is equivalent to decimation by k, followed by $z^{(-M/\mathbf{k})}$. The P (= $\log_2$(D)) product terms in Eq. 3.1 have the $(1+z^{-2^i})^2$ form that can be transformed to decimation by 2^i followed by $(1+z^{-1})^2$. This transformation is valid for cases where D = a^x (a, x are integers) and any number of sections.

There are several advantages to using the structure in Figure 3.2 for our application. First, the function $(1+z^{-1})^2$ requires only adders and hardwired shifts, resulting in a low-complexity implementation. Second, the wordlengths of the adders increase gradually along the chain (6-14 bits here) due to the feed-forward nature of the design. Finally, the feed-forward filter can now

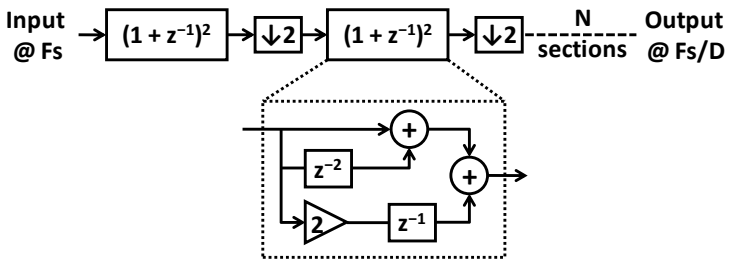

Figure 3.2: Throughput-optimized feedforward CIC architecture for a 2nd-order filter.

be parallelized to support the throughput requirement of the CIC. Parallelizing the output of $(1 + z^{-1})^2$ is advantageous, since it is followed by decimation-by-2; implying only alternate outputs are required by the next stage, and one of the output channels can be removed, resulting in greater computational efficiency. This version of the 2nd-order optimized CIC is used in the first stage of decimation of the receiver DFE. Next we look at reconfigurable FIR filters that are utilized for integer decimation as well.

3.2 Reconfigurable FIR Filters

An FIR filter is a versatile DSP block where several categories of frequency responses can be generated by varying the filter tap coefficients. For a flexible radio receiver system, we are interested in generating lowpass filters with varying passband, rolloff, and stopband attenuation characteristics which can be accomplished by configuring the FIR. The block diagram of a reconfigurable FIR decimation filter is shown in Figure 3.3. The filter can be tuned to vary its frequency response by adjusting the value of its tap coefficients $a_0, \ldots, a_{p-1}$. Zeroing out a subset of the coefficients starting from the last tap, a_{p-1}, reduces the order of the filter. This provides an opportunity for obtaining runtime power performance trade-offs with varying SNR conditions. In conditions of high SNR and weak interference the filter can be tuned to operate with minimal number of taps.

An example of this scenario is shown in the frequency responses of Figure 3.4. The two filters were designed for a decimation factor D = 2, hence the stopband of both responses begins at 0.5 with the rolloff beginning at 0.3. The passband ripple is also same for both filters, set at 0.1dB. The only difference is in the stopband attenuation; the 34-tap filter has a stopband attenuation of

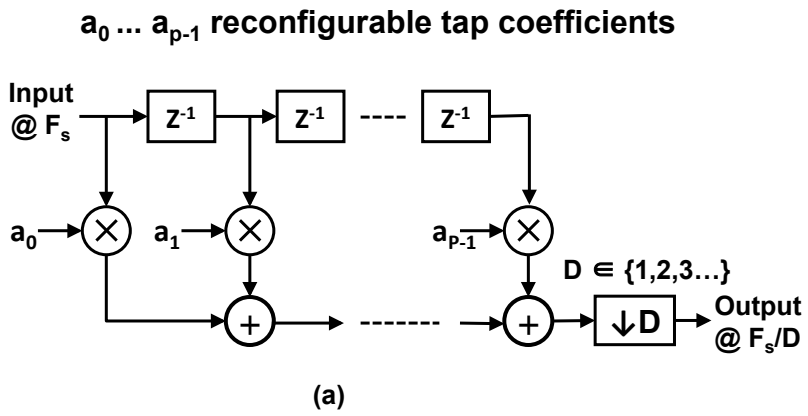

Figure 3.3: Reconfigurable FIR filter signal-flow-graph.

80dB while the 28-tap filter has an attenuation of 60dB. The frequency response can be controlled to vary the start of the stopband depending upon the decimation factor D that follows the FIR. For decimation factor D = 3, for example, the stopband would have to start at 0.33 instead of 0.5. The tap coefficients and filter order are chosen such that:

- The passband is nearly flat in the region occupied by the baseband signal.
- The filter cutoff frequency is greater than the location of the highest frequency content of the signal.

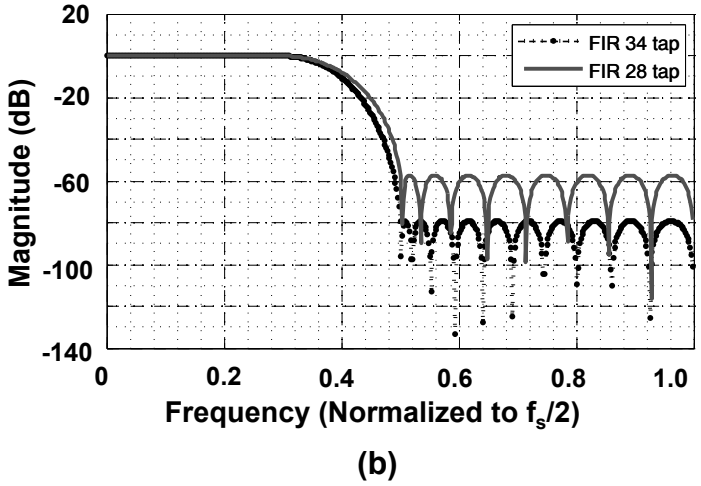

Figure 3.4: Frequency response of the FIR with variable tap coefficients.

- The stopband suppresses any noise or interference that aliases with the signal after decimation by at least 70dB with respect to the signal. The choice of 70dB was made in order to maintain 70dB SNR at the baseband, when this filter is the last stage of the DFE.

The reconfigurable FIR has distinct advantages over the CIC filter. We have finer control over the response of the FIR because all its tap coefficients can be varied. This ensures that the passband of the FIR can be designed to avoid the droop that we observed in the sinc response of the CIC. The penalty lies in the increased complexity of the FIR that needs multiplier units for its signal processing. Hence to minimize switching power consumption it is advantageous to place the FIR further down in the DFE chain where it can operate at lower sample rates.

Since the ADC sampling frequency f_{s1} is a function of f_{RF}, and the modem frequency f_{s2} is dictated by the target standard and signal bandwidth, f_{s1} and f_{s2} in general are not related by integer factors. The DFE, therefore, has to provide for flexible integer as well as fractional decimation. The integer decimation can be realized using a combination of CICs and FIRs, but the fractional decimation needs special interpolation filters that are discussed next.

3.3 Fractional Sample Rate Conversion

In a fully flexible receiver we need fractional sample rate conversion to convert the sample rate from a decimated version of f_{s1}, given by f_{s1}/K_1, to a multiple of f_{s2}, $K_2 f_{s2}$. We define the interpolation ratio IR as:

$$IR = \frac{f_{s1}}{K_2 K_1 f_{s2}} \tag{3.2}$$

The simplest way to perform this interpolation when IR is not an integer, is to express IR as the ratio p/q, upsample the data by q followed by downsampling it by p. The process is shown in Figure 3.5(a) for $p = 100$, $q = 101$, and $IR = 1.01$. But this method tends to use a high intermediate frequency, making it power hungry for large values of q. For example, an input signal sampled at

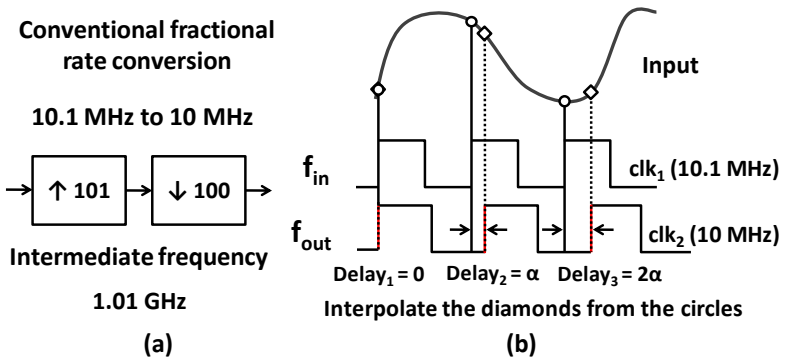

Figure 3.5: (a) Conventional fractional sample rate conversion (b) Digital Interpolation

f_{in} =10.1MHz has to be upsampled to 1.01GHz in the example shown in Figure 3.5(a). To avoid this problem, we can use digital interpolation techniques to re-sample the incoming signal; this concept is illustrated in Figure 3.5(b). The interpolator processes the incoming input samples at f_{in} (circles) to compute the output samples at f_{out} (diamonds). To implement the digital interpolation approach, the first step is to converge on an algorithm that finds reasonably accurate estimates for the output samples (diamonds) from the input samples.

It is useful here to consider the effect of fractional sample rate conversion in the frequency domain. This will give us insight into the various ways the interpolation can be accomplished. We take the example of a baseband signal sampled at f_{in}, which has to be re-sampled at f_{out}. Figure 3.6(a) shows the power spectral density of the original band-limited signal at f_{in}, which repeats at multiples of f_{in}. Assuming that the input $x(n)$ is 1 bit wide, the structure in Figure 3.6(b) shows a simple approach to interpolation that does not involve any mathematical computation. The original 1-bit digital signal $x(n)$, sampled at f_{in}, is captured by a register clocked at f_{out}, to generate the output $y(n)$. We can consider the output of the first register to be an analog signal $x_{ZOH}(t)$ that is generated by a zero-order-hold (ZOH) operation on the sampled data $x(n)$. The ZOH operation can be treated as a reconstruction filter that estimates the interpolated sample value by looking at the last sample value of the original data. The spectrum of the analog signal $x_{ZOH}(t)$ can be obtained by applying the ZOH sinc response to the digital spectrum of $x(n)$, as shown in Figure 3.6(a). The analog signal $x_{ZOH}(t)$ is then directly sampled by the second register at intervals of $T_2 = \frac{1}{f_{out}}$. The spectrum of $y(n)$ can be obtained by convolving the analog spectrum of $x_{ZOH}(t)$

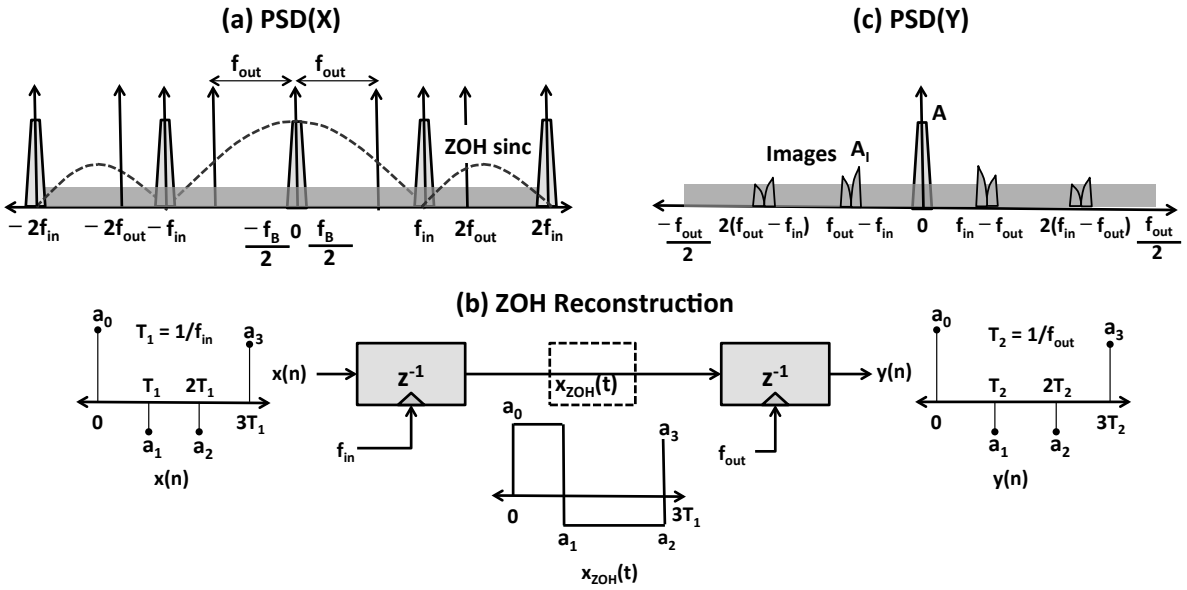

Figure 3.6: (a) Original baseband spectrum (b) Interpolation with ZOH as reconstruction filter (c) Attenuated spectral images after interpolation

with a frequency-domain impulse train spaced at f_{out} [OS89].

The resultant spectrum $PSD(Y)$ is shown in Figure 3.6(c). We find that the images at multiples of f_{in}, attenuated by the notches in the sinc response, appear in the baseband spectrum between $f_{out}/2$ and $-f_{out}/2$. Ideally, the magnitude A_I of these images should be below the permitted noise level at the MODEM input, determined by the SNR requirements of the target standard and also any additional filtering stages present after this interpolation stage. The magnitude A_I is mainly dependent upon the oversampling ratio of the baseband data, given by $P = f_{in}/f_B$. As the value of P increases, the attenuation provided by the sinc response for the images also increases. Hence such an interpolator will have to be placed very early on in the DFE chain where the oversampling ratio is high. For example, for a signal bandwidth of 20MHz, and an RF carrier of f_{RF} = 2GHz, the oversampling ratio right after A/D conversion and mixing is $P = 2/3f_{RF}/20e6 = 66.66$. This value of P provides about 45dB of attenuation for the first image.

Images at attenuation level of 45dB will degrade the SNR of the baseband signal and may even corrupt the in-band signal (in later stages of integer decimation) if not suppressed by the next stage of filtering. Hence, such an interpolator will need an FIR in the next stage. This approach is shown in Figure 3.7, where $P = f_{in}/f_B = 15$, and the image I is 30dB below the baseband signal. The

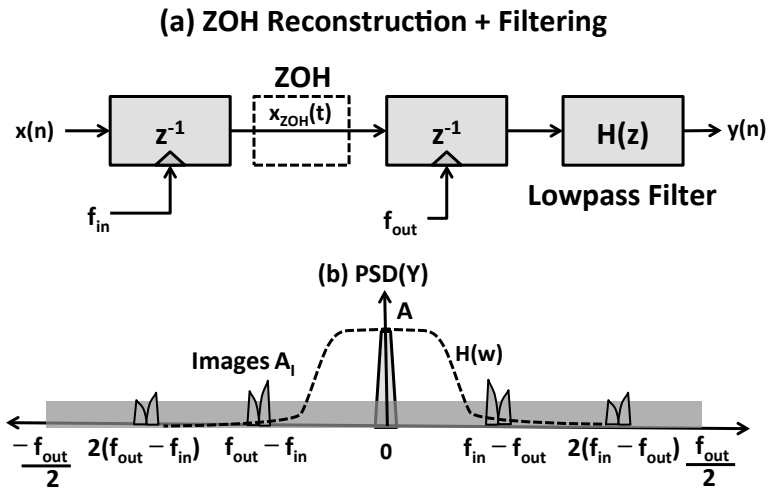

Figure 3.7: (a) ZOH Reconstruction followed by lowpass filter (b) Output spectrum

images can be lowered by an additional 15dB using a lowpass FIR after interpolation to frequency f_{out} as shown in Figure 3.7. A problem with this approach lies in the level of flexibility that can be allowed in the position of the images with respect to the baseband signal. For example, in Figure 3.7(b) the image position of $f_{out} - f_{in}$ is approximately halfway between DC and $f_{out}/2$. Hence an FIR with stop band edge starting at $f_{out}/4$ would suffice to attenuate the images. For a different oversampling ratio, P, and a different set of f_{out} and f_{in}, the position of the image can be closer to DC, for example at $0.3f_{out}/2$. In this case, the FIR stop band attenuation will have to start at $0.3f_{out}/2$. In this situation we would need a reconfigurable FIR filter, where the frequency response can be varied, to support both scenarios. This can be expensive to implement because not only will this require the use of general-purpose multipliers and stored coefficient values, but the FIR will operate at very high sample rates due to the need of high oversampling ratios for ZOH interpolation. This method also has a second disadvantage when it comes to generating the clock f_{out}. Since f_{out} will have to be a multiple of the baseband sample rate f_{s2}, we need to generate a very large multiple of f_{s2} in order to support the high oversampling ratio.

3.3.1 Polynomial Interpolation

The problems encountered in ZOH interpolation can be solved to some extent by polynomial interpolators that can accomplish fractional sample-rate conversion. The concept is illustrated in

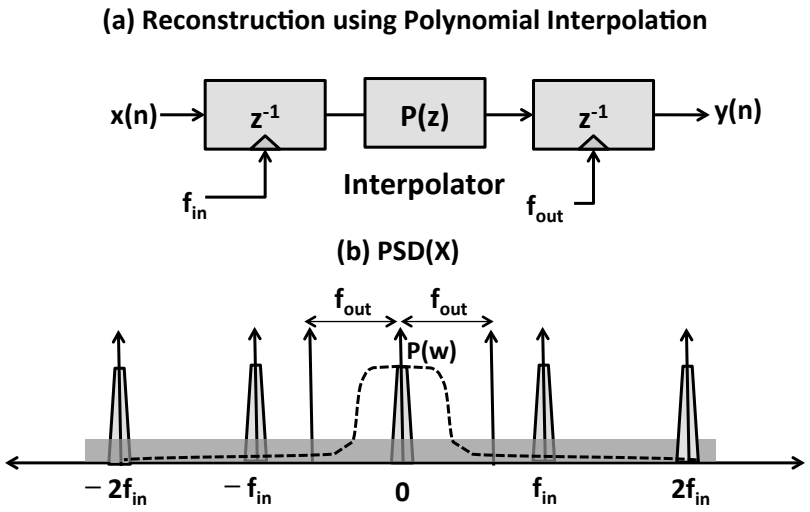

Figure 3.8: (a) Interpolation using polynomial interpolator, $P(z)$, for reconstruction (b) Attenuated spectral images after interpolation

Figure 3.8. In this approach we resort to explicit calculations of the output sample values at f_{out}, given the input samples at f_{in}. The calculations are based on standard interpolation approaches like linear interpolation, quadratic interpolation, or cubic-spline interpolation. The effect of the interpolator response, $P(w)$, will be to attenuate the spectral images at multiples of f_{in} in the frequency domain, as shown in Figure 3.8(b). In this work we use cubic polynomial interpolation since it attenuates the images by > 50dB. We use a low-speed FIR after the interpolator to further attenuate the spectral images. This method does not require a large oversampling ratio to ensure the 50dB+ image attenuation. The polynomial interpolation is implemented using the Farrow structure

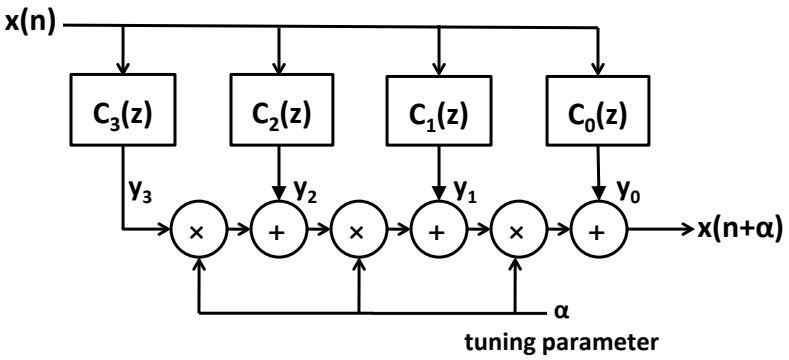

Figure 3.9: Farrow filter used for third order polynomial interpolation.

described in [Far88]. The block diagram of the Farrow filter is shown in Figure 3.9.

The interpolator transfer function, $P(z)$, implemented by the Farrow filter is given by Eq. 3.3. The polynomial $P(z)$ is in fact a representation of the Taylor's series interpolation in the z-domain.

$$P(z) = C_0(z) + \alpha C_1(z) + \alpha^2 C_2(z) + \alpha^3 C_3(z) \tag{3.3}$$

The third-order Taylor's series polynomial is expressed as in Eq. 3.4.

$$f(t+\alpha) = f(t) + \alpha\frac{f^{'}(t)}{1!} + \alpha\frac{f^{''}(t)}{2!} + \alpha\frac{f^{'''}(t)}{3!} \tag{3.4}$$

The transfer functions $C_1(z)$, $C_2(z)$, and $C_3(z)$ in Eq. 3.3 are FIR filters that approximate the first, second, and third order differentials $f^{'}(t)$, $f^{''}(t)$, and $f^{'''}(t)$. The function $C_0(z)$ is an allpass filter. All the FIRs in Eq. 3.3 were realized using 8-tap filters. The frequency response of the functions $C_0(z)$, $C_1(z)$, $C_2(z)$, and $C_3(z)$ are shown in Figure 3.10. The dashed line shows the response of the ideal all pass and differential functions, and the polynomials designed approximate

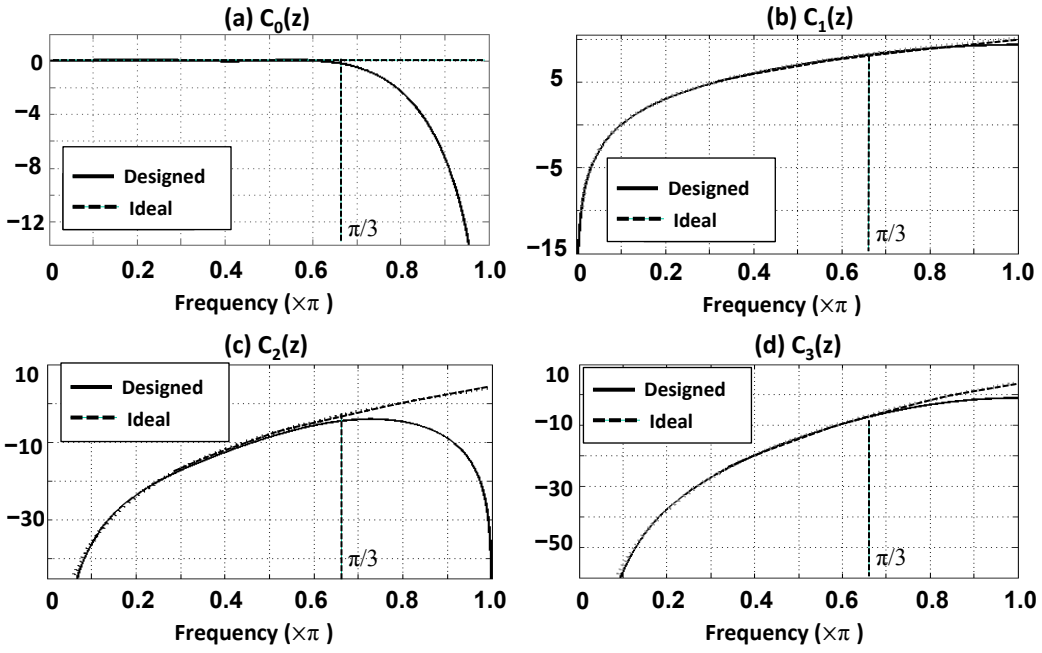

Figure 3.10: Frequency response of Farrow polynomials. The dashed line shows behavior of ideal allpass and differential functions.

their ideal behavior up to frequency $2\pi/3$. This is an adequate approximation, if we ensure that the signal of interest is band limited between $\pm 2\pi/3$. In order to fully approximate the differentials $f^{'}(t)$, $f^{''}(t)$, and $f^{'''}(t)$ over the entire bandwidth, we would need infinitely large number of taps in the FIRs, which becomes infeasible from an implementation perspective.

The symbol α in Eq. 3.3 represents the time delay between the input and the output sampling clock edges. This delay is computed in real time using the initial phase of both clocks and the difference between their respective time periods. Since the time period of the clocks is a known quantity, the user sets this difference value externally through a 24-bit word, sent to the system via a scan chain.

Erroneous samples appear at the interpolator output when α does not match the exact time delay between the two clocks. Miscalculation of α and subsequent sampling errors can be avoided if the initial phase difference between the two clocks is measured. Hence, a synchronization circuit, which facilitates relative phase calibration between the two clocks has to be added to the system. The concept of phase calibration is shown in Figure 3.11. The value V toggles at the positive edges of the output clock, f_{out}, and is latched at positive edges of the input clock, f_{in}. If we decimate, then the output clock period T_{out} is longer than T_{in}, and at some instant consecutive positive edges of T_{in} will lie within one period of T_{out}. When this happens, the value V latched by f_{out} fails to change, shown by consecutive V = 0 in Figure 3.11, and we know that the positive edges of both clocks are separated by less than $T_{out} - T_{in}$. The interpolator uses this information for phase difference calibration. For cases where $T_{out} < T_{in}$, the same approach can be followed by flipping f_{out} and f_{in}.

Hold time violations may occur when value V, toggling at f_{out}, is latched at positive edges of f_{in}. 2-stage synchronizers and circuit redundancy is used to minimize the probability of such

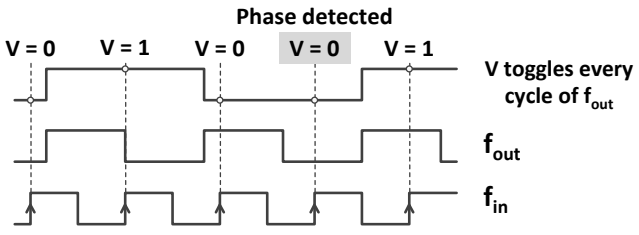

Figure 3.11: (a) Phase calibration between input and output clocks.

errors in the system. 2-stage synchronizers are a pair of cascaded flip-flops, where the first flip-flop captures the incoming asynchronous input at the output clock, f_{out}. Hold time violations are likely to occur at this interface and the data at the output of the first flip-flop is not reliable. Hence if other parts of the system tap this data, we are likely to end with errors across the system. However, if we allow one cycle of latency and tap the output at the second flip-flop, then we allow the data at the output of the first flip-flop to stabilize before it is captured by the second flip-flop in the next cycle. This scheme reduces the extent of errors due to timing violations.

The frequency response of the Farrow filter is shown in Figure 3.12 for f_{in} = 32MHz and f_{out} = 26.67MHz. The frequency response in Figure 3.12(a) was computed at f_{in} = 32MHz and an interpolation ratio of 5/6. The attenuation characteristics of the Farrow is easier visualized if we look at the response of the equivalent upsample-filter-downsample model shown in Figure 3.12(b). This model realizes the decimation factor of 5/6 by first up sampling the data by a factor of 5 to 160MHz, and then downsampling it by 6 to 26.67MHz. Although we still use the structure in Figure 3.9 for implementation purposes, the model in Figure 3.12(b) tells us how the interpolator suppresses the images that would fold into the baseband after the sample rate conversion. From the

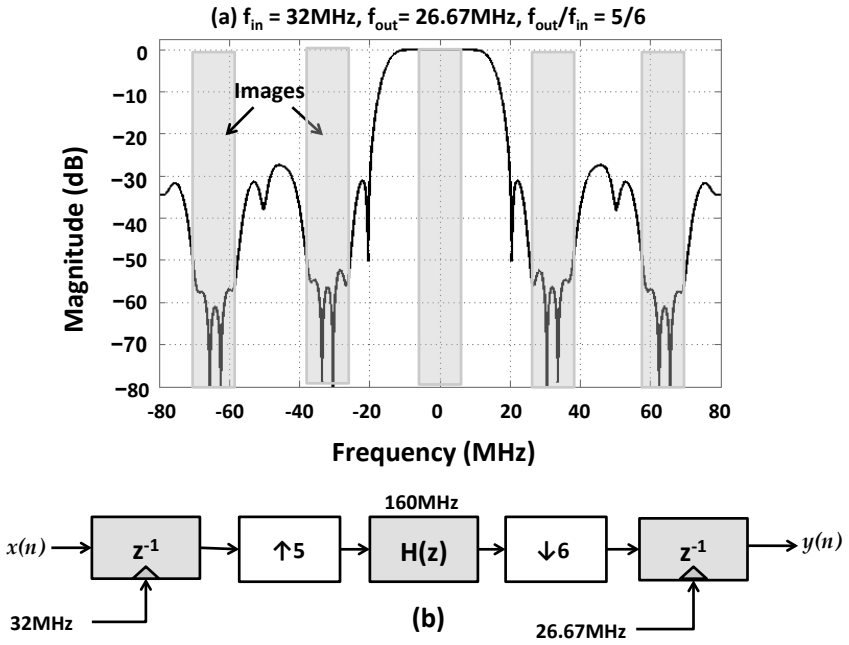

Figure 3.12: (a) Farrow frequency response at f_{in} = 32MHz and f_{out}=26.67MHz. (b) Equivalent upsample-filter-downsample model

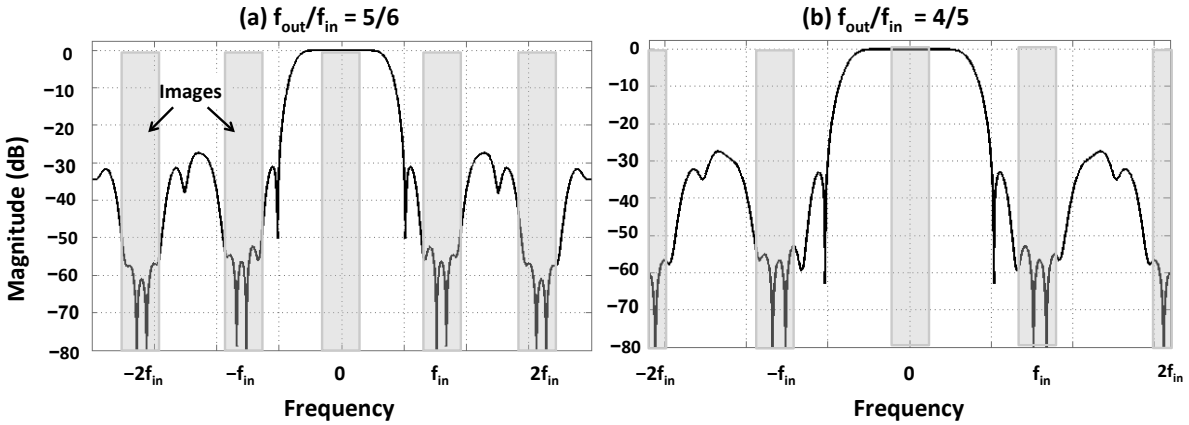

Figure 3.13: (a) Farrow frequency response at f_{out}/f_{in}=5/6. (b) Response at f_{out}/f_{in}=4/5.

figure, we can see that the images are attenuated by > 50dB. The advantage of using the polynomial interpolator over a reconfigurable FIR, lies in the fact that the frequency response and the sample rate conversion factor can be changed in real time by tuning the parameter α in Figure 3.9. For example, Figures 3.13(a) and (b) show the frequency response of the Farrow rate converter for two different conversion factors. In Figure 3.13(a), the conversion factor is 5/6, and the response spans a frequency range of $5f_{in}$, suppressing the images at multiples of f_{in}. In Figure 3.13(b), the response spans a range of $4f_{in}$ for a rate conversion factor of 4/5. Hence the Farrow is ideally suited for implementation in the flexible radio chains.

One of the limitations of the Farrow filter is the insufficient suppression of images w.r.t. to the SNR requirements at the MODEM input. For example, Figure 3.14(b), shows the result of interpolation by factor 4/5, where the dominant image is attenuated by approximately 66dB. But for SNR > 72dB we need higher suppression. This problem can be resolved by strategically configuring the frequency response of the next component in the DFE chain such that A_I can be suppressed to the required attenuation level. As previously discussed in this section, the dominant image is centered at $f_{in} - f_{out}$. We place an FIR after the Farrow to ensure that the the attenuation at the image location is at least 6dB to obtain 72dB suppression. In order to avoid over-constraining the FIR, it is useful if the image location is not too close to the signal of interest. For this reason we impose the additional constraint of $f_{in} - f_{out} \geq f_{out}/4$ to ensure that the dominant image is centered at locations $f_{out}/4$ and beyond. Although this approach constrains the choice of the fractional sam-

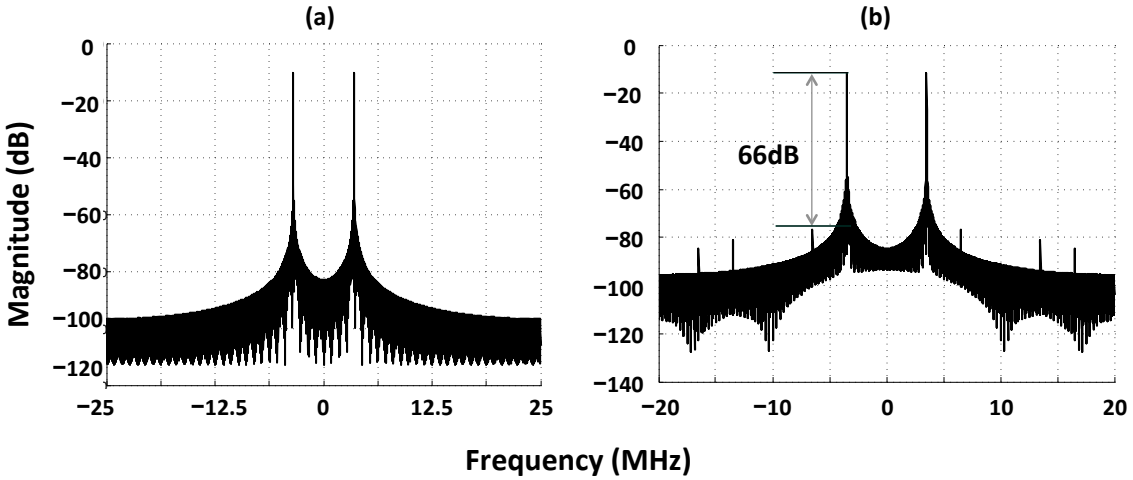

Figure 3.14: (a) Sinusoidal input to Farrow interpolator. (b) Farrow output after interpolation at rate 4/5.

ple rate conversion factor to some extent, in most cases it is possible to maintain this relation by effectively tuning the Farrow and the CICs preceding it. Now that we have looked at the major components that are part of the DFE chain, we move on to the system-level view of the flexible receiver architecture.

3.4 DFE Architecture

The proposed Rx DFE architecture is shown in Figure 3.15 [NCM11]. As discussed in Chapter II, the incoming RF signal centered at f_{RF} is under-sampled by the ADC at a frequency of $4/3f_{RF}$. The DFE is designed to process RF carriers up to 2.7GHz, which sets a maximum sample rate f_{s1} of 3.6GHz. After mixing, the down-converted signal has to be decimated to the modem frequency

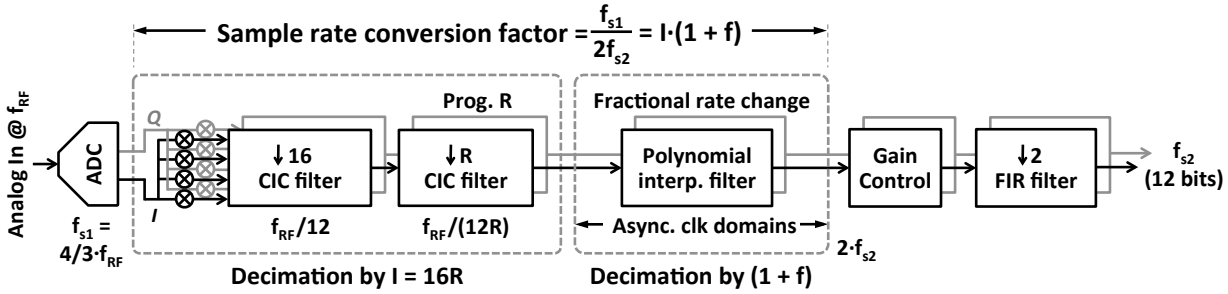

Figure 3.15: Receiver Digital Front-end Architecture.

f_{s2}. Since the decimation ratio f_{s1}/f_{s2} in general is not an integer, the DFE provides for flexible integer as well as fractional decimation. A combination of CIC, FIR, and polynomial interpolation performs decimation and lowpass filtering. The process is divided into integer (I) and fractional (1 + f) decimation (Figure 3.15).

The integer decimation process is split into two phases. First, a constant decimation by 16 is realized using a CIC filter. This filter brings the sample rate of the input signal from $4/3f_{RF}$ down to $f_{RF}/12$. A big challenge in implementing this filter is handling the GHz-input rate of $4/3f_{RF}$. The feed-forward parallel CIC architecture, discussed in Section 3.1, is used to make this high throughput feasible. The choice of decimation by 16 in the first stage was made in order to lower the sample rate enough so as to implement the subsequent programmable CIC with the traditional recursive structure. The next block in the DFE chain is a programmable integer (R) decimation filter that uses the recursive CIC architecture. This is a 5^{th}-order filter that ensures high noise and interference attenuation. The user can set the value of the integer decimation factor, varying from 1 to 64. The decimation factor is given by $R = \text{floor}(f_{RF}/(24f_{s2}))$.

The fractional decimation (1 + f) uses the 3^{rd}-order polynomial interpolation filter, discussed in Section 3.3. This filter must hand-off data between two asynchronous clock domains, $f_{RF}/(12R)$ and $2f_{s2}$. Two-stage synchronizers and redundancy are used to achieve this hand-off [RCN04]. A programmable gain-control block follows the fractional decimation. The final decimation-by-2, which brings the sample rate to f_{s2}, is done using reconfigurable FIR filters with 12-bit output providing attenuation of around 70dB to clear out adjacent interferers.

3.5 Simulation Results

The DFE architecture discussed in the previous section was programmed to receive an RF signal centered at 2.025GHz. The RF signal was undersampled and digitized at 2.7GHz ($4/3f_{RF}$). The spectrum after undersampling and ΣΔ modulation is shown in Figure 3.16(a) where the replica of the signal is formed at $1/3f_{RF}$ = 675MHz. The signal of interest is the sinusoid centered at 667MHz and the received signal has an interfering sinusoid at location 772MHz. The DFE has to decimate this signal to a baseband sample rate of 20MHz while also attenuating the interfering signal. The

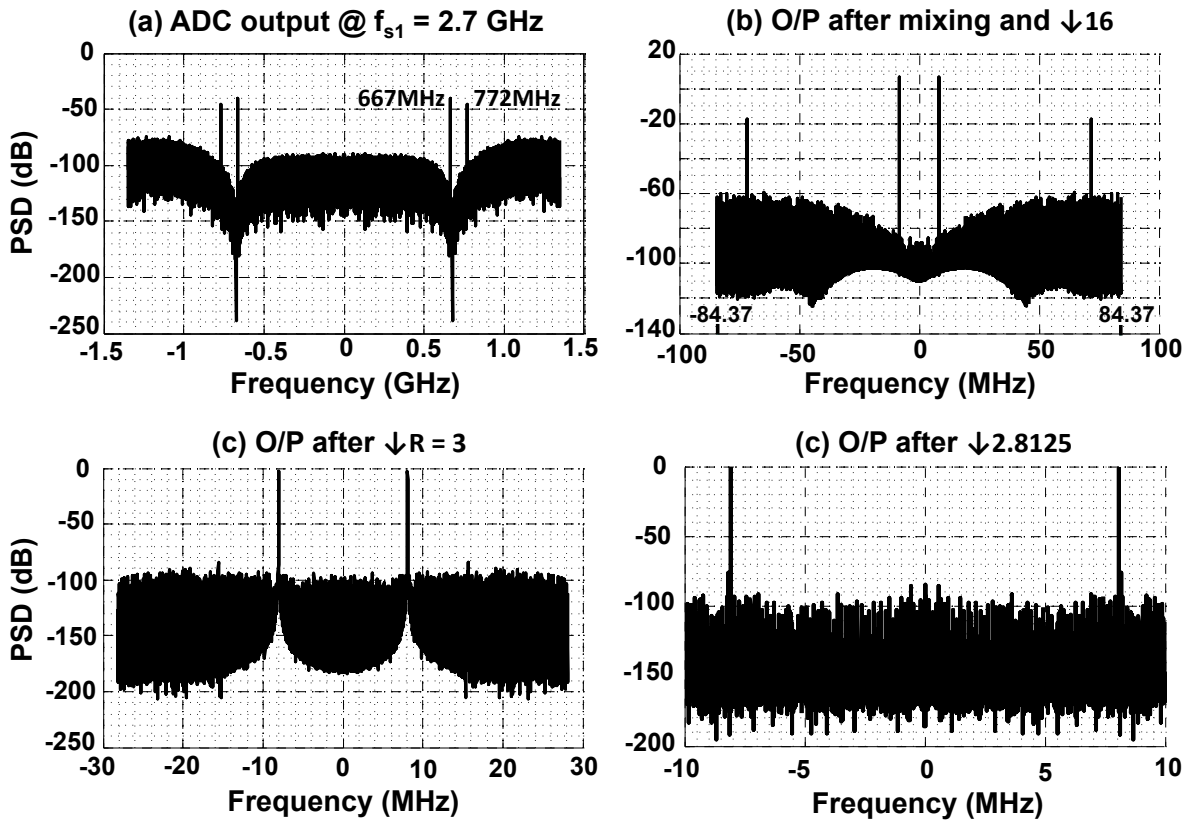

Figure 3.16: Spectrum of output at different stages of the receiver DFE.

spectrum of the signal after mixing and CIC decimation by 16 is shown in Figure 3.16(b). The mixing process brings the RF modulated signal to the baseband and the CIC decimation lowers the sample rate to 168.75MHz. The noise content in the input spectrum of 2.7GHz cannot be fully eliminated by lowpass filtering, which results in finite noise injection in the signal bandwidth after decimation. The decimation-by-16 CIC attenuates the interfering sinusoid by about 5dB.

The next block in the DFE chain is the programmable CIC decimator, which was tuned to a decimation factor of R = 3 in this configuration. This decimation brings the sample rate down to 56.25MHz. A choice of R = 2 would have resulted in a sample rate of 84.375MHz but this would mean that the polynomial interpolation filter, which is a relatively complex unit with a series of multiply-add operations, would have to function at a higher sample rate. This would result in increased power consumption and timing closure issues in the interpolator. Decimation by R = 4, on the other hand, would result in an output sample rate of 42.1875, which is only

about 3 times greater than the signal bandwidth and leads to unwanted signal attenuation due to the sinc response of the CIC. Hence the optimum choice for the decimation factor was R = 3 in this case. The resultant spectrum of the signal after decimation is shown in Figure 3.16(c). Polynomial interpolation reduces the sample rate by a factor of 1.40625 and is followed by a 40-tap reconfigurablee FIR that is configured for decimation by a factor of 2. The interfering signal was attenuated by the combination of programmable CIC and FIR filters. The output spectrum, which has an SNR of 73dB, is shown in Figure 3.16(d).

The DFE can be similarly programmed for other combinations of RF carriers, modem sample rates, and interference conditions. In scenarios where we have multiple options for the tuning parameters, the final configuration is the one that ensures the best power-performance trade-off. This will correspond to the configuration that operates the DFE components at the lowest switching power while also ensuring that the noise injection in the signal bandwidth remains below the target constraints.

Now that we have studied the architecture of flexible receiver front-ends, we will discuss the options of a corresponding transmitter front-end that is also implemented in a digitally-intensive manner.

CHAPTER 4

Digital Front-end for Flexible Transmitters

In Chapter I we looked at the options available for implementing an RF transmitter that modulates the baseband digital data to the RF carrier frequency. These were the outphasing, polar, and linear modulation schemes. With the intention of maximizing the signal bandwidth and supporting complex modulation schemes like 64-QAM that enable high data rate communication, we chose the linear modulator for our implementation.

The basic signal processing flow in a linear modulator is shown in Figure 4.1. The baseband signal from the modem sampled at rate f_{s2} is upsampled by the Tx digital front-end (DFE) to a higher sample rate f_{s1}. The DFE filters the images of the baseband signal occurring at multiples of f_{s2} to ensure that the images do not violate the emission mask requirements of the target standard. ΣΔ modulation is used after upsampling to lower the number of digital data bits input to the RFDAC. The ΣΔ modulator shapes the added quantization noise away from the baseband signal. The RFDAC is responsible for D/A conversion of the digital bits and mixing the baseband data with the carrier frequency, f_{RF}. A bandpass filter (BPF) is needed before the linear PA to suppress any noise content in the RF signal that violates the output emission mask or interferes with the noise

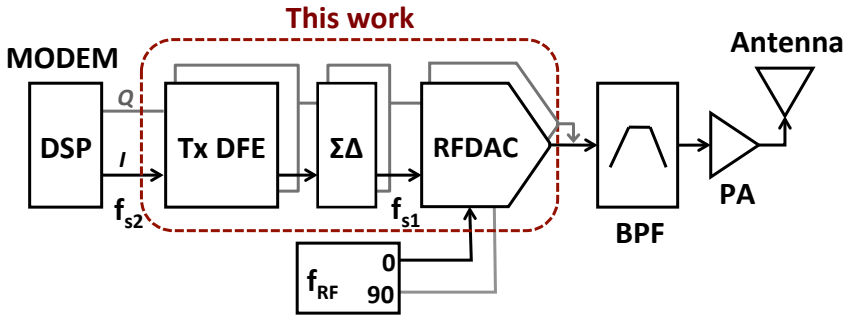

Figure 4.1: Linear RF modulator with a digital front-end (DFE), ΣΔ modulator, and RFDAC.

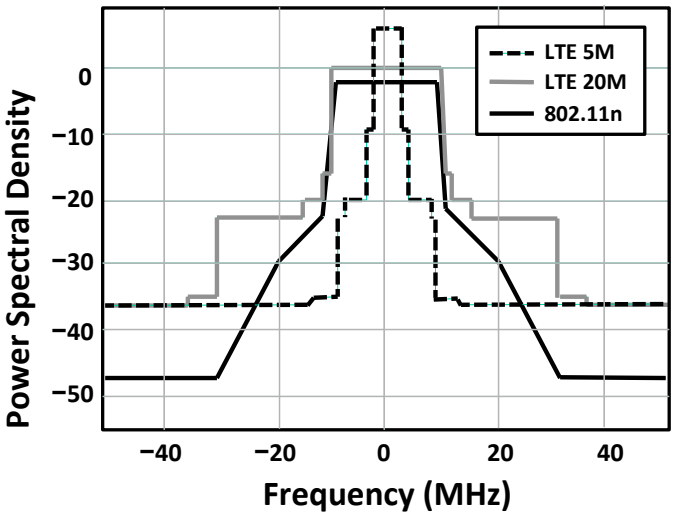

Figure 4.2: FCC emission mask for various cellular and WLAN standards.

level in the receive band in frequency division duplexing (FDD) schemes.

In this chapter we discuss the design of the digital components in the Tx chain. This involves a study of the digital front-end (DFE) and the ΣΔ modulator architecture. A major criticism of linear modulators springs from the high power consumption [PZH08] in the digital blocks that lowers the overall energy efficiency of the system. This is because for ΣΔ modulation to minimize the noise floor at low modulation orders (< 3), the sample rate f_{s1} has to be very high (5.4Gs/s in [PZH08] for a 20MHz signal) or the resolution of the D/A converters has to be large. For straightforward linear modulators without ΣΔ modulation, the resolution of the data converter is between 10-12 bits [ESK07], depending upon the wordlength at the output of the MODEM. The resolution for polar modulators is even higher due to more complex signal processing, and can go up to 17 bits [BMK11]. The power consumption in the digital components increases linearly with the sampling frequency, leading to overall higher power consumption in the Tx chain. In this work we optimize the ΣΔ modulator architecture by increasing its order, N, while simultaneously reducing the sample rate, f_{s1}, and the number of quantization levels at the output. This achieves an overall lower power implementation as well as low resolution in the RFDAC. The optimization techniques used will modulate signals from arbitrary baseband rates of f_{s2} to RF carriers up to 2.7GHz.

In order to compute the emission mask requirements, we look at the FCC masks for various

	WCDMA	802.11n	LTE (5M)	LTE (20M)
Channel Bandwidth	5MHz	20MHz	5MHz	20MHz
Modulation	HPSK	64QAM OFDM	4/16/64 QAM SC-FDMA	4/16/64 SC-FDMA
Peak-to-average ratio	3.5dB	10dB	2dB-8dB	2dB-8dB
ACPR	−33dBc @5MHz	−20dBc @20MHz	−27dBc @5MHz	−20dBc @20MHz
Emission Mask	−43dBc @7.5MHz	−45dBc @30MHz	−42dBc @12.5MHz	−36dBc @30MHz

Table 4.1: Specifications for multi-standard radio transmitters.

cellular and WLAN (802.11n) standards in Figure 4.2. The strictest requirement comes from the WLAN standard that requires 20dBc (relative to PSD level at RF carrier) in the adjacent channel and 45dBc in the alternate channel for a 20MHz bandwidth signal and 64QAM OFDM signal (peak-to-average ratio of 10dB). If the system is able to satisfy this requirement in the adjacent and alternate channels, then it can handle all other specifications. A table with detailed specifications of emission mask, adjacent channel power ratio (ACPR), modulation schemes, and signal bandwidths for representative standards are listed in Table 4.1. The peak-to-average power ratio (PAPR) affects maximum power delivered at the output of the transmitter for a fixed amount of current at the output of the DAC. The higher the value of PAPR, the lower is the power delivered at the output. This requirement is also toughest for the WLAN 64QAM OFDM signal. In the subsequent sections we look at the design of a modulator that can satisfy these specifications while maximizing the extent of DSP processing in the transmit chain.

4.1 Transmit Digital Front-end

A digital front-end interpolates the baseband digital signal, sampled at rate f_{s2}, to the RFDAC sampling frequency, f_{s1}. The sample rate f_{s2} is governed by the target radio standard and signal bandwidth. The RFDAC sample rate, f_{s1}, is a design variable that fundamentally affects the power and spectral performance of the modulator. We will discuss the choice of f_{s1} in detail in a later section of this chapter. For now, it is assumed that f_{s1} can lie in the range of several hundred MHz to a few GHz. Before we dive into the details of the DFE micro-architecture, it is useful to look at

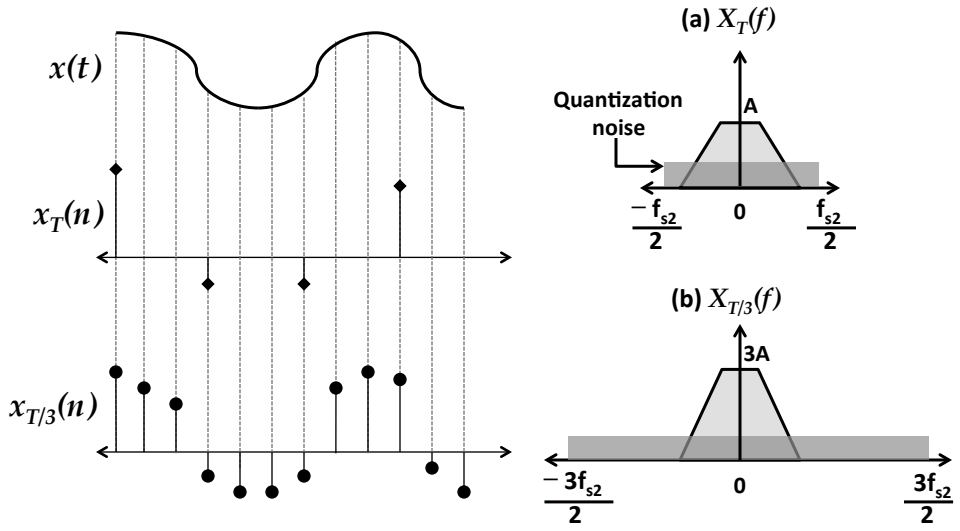

Figure 4.3: (a) $X_T(f)$ at sample rate f_{s2} (b) $X_{T/3}(f)$ at sample rate $3f_{s2}$.

the basics of the digital interpolation process.

4.1.1 Digital Interpolation

Let us take the example of a baseband signal, $x_T(n)$, which is sampled at 5MHz and has to be interpolated to 15MHz. This would correspond to an interpolation factor of 3. Let us assume that the original sequence $x(n)$ was obtained by sampling an analog signal $x(t)$ at a sample period of T. Also, let us assume that there exists a sequence $x_{T/3}(n)$ that is the result of sampling the input $x(t)$ at a sample period of $T/3$. The sequence $x_{T/3}(n)$ has 3 samples corresponding to every sample in $x(n)$. The sequence $x_T(n), x_{T/3}(n)$ and their discrete Fourier spectra are shown in Figure 4.3. The spectrum $X_T(f)$ of the signal $x_T(n)$ in Figure 4.3(a) has a frequency range of f_{s2}, while $X_{T/3}(f)$ in Figure 4.3(b), which is the spectrum of $x_{T/3}(n)$, has a range of $3f_{s2}$.

The objective of the interpolation process is to obtain a sequence $y(n)$ that approximates $x_{T/3}(n)$, given the samples of $x_T(n)$. Hence, to compute $y(n)$ we need to estimate the value of 2 additional samples for every sample in $x_T(n)$. The simplest way to interpolate is to approximate the additional samples with 0s as shown in Eq. 4.1. This corresponds to stuffing 2 zeros between adjacent samples of $x_T(n)$. The spectrum of the original signal $x_T(n)$ and the zero-stuffed version $x_{T/3,0}(n)$ are shown in Figure 4.4. The Fourier spectrum of $x_{T/3,0}(n)$ is derived in Eqs.4.1-

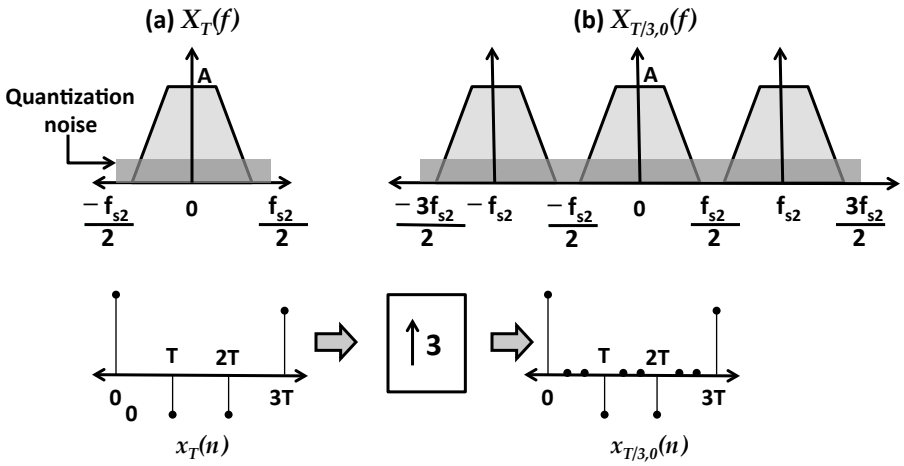

Figure 4.4: (a) Original signal sampled at f_{s2} (b) Zero-stuffed signal at $3f_{s2}$.

4.5 [SR73].

$$\begin{aligned} x_{T/3,0}(n) &= x_T(n/3),\ n = 0, \pm 3, \pm 6, \ldots \\ &= 0 \text{ otherwise.} \end{aligned} \tag{4.1}$$

The z-transform of $x_{T/3,0}(n)$ is given by:

$$\begin{aligned} X_{T/3,0}(z) &= \sum_{n=-\infty}^{\infty} x_T(n/3) z^{-n},\ n = 0, \pm 3, \pm 6, \ldots \\ &= \sum_{k=-\infty}^{\infty} x_T(k) z^{-3k},\ k = 0, \pm 1, \pm 2, \ldots \\ &= X_T(z^3) \end{aligned} \tag{4.2}$$

The frequency spectrum of $x_{T/3,0}(n)$ can be derived as:

$$\begin{aligned} X_{T/3,0}(e^{jw(T/3)}) &= X_T(e^{jw(T/3)*3}) \\ &= X_T(e^{jwT}) \end{aligned} \tag{4.3}$$

From Eq. 4.3 we can see that the spectrum $X_{T/3,0}(e^{jw(T/3)})$ is periodic with $w = 2\pi/T$. The original spectrum $X_T(e^{jwT})$ duplicates itself at $w = 0$, $w = 2\pi/T$, and $w = -2\pi/T$. This corresponds

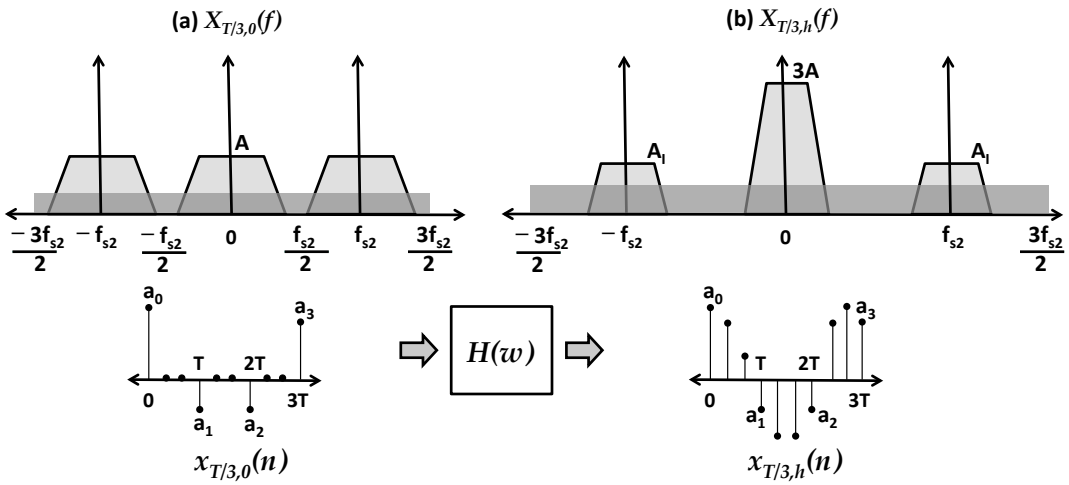

Figure 4.5: (a) Zero-stuffed signal at $3f_{s2}$ (b) Filtered signal at $3f_{s2}$.

to 0, f_{s2}, $-f_{s2}$ in the frequency axis, which explains the spectral images of $X_T(f)$ at f_{s2} and $-f_{s2}$ in Figure 4.4(b).

For the interpolated spectrum $X_{T/3,0}(f)$ to resemble the ideal spectrum of $X_{T/3}(f)$ (Fig. 4.3(b)), we need to suppress the spectral images at f_{s2} and $-f_{s2}$. These images can be suppressed if the zero-stuffed signal $x_{T/3,0}(n)$ is lowpass filtered to generate a new sequence $x_{T/3,h}(n)$ as shown in Figure 4.5. The lowpass filter $H(w)$ leaves the central signal replica at $f = 0$ unaltered while attenuating the images at f_{s2} and $-f_{s2}$. The resultant sequence $x_{T/3,h}(n)$ is the interpolated signal, where the additional 2 samples have been computed using the filtering (averaging) operation. The images are attenuated to level A_I as shown in Figure 4.5(b). The attenuation level A_I should ensure that the images remain below the spectral emission mask. The passband gain of $H(w)$ is set to 3 so that the original sample values a_0, a_1, a_2, $a_3, \ldots$ remain unscaled after interpolation. The passband gain scales the magnitude spectrum of the central signal by a factor of 3, as can be seen from the spectrum of $X_{T/3,h}(f)$. The power concentrated in the central signal ($f = 0$) increases when the images at f_{s2} and $-f_{s2}$ are attenuated.

4.1.2 Digital Front-end Architecture

The signal interpolation process in the Tx digital front-end (DFE) can therefore be regarded as a sequence of lowpass filtering operations that suppresses the spectral images of the baseband digital

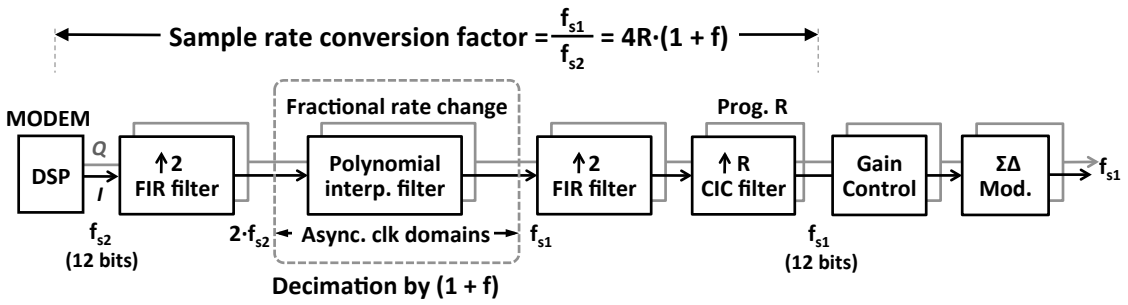

Figure 4.6: Linear RF modulator with a digital front-end (DFE), ΣΔ modulator, and RFDAC.

data at multiples of f_{s2}. As mentioned before, the interpolation occurs from the baseband rate of f_{s2} to the RFDAC sampling frequency of f_{s1}. With f_{s2} in the range of several MHz (5-30) and f_{s1} >800MHz, the interpolation factor f_{s1}/f_{s2} can be greater than 200. The first possibility is to perform this interpolation in a single stage, as was done by the upsampling and lowpass filtering operation in Figure 4.5(b). For an interpolation factor of 200 and a normalized signal bandwidth of 0.004π, the frequency response of $H(w)$ will be such that it has a cutoff frequency at $w = 0.004\pi$ and a stop band beginning at $w = 0.005\pi$. This would correspond to a rolloff of 0.001π resulting in a very constrained frequency response. For an FIR realization, such a filter can have in excess of 200 taps, leading to a computationally inefficient design. The interpolation process is therefore performed in stages and the overall interpolation factor is the product of the interpolation factors of the individual stages.

The block diagram of the Tx DFE is shown in Figure 4.6, with the FIR filter, the polynomial interpolator, and the cascaded-integrated-comb (CIC) filters responsible for interpolation. The I and Q components of the baseband input signal are first upsampled by a factor of 2 and then lowpass filtered by an FIR that suppresses the spectral images created during upsampling. The upsample-by-2 FIR could be implemented using half-band filters [CR83], but half-band filters are slightly constrained in their frequency response and have a fixed attenuation of 6dB at the frequency of $\pi/2$. In this upsample-by-2 FIR, we needed a higher attenuation at $\pi/2$ and therefore we chose to use a conventional FIR interpolator. Such an interpolator also gives us the freedom to generate a variety of frequency responses at different cutoff frequencies and attenuation when the filter is made programmable.

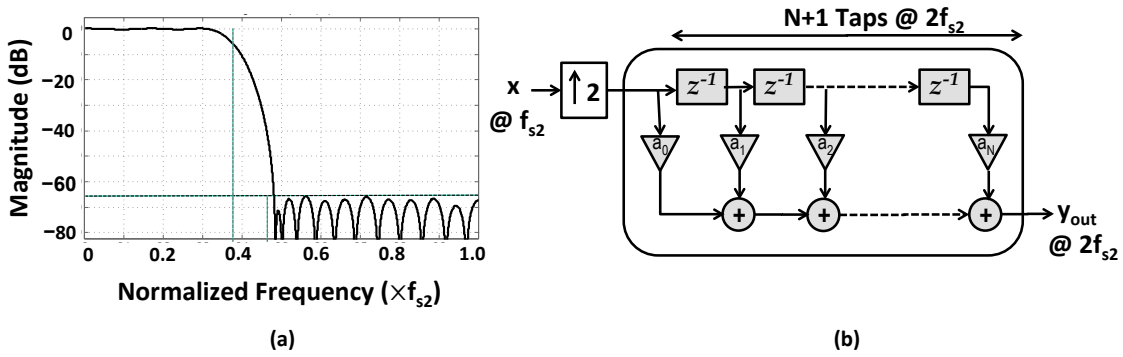

Figure 4.7: (a) Interpolation by 2 FIR response, (b) conventional FIR.

The frequency response of the 33-tap FIR used in this implementation is shown in Figure 4.7(a). The response has a stop band attenuation around 60dB and a passband ripple of 0.2dB. The ripples in the stop band become unequal after quantization of the filter coefficients to 14 bits of fractional accuracy, but this effect can be overlooked since the magnitude of the ripple is still below 60dB. The 60dB attenuation ensures that the transmitted power in the adjacent band remains below the target 45dBc level, leaving room for any spectral regrowth in the RFDAC or PA, which is dominant in the adjacent band. The conventional direct form realization of the FIR is shown in Figure 4.7(b). But this lowpass FIR interpolator does not take advantage of the upsampling operation preceding it. Hence all components in the filter operate at a frequency of $2f_{s2}$. A polyphase filter bank [Vai90], on the other hand, maximizes the computational efficiency of the filter from the knowledge that alternate samples input to the FIR are 0.

The polyphase implementation, shown in Figure 4.8, has two filter banks operating at sample rate f_{s2}, which is half the operating frequency of the conventional architecture. The filter banks compute two outputs $y(2m)$ and $y(2m+1)$ at every cycle of f_{s2}. The tap coefficients a_0, a_1, $a_2, \ldots,$ a_{33} are divided into even and odd groups. The output $y(2m)$ is generated using the even filter taps a_0, a_2, $a_4, \ldots,$ a_{32}, while the output $y(2m+1)$ is generated using the odd filter taps a_1, a_3, $a_5, \ldots,$ a_{33}. These two outputs are multiplexed together to generate the final output y_{out} at rate $2f_{s2}$. The lower frequency of operation results in significant power savings for this architecture (approximately half compared to Figure 4.7(b)). The power spectral density (PSD) of the FIR interpolator output is shown in Figure 4.8, where a 64-QAM modulated baseband signal sampled

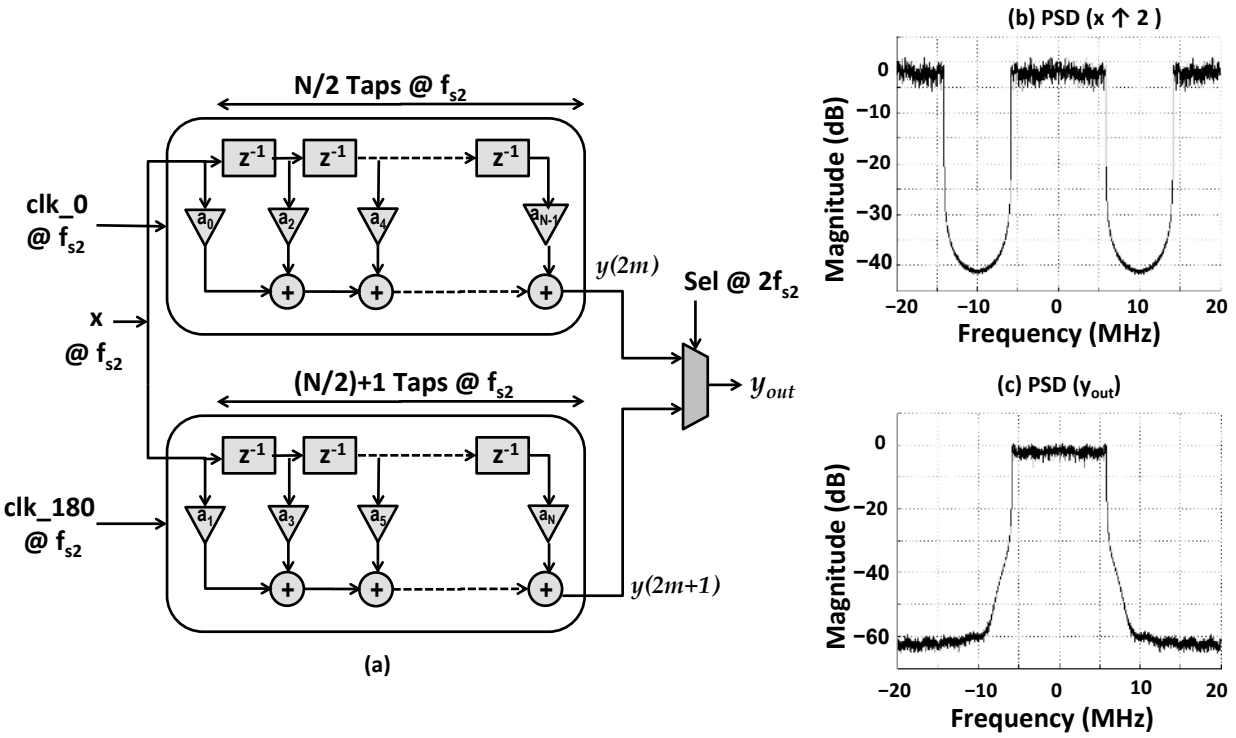

Figure 4.8: (a) Polyphase implementation of interpolation-by-2 FIR lowpass filter (b) PSD of zero-stuffed input to FIR (c) PSD of interpolated FIR output.

at 20MHz is interpolated to 40MHz. The PSD in Figure 4.8(b) is that of the zero-stuffed baseband data at 40MHz, and Figure 4.8(c) shows the filtered output. As expected, the FIR suppresses the image spectrum by 60dB.

4.1.3 Polynomial Interpolation

Proceeding with the signal processing along the Tx chain, the upsample-by-2 FIR is followed by a polynomial interpolation filter than resamples the signal from $2f_{s2}$ to a fraction of the RFDAC frequency, $f_{s1}/(6R)$. The interpolation ratio *IR* is given by $f_{s1}/(12 \cdot R \cdot f_{s2})$. This interpolator is similar in implementation to the polynomial interpolator discussed in Chapter III. The image attenuation provided by the polynomial interpolator should be such that the image attenuation level remains below the target 45dBc level and we also have to provide some margin since there is expected to be spectral regrowth and addition of noise during the ΣΔ modulation phase. Hence a target 50-60dB attenuation level was decided upon for the image bands being attenuated by the polynomial

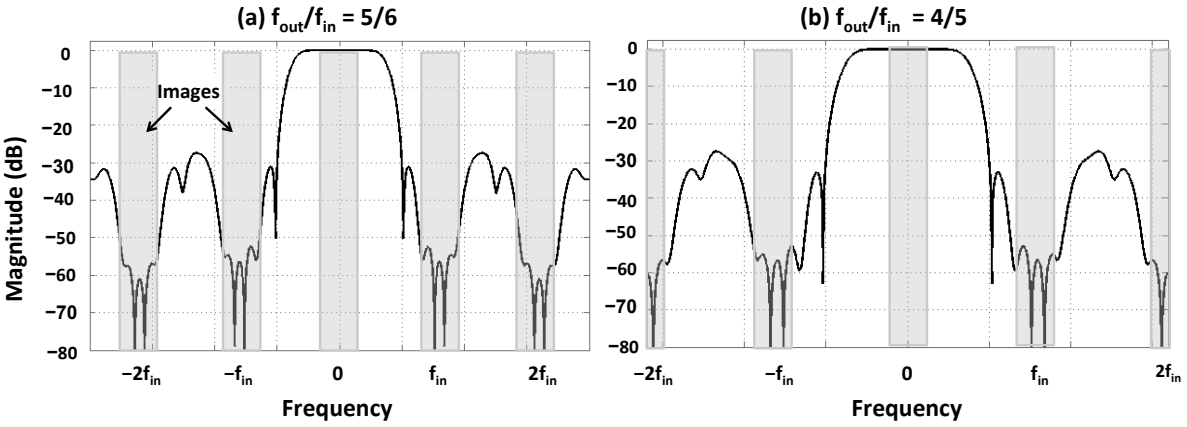

Figure 4.9: Farrow frequency response at different fractional interpolation ratios.

interpolator. We revisit the Farrow frequency response discussed in Figure 4.9 to analyze the attenuation characteristics.

As can be seen from Figure 4.9, the Farrow provides > 50dB attenuation at the dominant image locations, thus satisfying the system requirements. As we discussed in Chapter III, the location of the dominant image is centered at $f_{in}-f_{out}$, where f_{in} is the sample rate at the interpolator input while f_{out} is the sample rate at the interpolator output.

A second stage of upsample-by-2 FIR follows the fractional sample-rate converter. The frequency response of this FIR filter along with the expected image attenuation is shown in Figure 4.10. This FIR has a more relaxed transition band compared to the FIR used in the first stage of the DFE, with the passband edge at 0.4π and 50dB attenuation at frequency 0.8π. The relaxed

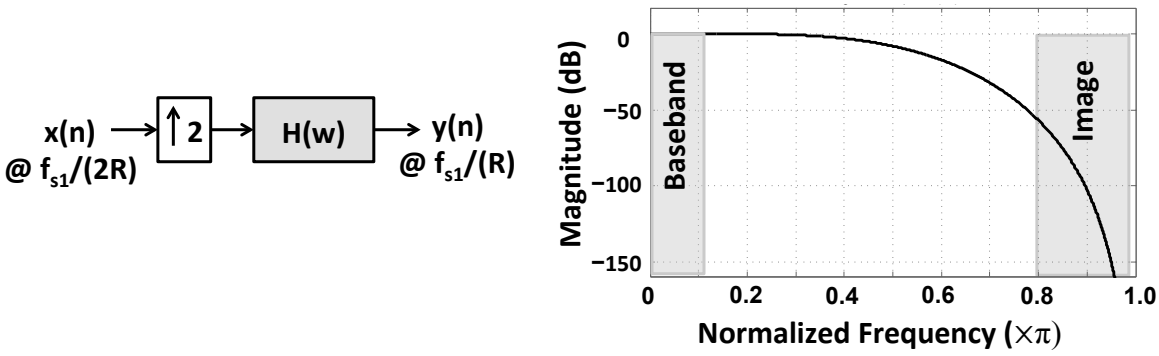

Figure 4.10: Conventional cascade integrated comb (CIC) interpolator.

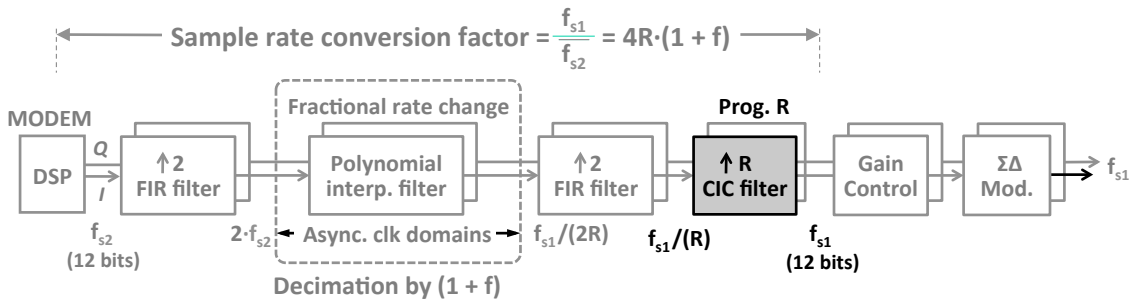

Figure 4.11: Linear RF modulator with a digital front-end (DFE).

transition is feasible since the baseband digital data has already passed through two stages of interpolation, and the oversampling ratio is close to 2. The grey rectangles in the frequency response in Figure 4.10 indicate the position of the baseband signal and its image during the filtering operation of by $H(w)$. The FIR function $H(w)$ was realized using 7 biquad sections in a direct form realization. This passband ripple for this filter is at to 0.01dB and minimizes the corruption of the in-band signal.

This FIR stage is followed by a reconfigurable cascade integrated comb (CIC) filter, shown in Figure 4.11, that achieves the final stage of interpolation by the variable factor R.

4.2 CIC Interpolation

CIC interpolators are a cascade of differentiators and integrators separated by a zero-stuffing upsampler. Figure 4.12 shows such an interpolator, which upsamples by a factor of R from an input sampling frequency of f_{s1}/R to an output frequency of f_{s1}. This structure cascades two differentiators to the left and two integrators to the right of the upsampler, making it a 2^{nd}-order CIC interpolator. The variable, M, is the differential delay in the differentiator and integrator loops.

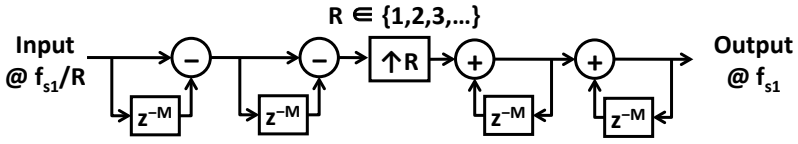

Figure 4.12: Conventional cascade integrated comb (CIC) interpolator.

Additional differentiator and integrator sections increase the order, N, of the filter and can obtain higher attenuation of the spectral images in the up-sampled data. Higher order filters with more than one section of differentiators and integrators can be pipelined; a pipeline registers between the sections reduces the critical path of the CIC to a single adder delay.

The frequency response of the CIC can be derived from its transfer function as shown below in Eqs. 4.4-4.9. The response is derived from the z-transform of the CIC (Eq. 4.4) by replacing all instances of z with e^{jw} to obtain the frequency response (Eqs. 4.5-4.9). The simplified magnitude response is shown in Eq. 4.9 to be a product of sine and secant functions.

$$H_{CIC}(z)|_{\mathrm{f_{s1}}} = \frac{(1-z^{-RM})^N}{(1-z^{-1})^N} \tag{4.4}$$

$$H_{CIC}(w)|_{\mathrm{f_{s1}}} = \frac{(1-e^{-jwRM})^N}{(1-e^{-jw})^N} \tag{4.5}$$

$$H_{CIC}(w)|_{\mathrm{f_{s1}}} = \frac{e^{-jwRMN/2}(e^{jwRM/2}-e^{-jwRM/2})^N}{e^{-jwN/2}(e^{jw/2}-e^{-jw/2})^N} \tag{4.6}$$

$$H_{CIC}(w)|_{\mathrm{f_{s1}}} = \frac{e^{-jwRMN/2}}{e^{-jwMN/2}} \frac{(e^{jwMR/2}-e^{-jwMR/2})^N}{(2j)^N} \frac{(2j)^N}{(e^{jw/2}-e^{-jw})^N} \tag{4.7}$$

$$H_{CIC}(w)|_{\mathrm{f_{s1}}} = \frac{e^{-jwRMN/2}}{e^{-jwN/2}} \frac{(sin(wRM/2))^N}{(sin(w/2))^N} \tag{4.8}$$

$$|H_{CIC}(w)| = \left| \frac{(sin(wRM/2))}{(sin(w/2))} \right|^N \tag{4.9}$$

This response is plotted in Figure 4.13 for R = 3 and various values of N and M. As is apparent from the figure, increasing the order, N, increases the attenuation of the images. At N = 2, the

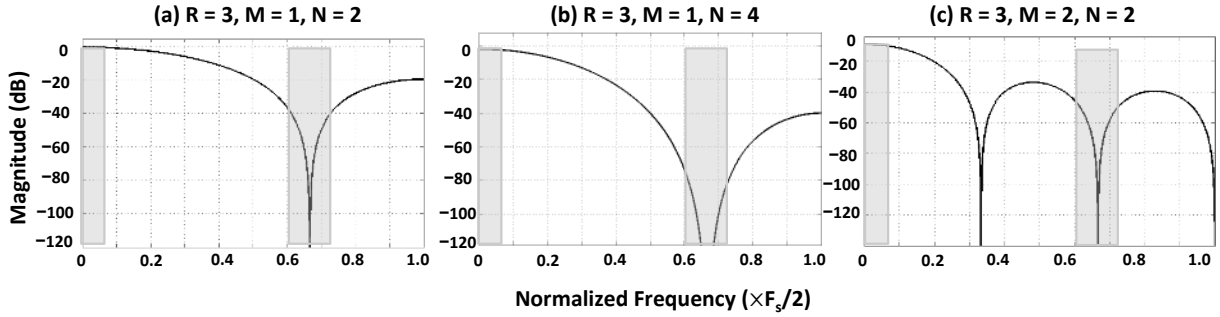

Figure 4.13: CIC Interpolator frequency response for various values of R, N, and M.

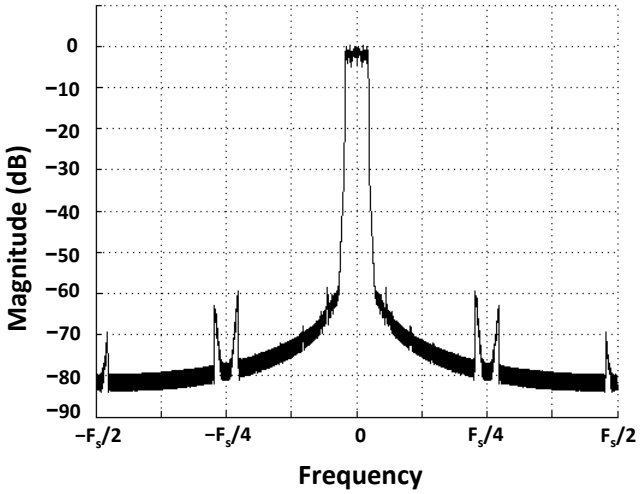

Figure 4.14: CIC Interpolator output frequency response for R = 4, N = 3, and M = 1.

image attenuation from Figure 4.13(a) is 40dB while at N = 4, in Figure 4.13(b), the attenuation increases to 80dB. Increasing the differential delay to M = 2 introduces an additional null at $f_s/3$, which can be used to suppress any images that occur in this strategic location. We use a reconfigurable 4^{th}-order CIC interpolator, shown in Figure 4.12, with a variable rate change factor R. This architecture is inexpensive to realize in hardware since the computations are restricted to additions and subtractions. The number of sections (order) in the filter can be varied by bypassing cascaded sections. The interpolator raises the sample rate of the baseband data from f_{s1}/R to f_{s1}. The spectrum of the CIC interpolated data for a baseband input is shown in Figure 4.14. The CIC setting was at R = 4, N = 3, and M = 1 for this example. We see that for this particular setting the images are attenuated by 60dB by the nulls of the CIC response.

Following the CIC stage is a programmable gain block, shown in Figure 4.11, that can tune the signal amplitude in steps of 6dB and digitally controls the output power at the end of the transmitter chain. This is followed by the digital ΣΔ modulator that is the last stage in the DFE chain.

4.3 ΣΔ Modulation

Ordinarily, the 12-bit signal sampled at f_{s1} would be directly sent for carrier multiplication and D/A conversion in a digital-to-RF converter (DRFC / RFDAC), as was done in [ESK07]. But achieving

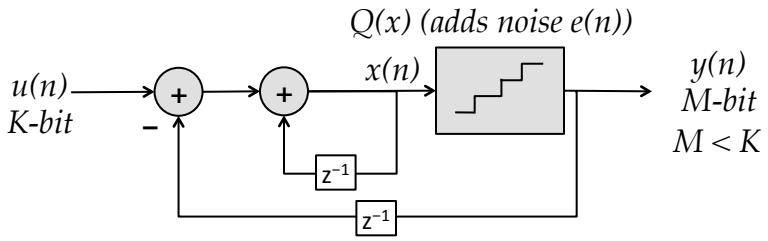

Figure 4.15: 1^{st}-order ΣΔ modulator with an M-bit quantizer.

12-bit resolution in the D/A converter at sample rates close to Gs/s can be extremely challenging. The work in [ESK07] demonstrates a 10-bit RFDAC that can handle signal bandwidths of up to 5MHz at a sample rate of 307.2Ms/s. For higher signal bandwidths, the D/A frequency will be higher and the DAC specifications more challenging. The primary cause for concern is the inherent mismatches between the unit cells of the D/A converter. The number of unit cells increases exponentially with the input bitwidth, and it becomes tougher to ensure that all the cells have identical voltage/current characteristics. We will examine the effect of mismatches in the D/A converter in the next chapter. In this chapter, it is sufficient to say that a reduction in wordlength is advantageous since it simplifies the design constraints on the RFDAC. The 12-bit digital signal is therefore compressed to a lower width using ΣΔ modulation.

A digital ΣΔ modulator compresses the wordlength of its input by shaping added quantization noise away from the signal of interest. The strategy works especially well when a maximum noise level has to be maintained in a well-defined bandwidth around the signal. Figure 4.15 shows an example of a 1^{st}-order noise shaping modulator that compresses a K-bit signal to an M-bit output while highpass filtering the quantization noise. The filtering behavior becomes apparent when we derive the transfer function of the modulator as shown in Eqs. 4.10−4.14.

$$Y(z) = X(z) + E(z) \tag{4.10}$$

$$X(z) = U(z) - z^{-1}Y(z) + z^{-1}X(z) \tag{4.11}$$

$$X(z)(1 - z^{-1}) = U(z) - z^{-1}Y(z) \tag{4.12}$$

$$Y(z) = U(z) + (1 - z^{-1})E(z) \tag{4.13}$$

$$Y(Z) = U(z) + H(z)E(z) \tag{4.14}$$

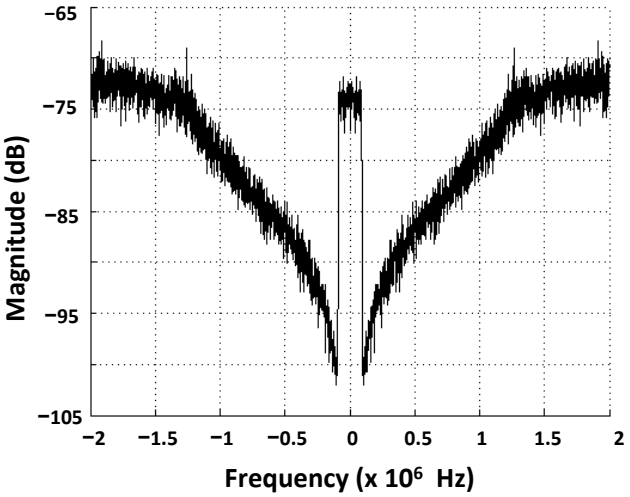

Figure 4.16: Spectrum of 1-bit output with 1st-order noise shaping.

The output $Y(z)$ is a combination of the input $U(z)$ and the highpass filtered quantization noise $E(z)(1-z^{-1})$. The order of modulation can be increased by using multiple integrator sections in cascade, each embedded within the ΣΔ feedback loop. This results in $(1-z^{-1})^N$ type of noise shaping, which is very popular since it is easy to implement in hardware, requiring only adders and delay elements. The noise shaping achieved for the 1^{st}-order function is illustrated in Figure 4.16 for an output bit width of $M = 1$. The single bit output completely eliminates any mismatch problems that may occur between the unit cells of the RFDAC. The noise level, however, is high, rising to about 20dB below the main signal level, which is unacceptable due to the emission mask specifications. In order to reduce the noise level, we can either increase the order of modulation, N, the number of bits, M, or the sampling frequency f_{s1}.

To increase the flexibility of design parameters, especially the choice of the noise transfer function, $H(z)$, we next look at a general topology of ΣΔ modulators, referred to as noise shaping coders [Nor93]. Figure 4.17 shows the topology of such a modulator where a K-bit signal $u(n)$ is compressed to an M-bit output $y(n)$. Word length compression is achieved through a quantizer $Q(x)$ that adds quantization noise $e(n)$ to its input $x(n)$. This quantization noise is shaped outside the band of interest by the noise transfer function (NTF), $H(z)$. Eqs. 4.15−4.17 derive the transfer

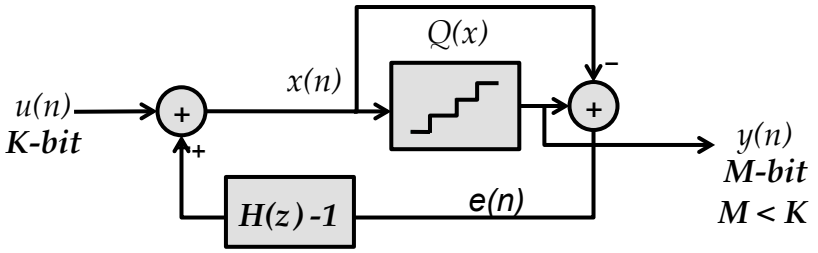

Figure 4.17: Noise shaping coder topology of a ΣΔ modulator.

function, $Y(z)$, of the modulator output.

$$X(z) = U(z) + (H(z) - 1)E(z) \tag{4.15}$$

$$E(z) = Y(z) - X(z) \tag{4.16}$$

$$Y(z) = U(z) + E(z)H(z) \tag{4.17}$$

The function $H(z) - 1$ should always be of the form $h_1 z^{-1} + h_2 z^{-2} + h_3 z^{-3} + \ldots$ to ensure that the feedback ΣΔ loop is not delay free. This would mean that the first coefficient h_0 of the function $H(z)$ must always equal 1. The primary variables in the design of such modulators are the output bitwidth, M, the sample rate of the modulator, f_{s1}, and the noise transfer function, $H(z)$. The quantization noise spectrum after modulation depends on all three variable in the following manner:

- M - As the output bit width of the quantizer $Q(x)$ increases, the noise power in $e(n)$ reduces. As M increases, the RFDAC input bit width also increases leading to higher resolution in the D/A converter.
- $H(z)$ - The shaped noise, $H(z)E(z)$, is a function of the noise transfer function, $H(z)$.
- f_{s1} - The shaped quantization noise spreads in the frequency range of f_{s1}. Hence, higher the value of f_{s1}, lower the noise level in the shaped spectrum.

In order to meet the spectral emission mask requirements, we need to ensure that the quantization noise level remains below 45dB in the frequency region close to the baseband signal. The

noise transfer function, $H(z)$, is hence a highpass filter that shapes the added noise towards the edges of the frequency spectrum. The highpass filtered noise can be attenuated by a small order RF bandpass filter after the D/A conversion process. A careful choice of the variables M, f_{s1}, and $H(z)$ will help us meet the emission mask while also ensuring that we operate at low levels of power consumption in the transmitter chain. We next look at the tunability of each of these variables in detail.

4.3.1 Quantizer Output Width (M)

Although increasing the quantizer bitwidth seems like an attractive solution, this approach has its share of problems. As mentioned before, higher values of M lead to mismatches in the unit cells of the RFDAC that follows the $\Sigma\Delta$ modulator. The mismatches in turn will result in linearity degradation at the output of the DAC. Hence, we choose to minimize the value of M as far as possible. In order to completely eliminate the mismatch problems between the unit cells of the RFDAC, the work in [FFS09] uses a 1-bit quantizer (M = 1) with a 3^{rd}-order NTF, $H(z)$, and a sample rate f_{s1} = 4Gs/s for a 1GHz RF carrier. The high oversampling ratio w.r.t. to the RF carrier is needed to spread the large amount of quantization noise across a wide frequency range. The 4Gs/s sample rate leads to excessive power dissipation of 120mW in the transmitter. This approach also suffers from the need of a very high selectivity bandpass filter after the RFDAC. The

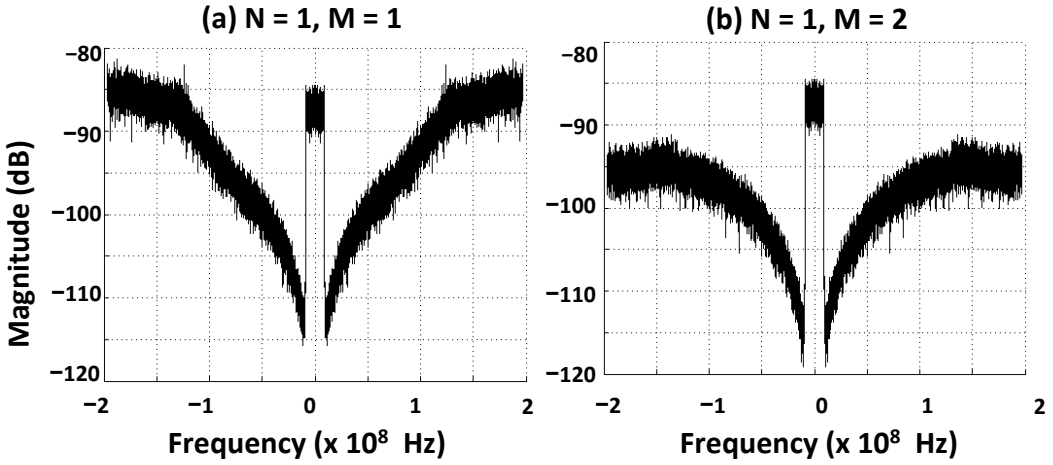

Figure 4.18: Effect on noise shaping characteristics with increasing quantizer bits (M).

filter is required, since the output bit width of $M = 1$ leads to excessive out-of-band quantization noise that has to be eliminated by a sharply selective filter. We will discuss more details about the bandpass filter in a later section of this chapter. Hence using a single-bit modulator becomes practically infeasible for this system and we adopt the use of multi-bit modulators. Figure 4.18 compares the spectrum of a 1st-order noise shaped, 1-bit wide output (Figure 4.18(a)) with that of a 1st-order noise shaped, 2-bit wide output (Figure 4.18(b)). We find that the noise level adjacent to the signal bandwidth has lowered by approximately 5dB with a single bit increase.

4.3.2 Order of $\Sigma\Delta$ Modulation (N)

Increasing the order of modulation results in a sharper notch at DC in the highpass frequency response of the noise function, $H(z)$. This in turn results in a lower noise level around the signal of interest, which helps in meeting the emission mask specifications about the signal. The effect on the noise shaping characteristics with increasing order (N) is shown in Figure 4.19. Figure 4.19(a) shows the spectrum of a 1st-order modulator ($H(z) = 1 - z^{-1}$) output with a 1-bit quantizer. Compared to this, the noise level adjacent to the signal is approximately 7dB lower for the 2nd-order modulator ($H(z) = (1 - z^{-1})^2$) in Figure 4.19(b). The out-of-band quantization noise level rises sharply in Figure 4.19(b), which means that the bandpass filter following the RFDAC will need a sharply selective response to clear out the quantization noise. This is one of the drawbacks of

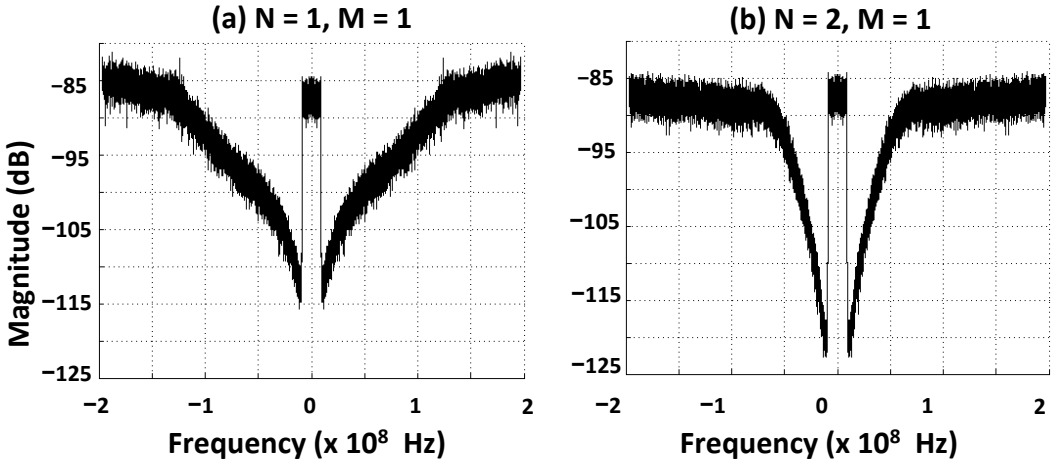

Figure 4.19: Effect on noise shaping characteristics with increasing order of modulation (N).

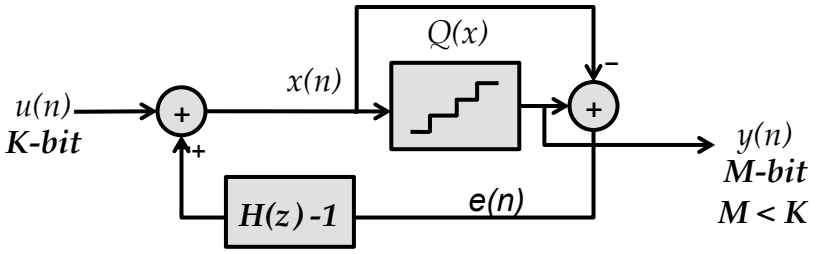

Figure 4.20: Noise shaping coder topology of a ΣΔ modulator.

increasing the modulation order.

Increasing the modulation order indefinitely also leads to quantizer overload and instability in the ΣΔ loop. To better understand this issue, we revisit the ΣΔ loop topology, shown in Figure 4.20. From the signal processing flow in the figure, we can see that the maximum magnitude of the quantizer input $x(n)$ is determined by the maximum input signal magnitude, $\max(|u(n)|)$, and the maximum magnitude of the feedback filter output. This relation has been expressed in Eq. 4.19, where the term $h_0e(n) + h_1e(n-1) + \ldots + h_Ne(0)$ is the convolution between H(z) and E(z) in the time domain. Since we know that $h_0 = 1$, Eq. 4.19 can be simplified to the form in Eq. 4.20. Quantizer overloading occurs when the highest magnitude of $x(n)$ exceeds the highest level, L, in the quantizer, leading to instability and distortion in the ΣΔ loop. To ensure that the loop remains stable, we must restrict the magnitude of $x(n)$ to L. This can be done by scaling the magnitude of the input $u(n)$, but the signal power reduces linearly with such scaling. The other alternative is to restrict the magnitude of the filtered output. The upper bound on this magnitude can be computed as in Eq. 4.21.

$$X(z) = U(z) + (H(z) - 1)E(z) \tag{4.18}$$

$$max(|x(n)|) \leq max(|u(n)|) + max(|h_0e(n) + h_1e(n-1) + \ldots + h_Ne(0) - e(n)|) \tag{4.19}$$

$$max(|x(n)|) \leq max(|u(n)|) + max(|h_1e(n-1) + \ldots + h_Ne(0)|) \tag{4.20}$$

$$max(|h_1 e(n-1) + \ldots + h_N e(0)|) \leq (\sum_{n=0}^{n=N} |h(n)| - 1) \times max(|e(n)|) \tag{4.21}$$

$$max(|x(n)|) \leq max(|u(n)|) + 0.5 \times (\sum_{n=0}^{n=N} |h(n)| - 1) \leq L \tag{4.22}$$

Assuming a quantization interval of 1, the error $e(n)$ will vary between -0.5 to $+0.5$. Hence Eq. 4.19 can be simplified to the form in Eq. 4.22. The primary bottleneck in satisfying the condition in Eq. 4.22 is to restrict the 1-norm of the coefficients of the noise shaping filter, given by the term $\sum_{n=0}^{n=N} |h(n)|$. As the order of the filter, N, increases, the 1-norm also increases. For example, the function $H(z) = 1 - z^{-1}$ has a 1-norm of 2, while $H(z) = (1 - z^{-1})^2$ has a 1-norm of 4. This would mean that when the 1-norm of $H(z)$ increases, we need to increase the value of L in order to support the higher dynamic range of $x(n)$ to ensure stability. This is unattractive due to an increased number of levels in the RFDAC. Hence a large part of the optimization lies in the design of a noise transfer function, $H(z)$, that has a low 1-norm and also satisfies the noise shaping characteristics necessary to keep the noise level adjacent to the signal below the emission mask.

4.3.3 Sample Rate (f_{s1})

Due to the nature of the shaped quantization noise in the ΣΔ spectrum, a distinct relation, shown in Eq. 4.23, must exist between the sample rate f_{s1} and the RF carrier f_{RF}.

$$f_{RF} = \frac{n}{2} f_{s1}, \quad n \in \{1, 2, 3, \ldots\} \tag{4.23}$$

Such a relationship is necessary to avoid aliasing of the quantization noise spectrum into the signal of interest after the baseband signal is up-converted to the RF carrier frequency. When the spectrum of the baseband signal, repeating at intervals of f_{s1}, is up-converted to f_{RF}, two copies of the baseband spectrum are created at f_{RF} and $-f_{RF}$. If the relation in Eq. 4.23 is not satisfied, then the two copies alias with each other corrupting the baseband signal. An example of aliasing is shown in Figure 4.21(a), where $f_{RF} = 4/3 f_{s1}$. We find that the quantization noise from the spectrum centered at $-f_{RF}$, shown in grey, aliases with the signal positioned at f_{RF} and vice versa. When the relation in Eq. 4.23 is satisfied then the notches in the quantization noise spectrum for both copies

at f_{RF} and $-f_{RF}$ align with the position of the RF modulated data after up-conversion, as shown in Figure 4.21(b) for $f_{RF} = 3/2f_{s1}$.

In reality, digital noise leaks from the DAC input into the output spectrum at multiples of $f_{s1}/2$. The use of differential current cells in the RFDAC suppresses the noise at even multiples of $f_{s1}/2$, but the noise at the odd multiples of $f_{s1}/2$ still remains. To avoid any SNR degradation due to this noise, the relation in Eq. 4.23 can be modified to the form in Eq. 4.24, that restricts the RF carrier frequency to even multiples of $f_{s1}/2$.

$$f_{RF} = nf_{s1}, \quad n \in \{1,2,3,\} \tag{4.24}$$

Hence the frequency of operation of the ΣΔ modulator and consequently the RFDAC becomes a function of f_{RF}. Assuming that the integer n remains constant in the implementation, our aim will be to maximize n as far as possible so as to minimize the frequency of operation of the DFE core. The effect on the noise shaping characteristics with doubling of the sampling frequency, f_{s1}, is shown in Figure 4.22 for a 2nd-order modulator with a 1-bit output. The original spectrum in Figure 4.22(a) has noise level at around -107dB in the adjacent signal band, while the upsampled signal has a noise level at -115dB. This corresponds to an 8dB lowering of the noise level with

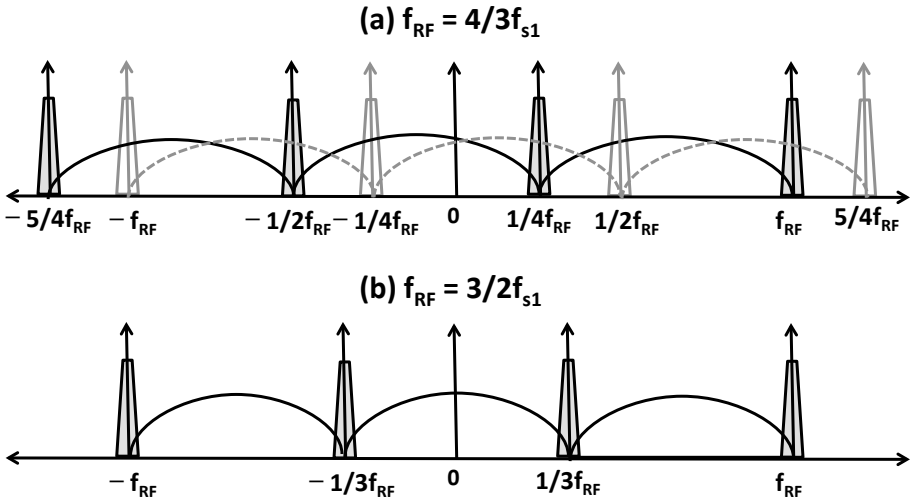

Figure 4.21: (a) Aliasing of quantization noise with signal after up-conversion to f_{RF} (b) up-conversion without aliasing.

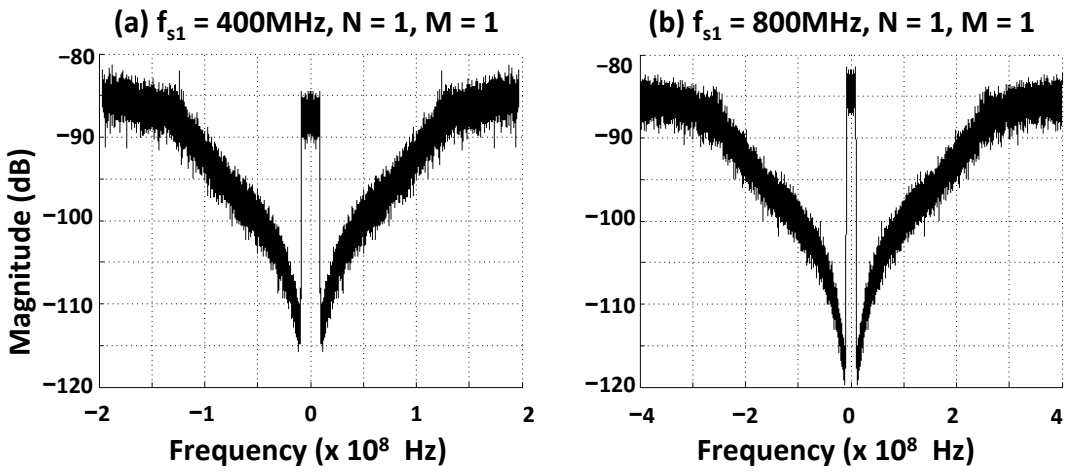

Figure 4.22: Effect of increasing sample rate on noise shaping characteristics.

doubling of the sampling frequency. Hence the sample rate increase is an attractive solution when it comes to satisfying the noise requirements of the transmitted spectrum. We briefly look at the frequency response of the RFDAC in the next section in order to complete the study of all components affecting the shape of the transmitted spectrum.

4.4 RFDAC Architecture

Although Chapter V is devoted to the architecture and circuit of the RFDAC, we briefly discuss the RFDAC frequency response in this chapter, since it aids in filtering the transmitted signal. This filtering helps relax the expected attenuation characteristics of the bandpass filter following the RFDAC. A DAC operating at an input sample rate of f_{s1}, has a sinc frequency response with the first nulls positioned at f_{s1} and $-f_{s1}$. Subsequent nulls in the response occur at multiples of f_{s1}. The nulls attenuate the replica of the baseband signal that repeat at multiples of the input sample rate. The extent of attenuation is determined by the oversampling ratio, given by f_{BW}/f_{s1}, where f_{BW} is the signal bandwidth. An example of this filtering operation is shown in Figure 4.23(a), where the DAC response is superimposed over a noise-shaped signal. We find that although the DAC sinc response attenuates the signal images, it does little to filter out the out-of-band quantization noise that is a result of $\Sigma\Delta$ modulation. In order to relax the specifications of the bandpass filter that follows the RFDAC, we need the DAC to attenuate the quantization noise to some extent.

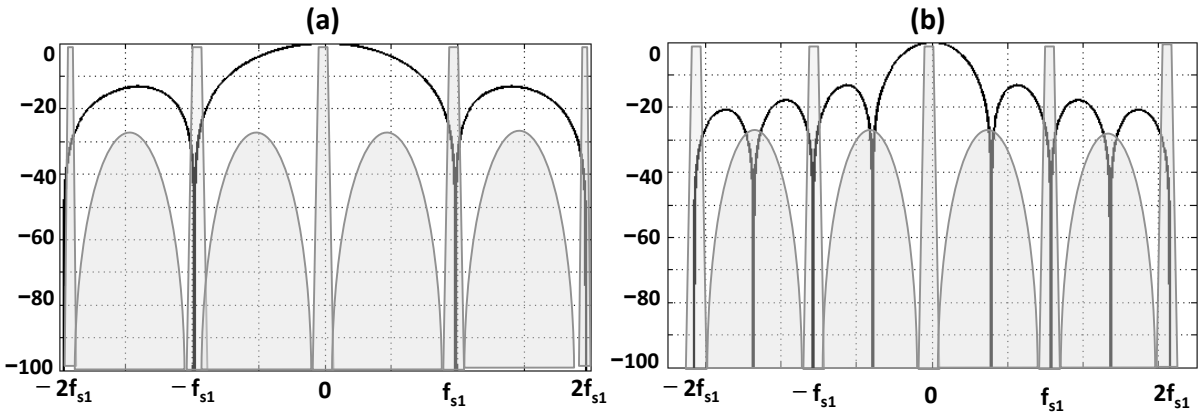

Figure 4.23: (a) Sinc filtering in RFDAC (b) Sinc filtering in 2-channel time-interleaved RFDAC.

This can be achieved through the use of a time-interleaved RFDAC architecture. The time-interleaved DAC accepts two interleaved input channels, each at rate $f_{s1}/2$. The two channels are synchronized to clocks that are a 180° out of phase. This results in nulls at multiples of $f_{s1}/2$ instead of f_{s1}, which means that the additional nulls introduced by the time-interleaved DAC are able to attenuate the out-of-band quantization noise at these locations. An example of this scenario is shown in Figure 4.23(b), where the time-interleaved DAC response is superimposed over the noise-shaped signal. We can see that the nulls at the Nyquist frequency ($f_{s1}/2$) suppress the quantization noise when it is at its highest magnitude. An apparent extension of this approach can be to introduce additional nulls in the DAC response by the use of more than 2 time-interleaved channels. We did not choose to do this, however, since with the introduction of additional nulls, the central lobe of the DAC response compresses. This would mean that the sinc response will introduce increased drooping in the main signal. The other problem lies in the timing alignment of the time-interleaved channels at the output, which becomes more difficult with increasing number of channels. Hence, we adopted the 2-channel time-interleaved DAC architecture in this work.

4.5 Proposed Solution

The problem now crystallizes to finding an optimal noise transfer function (NTF) that ensures stability of the $\Sigma\Delta$ loop while minimizing the number of quantization levels, $2L+1$, and sample

rate, f_{s1}. In this work we use the approach suggested in [Nor93], which proposes the use of customized FIR noise transfer functions for arbitrarily shaped magnitude response. We used this approach to design an NTF that would maximize the bandwidth over which the shaped noise level remains below the emission mask for a given sample rate and quantization levels. The signal-flow graph for the modulator with the FIR function embedded in the loop is shown in Figure 4.24. The transfer function, $H(z)$, is designed in such a way that it maximizes the bandwidth over which we meet the emission mask while also minimizing the 1-norm of its coefficients so that the modulator loop remains stable. According to the work in [Nor93], the filter norm is minimum when the filter approaches minimum phase characteristics in its phase response. It is important to note that although we care about the magnitude characteristics of $H(z)$, we need not constrain the phase response, since the phase of the shaped noise does not affect the Tx specifications. Hence an FIR with nonlinear phase response becomes acceptable as the NTF.

The algorithm used for determining the optimal NTF is stated below:

- Step 1: Fix the characteristics of the RF bandpass filter (BPF) that follows the RFDAC. Decide on its order and bandwidth. Fix the sample rate f_{s1}. Fix the quantization step interval, Δ, in the $\Sigma\Delta$ loop quantizer and the maximum magnitude of the input $u(n)$.

- Step 2: Determine the corresponding noise-shaping characteristics (magnitude response) of

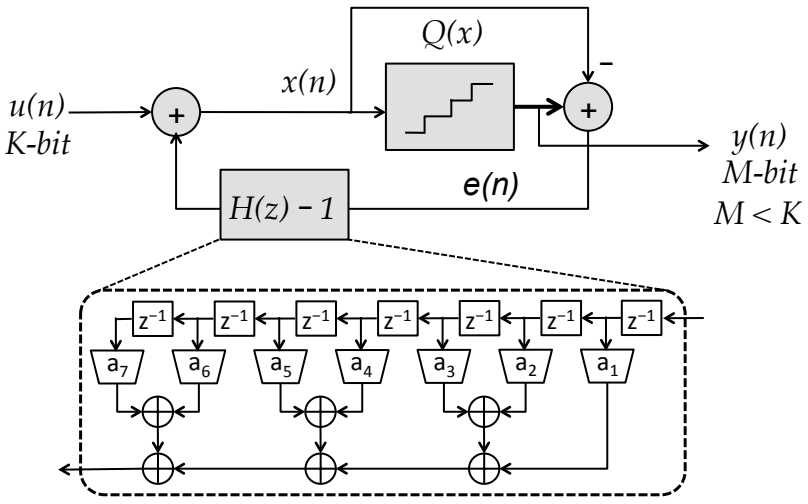

Figure 4.24: Non recursive realization of the NTF, $H(z)$, using an FIR.

the noise shaping function, $H(z)$, of the modulator.

- Step 3: Design a minimum-phase FIR filter that satisfies the characteristics determined in Step 2. Use the METEOR [SPK92] tool for this design.

- Step 4: Compute the 1-norm of this FIR. Determine the number of quantization levels, $2L+1$, for the $\Sigma\Delta$ loop to be stable.

- Step 5: Determine the noise-shaped spectrum using $H(z)$ and L computed in steps 3 and 4. Incorporate the filtering effect of the 2-channel time-interleaved DAC and the bandpass filter when evaluating the spectrum.

- Step 6: Determine if the spectrum satisfies the emission mask requirement and/or Rx band noise specifications in an FDD system. If no, then increase the input magnitude until the spectrum satisfies the emission mask. Recompute the number of quantization levels, $2L+1$, required to keep the loop stable.

- Step 7: If the computed value of $2L+1$ is acceptable, then stop. If no, then return to Step 1, change the initial parameters, and repeat Steps 2-6.

The BPF was assumed to be a Bessel filter [JS07], and its frequency response was extracted in MATLAB for various orders and bandwidths. METEOR is a tool developed at Princeton and can be used to design minimum-phase filters for arbitrarily shaped frequency responses. Due to the previously established constraint of $f_{RF} = nf_{s1}$ (Eq. 4.24), it was easiest to fix the value of n for a given f_{RF}.

We began with the choice of $f_{RF} = 4f_{s1}$, which would restrict the sample rate between 500Ms/s to 675Ms/s for RF carriers between 2GHz to 2.7GHz. This proved to be a very tight constraint, since it was difficult to minimize the noise level below the target level of 45dBc in the adjacent signal bands, without increasing L to values beyond 60 or having a very sharply selective bandpass filter. This can be seen from the few representative architectures shown in Figure 4.25(a), where the modulators either required quantization beyond 60 levels (markers on the right of the graph) or bandpass filters with bandwidth in the tune of 100MHz (markers on the bottom-left of the graph)

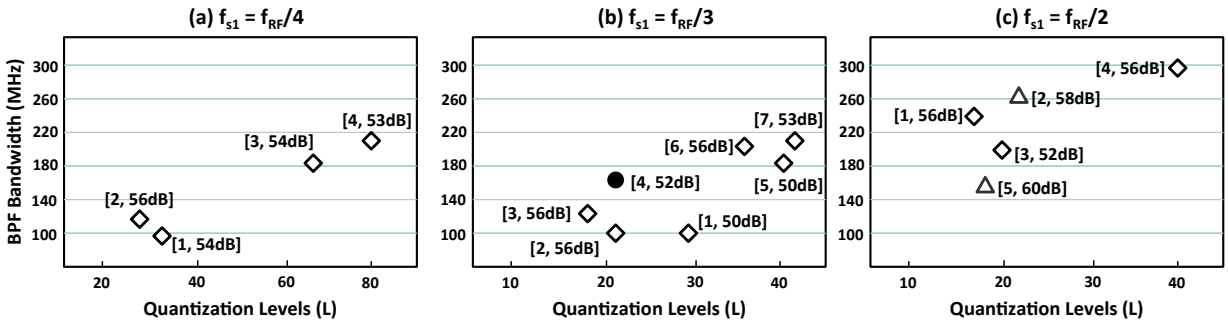

Figure 4.25: Sigma-delta modulator architectures generated using the iterative flow.

for an RF carrier of 2.0GHz. The text $[a, b]$ next to the markers on the graph represent the index and the ACPR performance of the architecture respectively. We next proceeded with the sample rate of $f_{s1} = f_{RF}/3$. This was a more acceptable choice, since the sample rate now varied between 667Ms/s to 900Ms/s, which is sufficient to reduce the noise level below the emission mask with quantization levels ranging between 20-40. The bandpass filter requirement for this sample rate was dominated by the noise level requirement in the Rx band, which is 190MHz away from the Tx signal in the first LTE band. We found several architectures with low quantization levels between 20-40 and reasonable filter bandwidths between 150-200MHz, as can be seen in the graph in Figure 4.25(b). This corresponds to a Q of 10-13 for an RF carrier of 2.0GHz.

The final round of iterations was performed for a sample rate $f_{RF} = 2f_{s1}$, which yielded architectures with larger filter bandwidths as can be seen in the graph in Figure 4.25(c). This sample rate however was about 50% larger than the previous sample rate of $f_{RF}/3$ and leads to higher power consumption and timing infeasibility issues in the modulator. The triangles in the graph of Figure 4.25(c) represent architectures that became timing infeasible in the 65nm technology we targeted for implementation. This choice of sample rate will become feasible and attractive when we move to modulator implementations in more advanced technology, where high-speed implementations can be realized with lower power consumption.

Eventually we proceeded with an implementation that would enable low power consumption as well as low D/A complexity. This selection is shown by the solid circle in the graph in Figure 4.25(b), which has 21 levels of quantization at a bandpass filter bandwidth of 160MHz. This modulator was realized using a 7^{th}-order FIR as the noise transfer function $H(z)$. The impulse

response of the FIR along with the pole-zero plot is shown in Figure 4.26. From the pole-zero plot in Figure 4.26(b), we can see that $H(z)$ has seven poles at the origin, and all the zeros of the transfer function lie within the unit circle, making $H(z)$ a minimum phase filter. The 1-norm of the filter coefficients is 10.823. An input with a highest magnitude of 5 was applied to this modulator. According to the stability conditions derived in Eq. 4.22, the maximum magnitude of $x(n)$ for this modulator will be limited to $0.5 \times (10.823 - 1) + 4.5 \leq 10$. Assuming a quantization interval of 1, $x(n)$ ranges between -10 to $+10$, which means that we need 21 levels to accommodate the dynamic range of $x(n)$ in the quantizer. Hence the output width, M, of the quantizer was set to 5, to accommodate the 21 levels in binary form.

4.6 Simulation Results

The choice of the 7th-order NTF, shown in Figure 4.26, and 21 levels of quantization results in the noise shaping characteristics shown in Figure 4.27(a). The dashed line represents the 45dBc emission mask, which was our design target. We find that the emission mask is satisfied up to a bandwidth of 160MHz for the noise-shaped signal after $\Sigma\Delta$ modulation at f_{s1}=667Ms/s for f_{RF}=2.0GHz. The spectrum after the RFDAC, which performs the D/A conversion and mixing with the RF carrier, is shown in Figure 4.27(b). The nulls in the spectrum at a distance of $f_{s1}/2$ from f_{RF} are a result of the time-interleaved DAC frequency response. It can be observed that the quantization

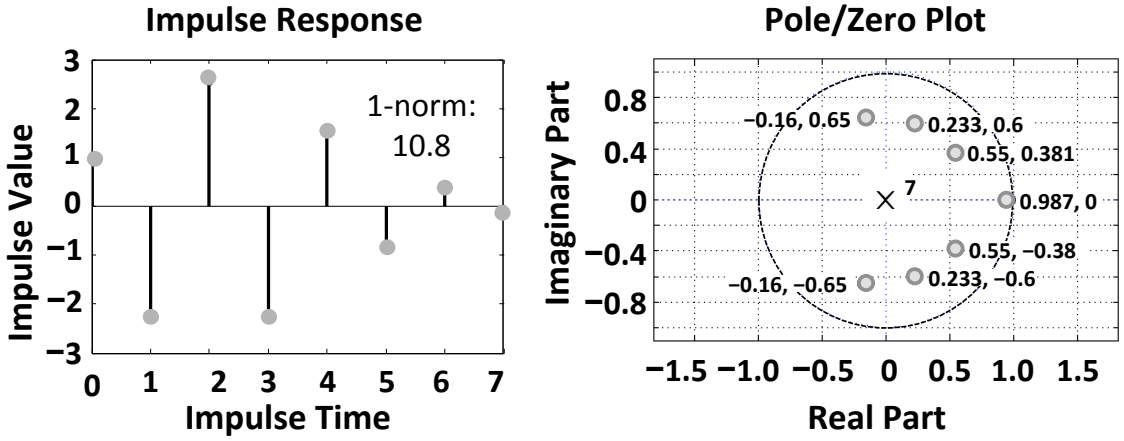

Figure 4.26: Impulse response and pole-zero plot of the 7th-order minimum phase FIR filter.

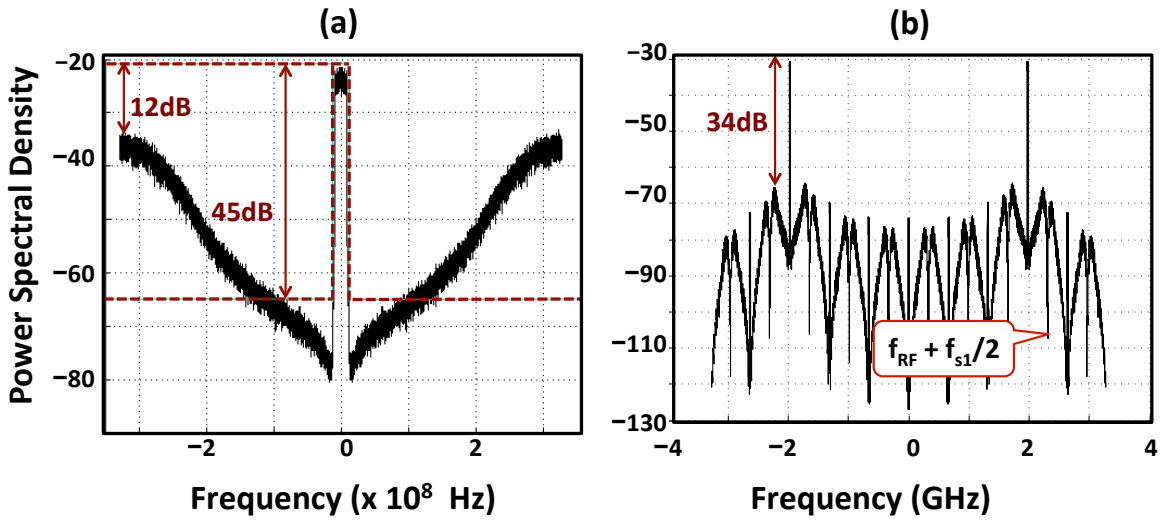

Figure 4.27: (a) Noise shaping with 7th-order FIR NTF (b) 2-channel time-interleaved DAC filtering on up-converted RF signal at f_{RF} = 2.4GHz.

noise at its peak locations ($f_{RF} \pm f_{s1}/2$) has been suppressed by 22dB compared to the spectrum in Figure 4.27(a).

In order to use this transmitted signal in the LTE standard, we operate the 2.0GHz signal in the TDD (time division duplex) mode (TDD band 34 for LTE). The LTE spectrum mask demands a 37dBc emission mask across the entire spectrum for a 20MHz signal. A 2nd-order RF bandpass filter with a bandwidth of 400MHz can clear out the additional quantization noise that violates the

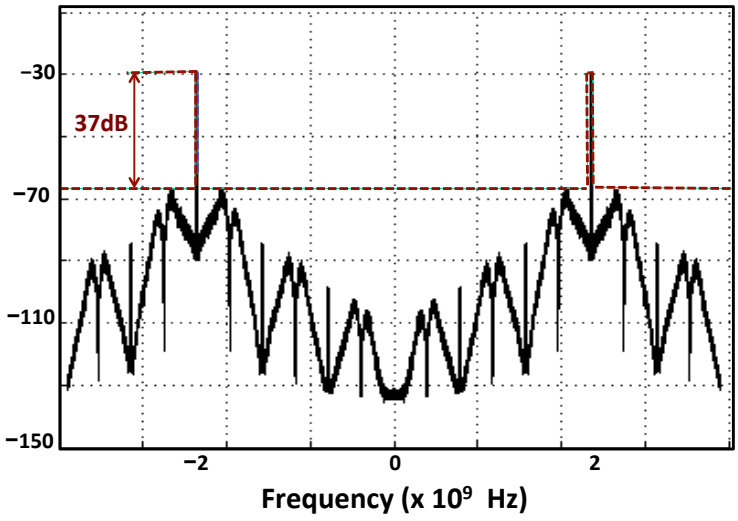

Figure 4.28: Spectrum of RF modulated signal after 2nd-order bandpass filtering.

emission mask. This selectivity, Q, of the bandpass filter is given by the ratio of the RF carrier and its bandwidth, which in this case equals 5. A passive LC filter with this order of selectivity can be easily implemented on-chip [JS07] for this Q value, thus providing a fully integrated solution for the RF transmitter in this mode. The spectrum of the bandpass filtered signal is shown in Figure 4.28.

A second example of operating the RF modulator in an FDD (frequency division duplex) system is presented next. In this example, the 20MHz signal has to be modulated to f_{RF} = 1.98GHz, which corresponds to Band 1 of the LTE FDD specifications. The downlink (Rx channel) lies 190MHz away at frequency 2.17GHz. In FDD systems, the constraints on noise power levels are tighter, since we not only have to satisfy the 37dBc emission mask but also ensure that the noise in the Rx channel is around −130dBm/Hz. This level of noise in the Rx band ensures that after duplexer attenuation (50dB), the noise level reaches the Rx noise specifications of −180dBm/Hz.

The spectrum of the signal after passing through the time-interleaved RFDAC is shown in Figure 4.29(a). The noise density in the Rx band at frequency 2.17GHz equals −128dBm/Hz for a 0dBm output signal. After the power amplifier stage, for an output power level of +20dBm, the noise density rises to −108dBm/Hz. Hence, the bandpass filter following the RFDAC will have to provide about 22dB of attenuation at 2.17GHz in order to satisfy the Rx noise requirements. A 6th-

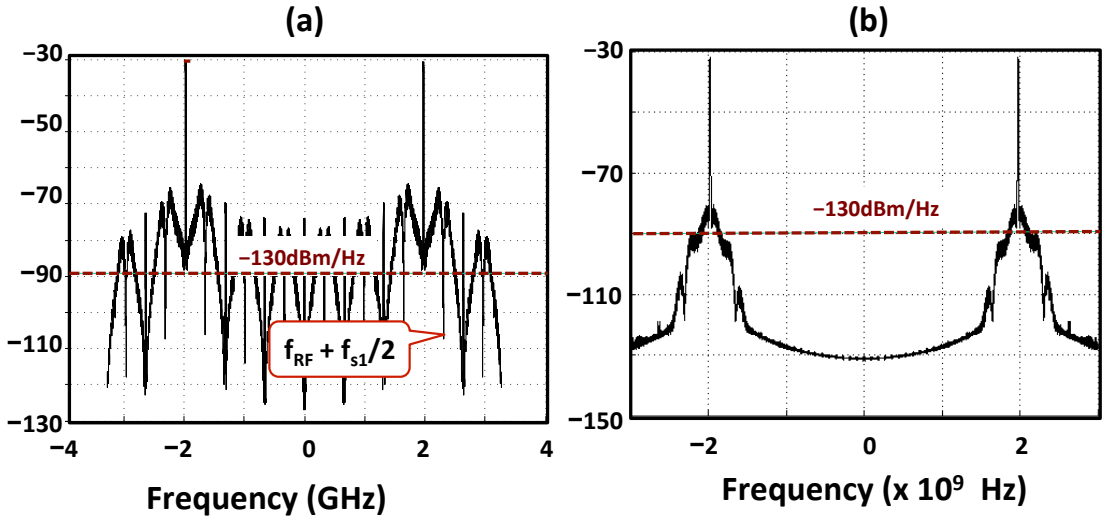

Figure 4.29: (a) Effect of 2-channel time-interleaved DAC filtering on upconverted RF signal at f_{RF} = 1.98GHz. (b) Spectrum of RF modulated signal at f_{RF}=1.98GHz after 3-pole bandpass filtering.

order BPF with a bandwidth of 160MHz can provide this attenuation. This bandwidth corresponds to a selectivity, Q, of 12.5 for the filter at the RF carrier of 1.98GHz. The bandpass filtered signal satisfying the 37dBc emission mask is shown in Figure 4.29(b). This filtered signal will be able to operate in FDD mode since the noise level is less than -130dBm/Hz at 2.17GHz (190MHz away from f_{RF}=1.98GHz).

In the next chapter we discuss the architecture and circuit level details of the final component in our Tx chain, which is the RFDAC.

CHAPTER 5

Mixing and Data Conversion in Radio Transmitters

Once the digital baseband signal has been oversampled, filtered, and compressed using ΣΔ modulation, it needs to be converted into analog form and modulated to the RF carrier frequency. An RFDAC performs both these tasks by first mixing the digital data bits with the sine and cosine of the RF carrier, and then converting this digital waveform into an analog signal using current-steering cells [ESK07]. The schematic of an individual DAC cell is shown in Figure 5.1, where the inputs are the data bit, b_i, sampled at rate $f_{s1}/2$, and the signals LO and $\overline{LO}$ are clock signals at frequency f_{RF} separated by a phase difference of π.

The input bit, b_i, is clocked at rate $f_{s1}/2$, since we use a time-interleaved DAC that accepts 2 channels of data at rate $f_{s1}/2$ (synchronized to clocks with phase difference of π). The data is first stored in a register that generates complementary outputs d_i and $\overline{d_i}$. The digital bits d_i and

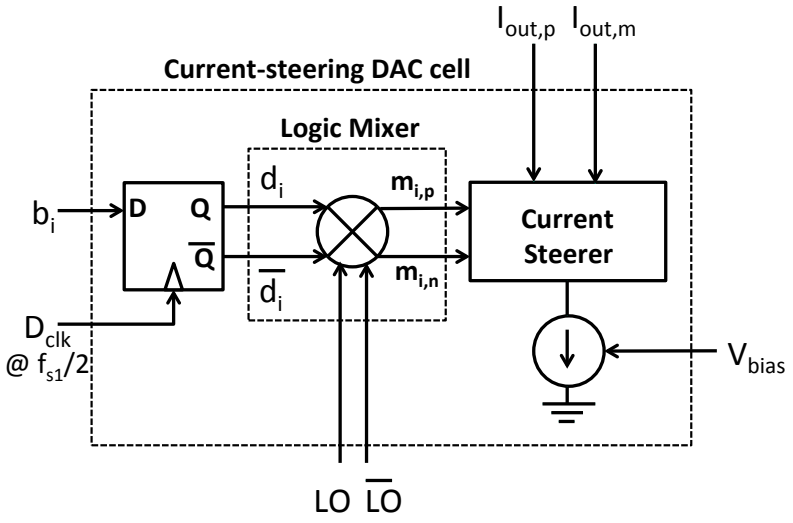

Figure 5.1: RFDAC unit cell schematic.

$\overline{d_i}$ are then mixed with the LO and $\overline{LO}$ signals in the logic mixer, generating the outputs $m_{i,p}$ and $m_{i,n}$. The mixer output is fed to the current steerer that steers the current from the the tail current source, I, in either the direction of $I_{out,m}$ or $I_{out,p}$ depending upon the polarity of $m_{i,m}$ and $m_{i,p}$. The magnitude of current flowing through the tail current source can be controlled using the bias signal, V_{bias}, which is also an input to current cell.

The schematic in Figure 5.2(a) shows the mixing operation in the logic mixer. The mixer is realized using CMOS pass-transistor logic that effectively performs the XNOR and XOR operations needed for multiplication of the data bits with the RF carrier. The output $m_{i,p}$ is at logic high when both LO and d_i are at the same logic levels, either logic high or low. The same is true for $m_{i,m}$ but with $\overline{d_i}$ replacing d_i. The inverters at the output provide a buffer between the mixer and the current-steerer in the next stage.

The current steerer, shown in Figure 5.2(b), accepts the signals $m_{i,p}$ and $m_{i,m}$ as differential inputs to the transistors M_1 and M_2. The magnitude of the output current from the cell is set by the bias of the tail current source transistor, M_3. The cascode devices M_4 and M_5 ensure that the output impedance of the current cell remains high by shielding the input transistors, M_1 and M_2, from the output node. The cascode devices are realized using thick oxide transistors, since their bias voltage, V_{cas}, can exceed the nominal rating of 1.2V. When the input $m_{i,p}$ is high and $m_{i,m}$ is low, the transistor M_1 turns on and M_2 is off. The entire current I flows through M_1, which

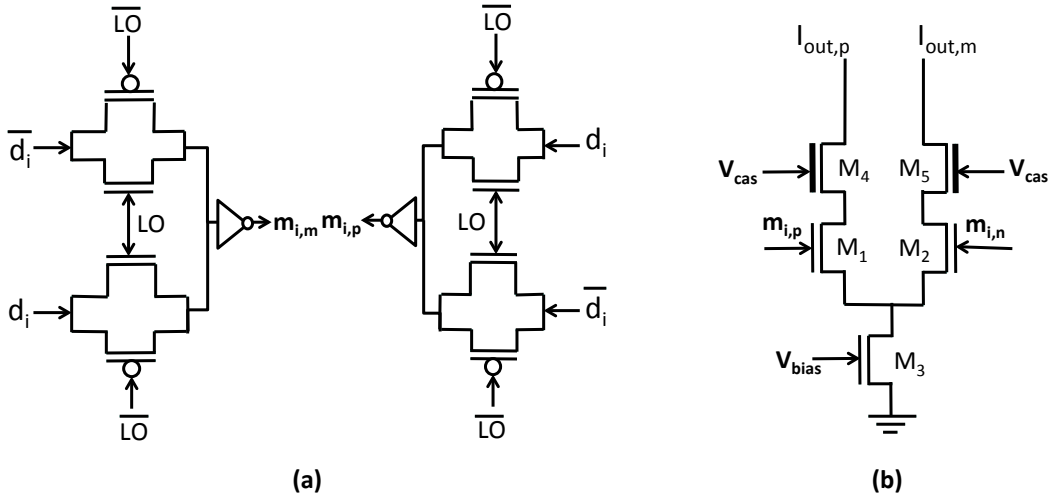

Figure 5.2: (a) Schematic of logic mixer (b) Current steering operation.

corresponds to $I_{out,p} = I$ and $I_{out,m} = 0$. The opposite scenario of $I_{out,p} = 0$ and $I_{out,m} = I$ occurs when $m_{i,m}$ is high and $m_{i,p}$ is low. Since the DAC cell output is in current form, the current contribution from different cells can be added by shorting their output branches together.

The weighted contributions of the individual binary bits of the SDM output $\{b_4,b_3,b_2,b_1,b_0\}$ can be combined within the DAC using different approaches. The straightforward method is to implement a binary-weighted DAC with 5 current cells, C_4, C_3, C_2, C_1, C_0, as shown in Figure 5.3, each driven by the corresponding bit, b_i. The magnitude of current from cell C_i is linearly dependent upon the weight of the bit b_i in the input word. This implies that the output current magnitude from cell C_0 to C_4 will have to be binary weighted, as shown in Figure 5.3. From the figure we can see that when the input b_i is high, then the current in cell C_i is steered onto the branch I_{out}, otherwise the current sinks to ground. When the current from all the cells combine, we obtain an output current proportional to the input binary word in the I_{out} branch.

The behavior of the current cells when generating their corresponding output currents can be non-ideal. This implies that the current from each cell consists of the desired current I-32I, and a random variable Δ_i, which is a result of implementation specific non-idealities like mismatches in the layout, process variations, and nonuniform voltage distribution. These non-idealities make the output current deviate from its desired value. If all the current sources are ideal then the random variables Δ_0-Δ_4 are zero. But the probability of mismatches between the current sources is high

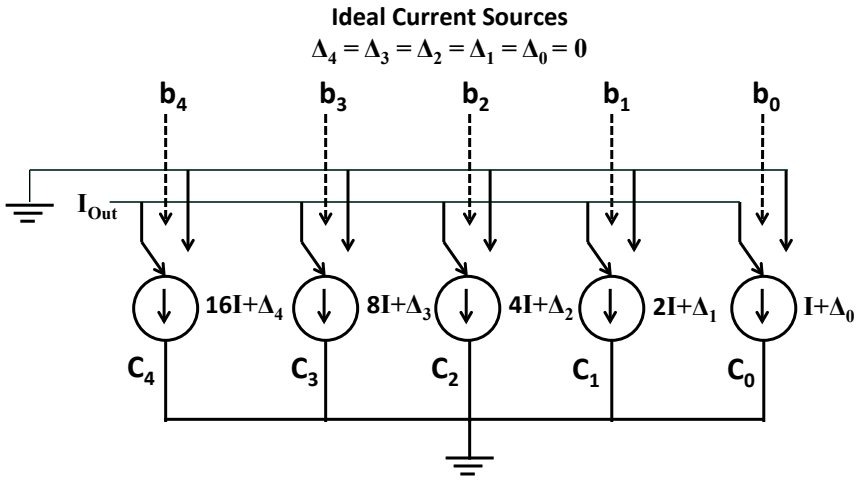

Figure 5.3: Binary-weighted DAC topology with unequal weighted current sources.

and since the cells are not identical to begin with [Gal10], the random variables Δ_i do not have the same statistical properties, leading to poor matching between the current cells. This ends up affecting the linearity of the converter, resulting in poor ACPR and EVM (error vector magnitude) performance at the transmitted output.

A second solution is to reduce the probability of mismatches between the converter cells by using unary current cells. This implementation corresponds to an array of identical unit cells, each producing a fixed current, I, that are combined together by shorting the output current from all the cells. For this topology, the binary word, $b_4b_3b_2b_1b_0$, is converted to a unary word $u_{19}u_{18}\ldots u_0$ using thermometer decoding. We use 20 bits to represent the unary word since the SDM output lies between -10 to $+10$ with a quantization interval of 1. The unary bits represent a value of +0.5 when set (logical value of 1) and a -0.5 when unset (logical value of 0). This scheme ensures that the 20 bits capture all integer values between -10 to $+10$ using differential current outputs from the RFDAC. As shown Figure 5.4, the unary bits drive the individual cells in the current cell array. The current is steered towards the branch $I_{out,p}$ when bit u_i is set, and towards $I_{out,m}$ when unset. The final voltage output is obtained by taking the difference of both currents across the output load.

The possibility of mismatches between unary cells is smaller than the binary-weighted case since the cells are identical in design. Also, the random mismatch variable Δ_i will have nearly

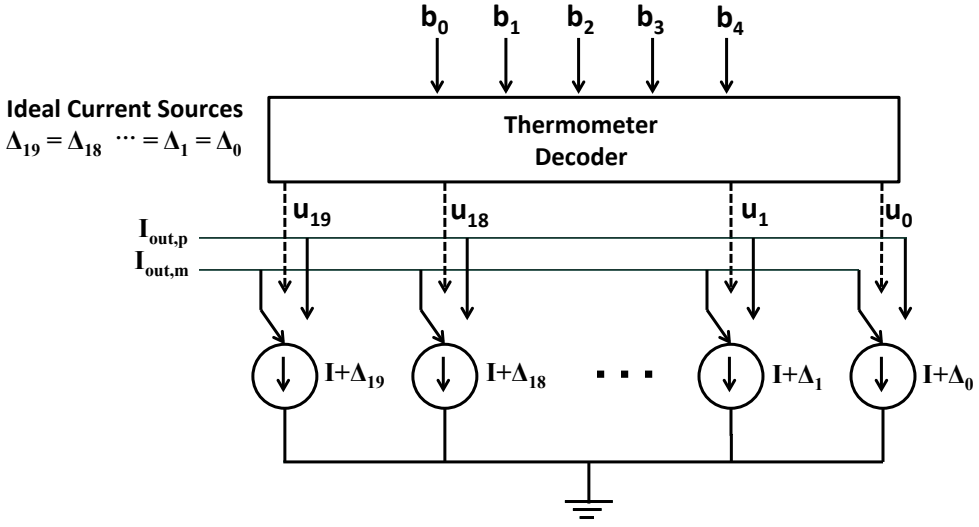

Figure 5.4: Unary DAC with equal weighted current sources.

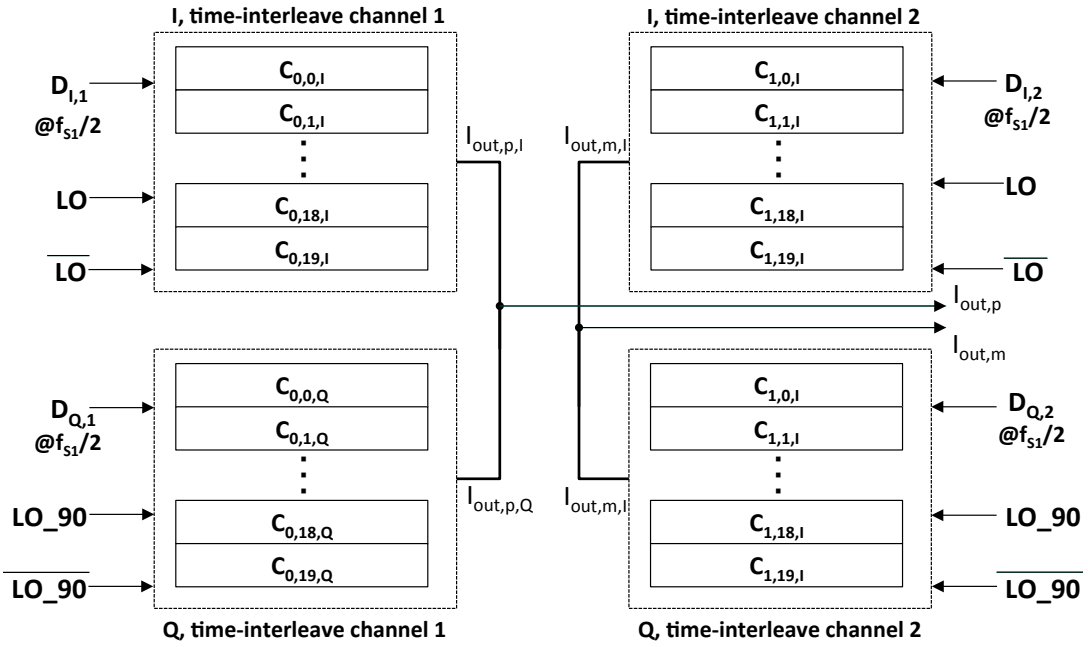

Figure 5.5: Time-interleaved RFDAC architecture with I and Q channels.

identical statistical properties across all the cells. We chose to use the unary topology because of its improved performance when it comes to matching and also because the number of quantization levels in the DAC input was small. The RFDAC architecture used in our implementation, shown in Figure 5.5, has 4 sub-units, 2 for the time-interleaving in the I channel and another 2 for the Q channel. Each of the 4 sub-units has 20 unit cells to accommodate the dynamic range of the input $D_{i,j}$, clocked at $f_{s1}/2$. The unit cell $C_{i,j,I/Q}$ in Figure 5.5 refers to the j^{th} (j = {0...19}) cell in the i^{th} (i = {1,2}) time-interleaved channel of the I or Q channel. The current from all the cells is summed by shorting the outputs to generate the final differential output $I_{out,m}$ and $I_{out,p}$. The LO and $\overline{LO}$ signals are input to the I channel while the 90-degree phase shifted version, LO_90 and $\overline{LO\_90}$, are used by the Q channel for RF modulation.

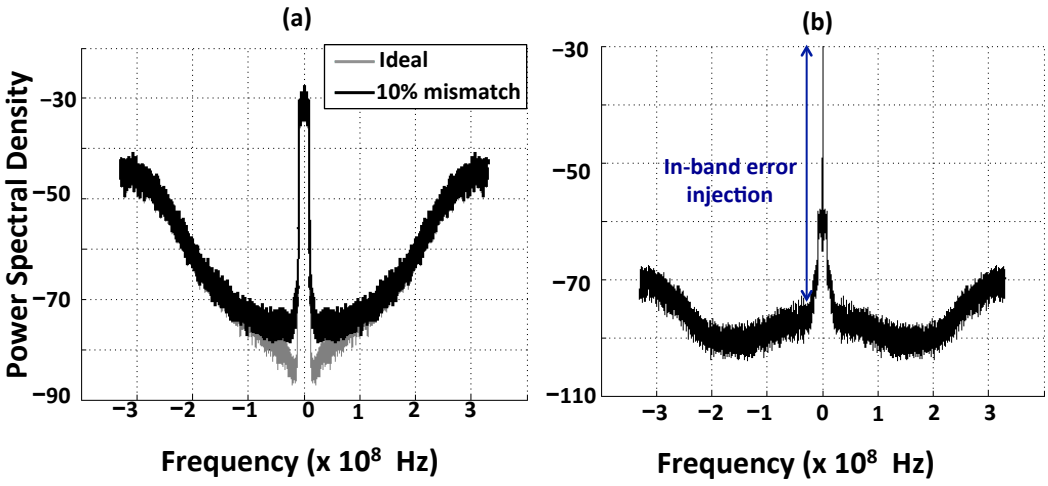

Figure 5.6: Effect of mismatches in the unary current sources of the RFDAC.

5.1 Dynamic Element Matching (DEM)

A study on the effect of mismatches between the current cells on the DAC output spectrum was done by modeling the current steering cells in MATLAB and incorporating zero-mean Gaussian random variables Δ_i into each cell. The resultant output spectrum with 10% mismatches and the corresponding error spectrum is shown in Figure 5.6. We find that the mismatches introduce both in-band noise as well as ACPR degradation in the spectrum. The in-band noise affects the EVM (error vector magnitude) performance of the transmitter. The noise shaping characteristics around the transmitted signal became worse with the notch at the baseband getting shallower as the extent of mismatches increase.

We used dynamic element matching techniques [Gal10] after thermometer decoding of the binary input in order to reduce the performance degradation that results from static current mismatches. Dynamic element matching is based on the principle of altering the usage pattern of the DAC elements so that the resultant error becomes less correlated with the input data sequence. In order to understand this concept, we take an example of how static mismatches introduce errors in the DAC output. Figure 5.7 shows the current distribution of 4 unit cells in a data converter. Each of these cells must ideally produce a current of 0.5 unit that is steered either in the direction of $I_{out,m}$ or $I_{out,p}$ depending on the unary control inputs, u_i. The actual current deviates from this

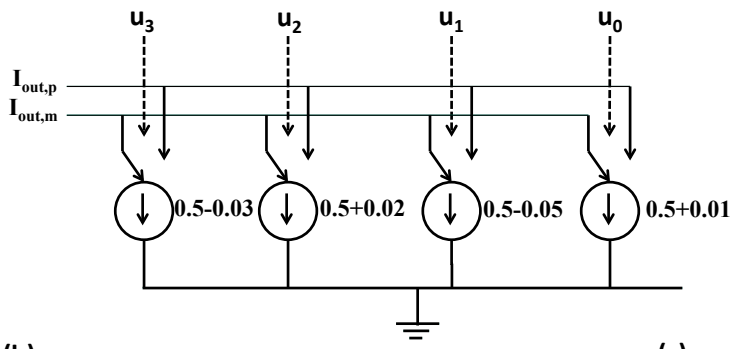

(b)

u_0	u_1	u_2	u_3	$I_{out,p}$	$I_{out,m}$	$I_{out,p} - I_{out,m}$	Ideal O/P
1	1	1	0	1.48	0.47	1.01	1
1	1	1	0	1.48	0.47	1.01	1
1	1	0	0	0.96	0.99	-0.03	0
1	1	0	0	0.96	0.99	-0.03	0

(c)

u_0	u_1	u_2	u_3	$I_{out,p}$	$I_{out,m}$	$I_{out,p} - I_{out,m}$	Ideal O/P
1	0	1	1	1.5	0.45	1.05	1
1	1	1	0	1.48	0.47	1.01	1
1	1	0	0	0.96	0.99	-0.03	0
1	0	0	1	0.98	0.97	0.01	0

Figure 5.7: (a) Non-ideal model of the current-steering DAC (b), (c) Table showing ideal and non-ideal outputs of the DAC for identical input stimuli.

value of 0.5 due to implementation related mismatches. The table in Figure 5.7(b) shows an example of test patterns for the vector $\{u_0, u_1, u_2, u_3\}$ along with the ideal and actual output in each case. The input pattern of $\{1, 1, 0, 0\}$ results in a highly correlated error pattern of $\{0.01, 0.01, -0.03, -0.03\}$. Spectrally, the error will be concentrated in the signal bandwidth, leading to EVM degradation. This effect was observed in Figure 5.6(b) where we saw that the error introduced in the spectrum was centered around the input baseband signal.

Dynamic element matching techniques take advantage of the fact that all the unit cells in a unary DAC contribute equally to the output current and therefore altering the input pattern will not make a difference with the output value as long as the number of 1's and 0's in the input vector remain unaltered. An example of this scenario is shown in the table in Figure 5.7(c), where both the first and second input vector ideally represent an output value of 1. For the first test vector, cells u_0, u_2, and u_3 are 1 while u_1 is 0. This results in an output of 1.05. For the second vector, cells u_0, u_1, and u_2 are 1 while u_3 is 0, resulting in an output of 1.01. The error pattern of $\{0.05, 0.01\}$ is not directly correlated with the input $\{1, 1\}$ due to change in cell selection. For the same vectors in Figure 5.7(b), the output was constant at 1.01 since the same set of cells were used as 1s and 0s

for both vectors. Hence, by altering the usage of the unit cells, we have reduced the correlation of the error pattern from that of the input.

A DEM encoder takes the unary vector from the thermometer decoder and rearranges the input pattern every cycle while maintaining the same numbers of 0s and 1s. There are several ways the input can be altered, each associated with a distinct spectral distribution of the error. For example, random scrambling using a pseudorandom generator results in a uniform distribution of the error across the entire frequency band of f_{s1}. In this work, we chose the data weighted averaging (DWA) approach [Bai95] since it is easy to implement and achieves 1st-order noise shaping of the error introduced by the mismatches. The schematic of this approach is shown in Figure 5.8. The DWA scrambler uses an index pointer, $Ptr(n)$, generated by an integrate operation on the input, $x(n)$. The input $x(n)$, ranging in value from $-L$ to L, is first converted to a positive quantity, $v(n)$, by addition of L, as shown in the schematic. The signal $v(n)$ then passes through an integrator, the output of which is clipped to $2L$ in order to generate the pointer $Ptr(n)$. This clipping is done since the number of bits in the unary vector is $2L$, which in our case was 20. The pointer, $Ptr(n)$ is input to a barrel shifter which circular shifts the unary input vector by the number of bits set by the index pointer. A proof of this approach leading to 1st-order noise shaping was presented in [Bai95].

The spectrum of the baseband data after dynamic element matching is shown in Figure 5.9(a). The spectrum has been compared with the spectrum obtained without the use of matching techniques when a 10% mismatch exists between the current cells. We find that the ACPR improves

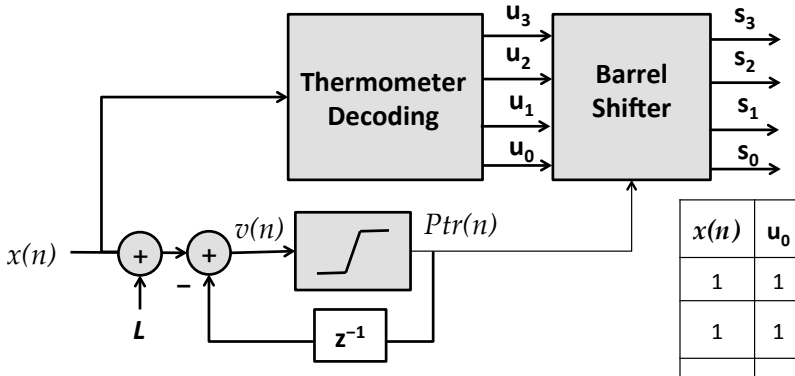

x(n)	u0	u1	u2	u3	s0	s1	s2	s3
1	1	1	1	0	1	1	1	0
1	1	1	1	0	0	1	1	1
0	1	1	0	0	0	0	1	1

Figure 5.8: Data weighted averaging approach for dynamic element matching.

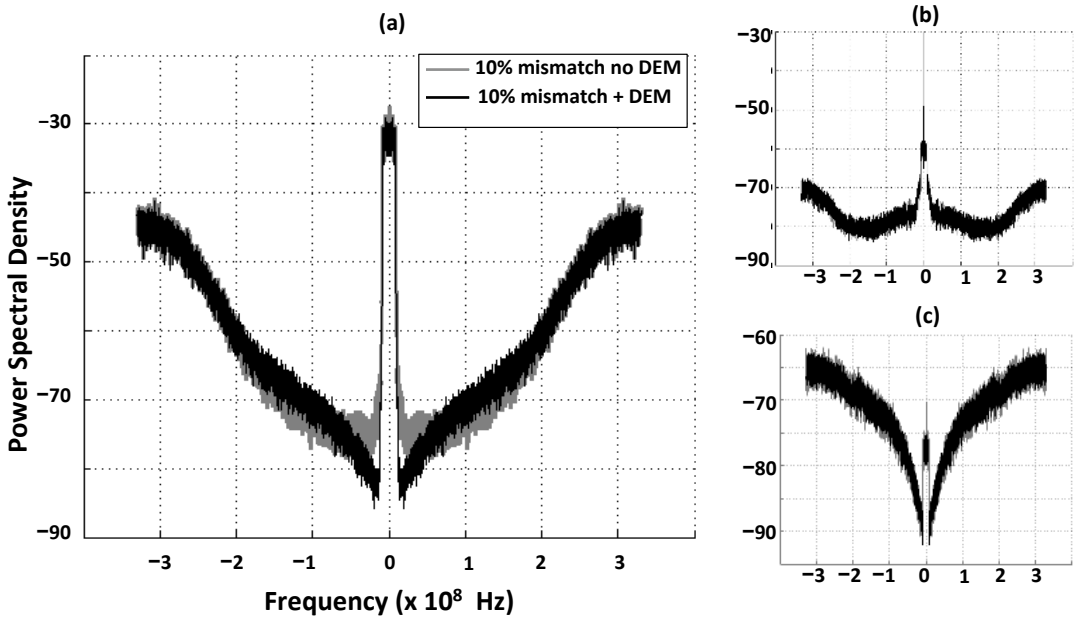

Figure 5.9: (a) Effect of dynamic element matching (b) Error spectrum without matching (c) Error spectrum after matching.

by approximately 8dB after the use of data weighted averaging. The error spectrum after dynamic element matching, shown in Figure 5.9(c), is 1st-order noise shaped. This spectrum is a significant improvement over the error spectrum in Figure 5.9(b) where DEM was not employed. Not only is the DC offset attenuated, but the average noise level in the signal bandwidth reduces from 60dB to 75dB.

Now that we have looked at the block diagram and circuit components of the RFDAC, in the next section we look at the waveforms and output spectrum generated after its circuit-level simulations.

5.2 Simulation Results

The RFDAC unit cell was simulated at a clock rate of $f_{s1}/2$ = 450Ms/s corresponding to the maximum f_{RF} of 2.7GHz. The simulated waveforms are shown in Figures 5.10 (a) and (b). The data input, D, triggers at the negative edge of the clock signal, Clk, and is latched at its positive edges. The latched data is then mixed with the LO and $\overline{\text{LO}}$ signals before being converted to analog cur-

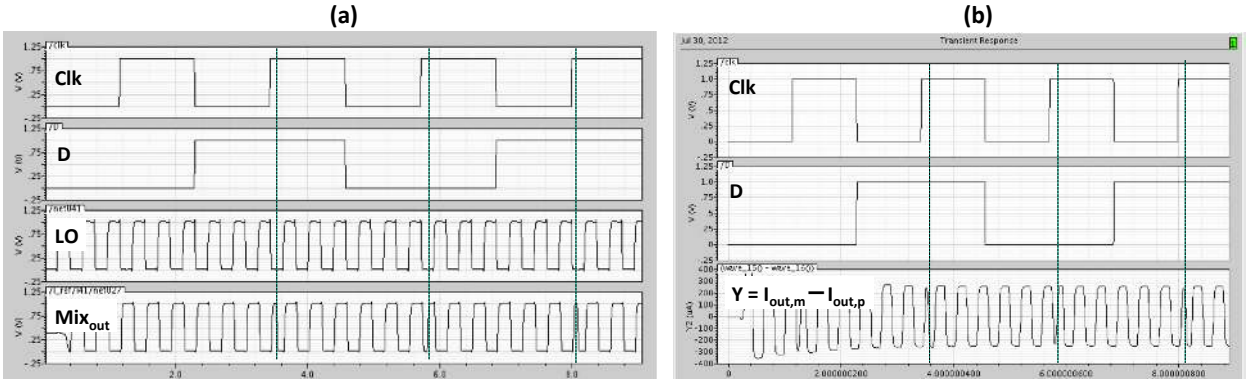

Figure 5.10: (a) Simulated waveforms from mixer output (b) Waveforms from current steerer output.

rent waveforms $I_{out,p}$ and $I_{out,m}$. The difference between the differential currents $I_{out,p}$ and $I_{out,m}$ is the final output, Y, that captures the RF modulated data. The mixing operation of the LO signals with the data can be observed from the phase shifts in the output, Mix_{out}, in Figure 5.10(a). The phase shifts are highlighted by the dashed lines in the figure. These phase shifts result when the signal D changes polarity from its previous state. As required by the mixing operation, the output should be positive when both LO and D are of the same polarity, and negative otherwise. This effect is observed in the simulated waveforms in all three cases of the data changing polarity. In the first case, the LO signal is low and D changes from low to high. The output signal, Y, hence changes from a high to low since D and LO are of opposite polarity. The opposite scenario occurs in the second case when D changes from a low to high and LO is low. The output, Y, changes from a low to high since D and LO are of the same polarity. The output Mix_{out} and its complement are then sent to the differential inputs of the current steerer, which was shown in Figure 5.2(b). The final output Y captures the changes in signal D in the form of phase shifts in the current waveform. These phase shifts are highlighted via dashed lines in Figures 5.10(b).

The output current from all the DAC unit cells are summed up by shorting their individual $I_{out,p}$ and $I_{out,m}$ branches. The final output is obtained by taking the difference of the currents flowing in the two branches. The spectrum of the RF modulated data after processing by the RFDAC is shown in Figures 5.11. This spectrum was obtained after simulating the time-interleaved RFDAC with a continuous vector of 64QAM baseband modulated data. The RFDAC up converts the baseband data to the RF carrier while also performing D/A conversion in the form of an analog output

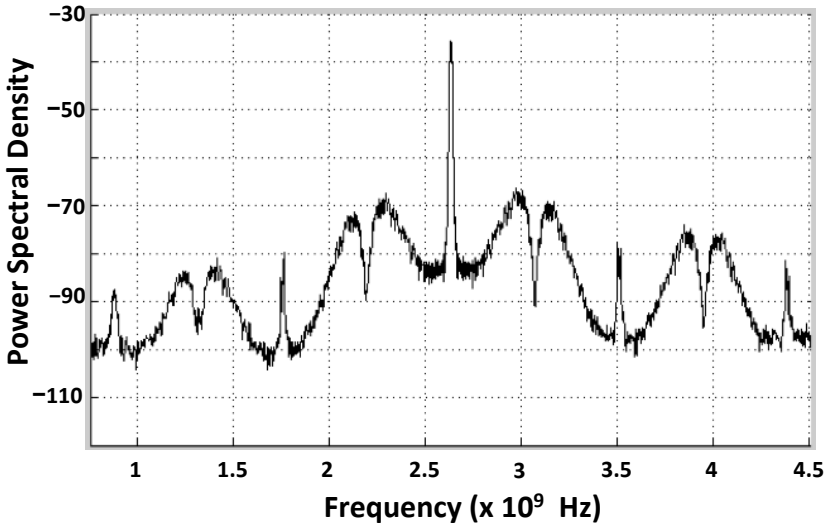

Figure 5.11: Spectrum obtained after circuit simulations of the RFDAC.

current. The spectrum deviates from the expected ideal spectrum we saw in Chapter IV and we find that the notches have flattened out slightly after the processing in the RFDAC. This results in a loss of 2dB in the ACPR which reduces to 46dB from the ideal expected value of 48dB.

Now that we have looked at the design of the receiver and transmitter chain, we will discuss the measured results obtained from the prototyped architectures in the next chapter.

CHAPTER 6

Experimental Results

The receiver and transmitter architectures discussed in the previous chapters were prototyped in 65nm CMOS technology and we look at their measured performance in this chapter. On the receiver side, the prototyped digital front-end (DFE) accepts high-speed digitized RF data, and performs mixing, down-conversion, and filtering before readying the data for handoff to the modem. The signal processing steps in the chip are the same as outlined in Chapter III. The ADC data was generated using MATLAB models and fed to the DFE chip. The DFE was implemented in an active area of 0.4mm^2. The chip photo is shown in Figure 6.1(a). The DFE was made programmable through tuning parameters that were sent to the chip via scan chains. An FPGA board was used to provide closed-loop testing of the chip with maximum flexibility. The test setup, shown in Figure 6.1(b), is based around the ROACH FPGA board, which is a general-purpose instrumentation platform [ROA]. ROACH provides a testing environment [MCR07] where MATLAB routines are coupled with an on-board embedded CPU for run-time access to shared memory on the FPGA.

The Xilinx Virtex-5 FPGA on the ROACH provides 80 high-speed differential GPIOs on a Tyco Z-Dok+ board-to-board connector. The high-speed input signal clock and the low-speed baseband clock were both provided by the FPGA to the chip. Test vectors representing ADC data

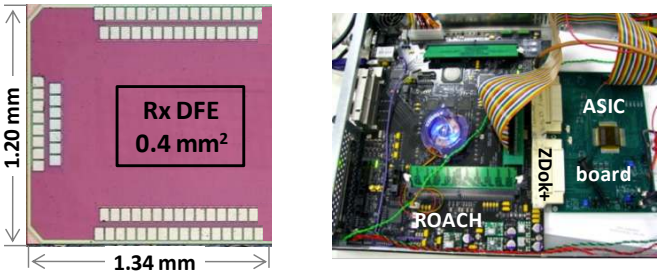

Figure 6.1: (a) Rx DFE die photo (b) Measurement setup.

and scan chain configuration bits were loaded into the QDR memory via MATLAB. The vectors were then read from the FPGA using the high-speed clock and output to the chip. For signal integrity considerations, this high-speed data uses low-voltage differential signaling (LVDS) from the FPGA through the Z-Doks before being converted into single-ended signals on the ASIC test board. Decimated data from the chip was sent on the low-speed clock back to the FPGA, where it was buffered in on-chip BRAMs. The results were then read into MATLAB for analysis.

In the implemented DFE, the first-stage CIC takes 8 parallel data streams as input. This reduces the sampling frequency of each stream to $(4/3)f_{RF}/8 = f_{RF}/6$, which is in the range of several hundred MHz (450MHz, for a 2.7GHz RF input). The input data stream had to be parallelized due to speed constraints imposed by the FPGA testing infrastructure. 40 input I/O pads were used to support 8 parallel channels of input data. These pads would not be required when the DFE takes the input from an on-chip ADC. In the next section we look at the measurement data obtained from the chip.

6.1 Rx DFE Measurements

The Rx DFE was tested with a 16MHz bandwidth input modulated at 2.025GHz. The digitized RF data, shown in Figure 6.2(a), was obtained from the ΣΔ ADC model discussed in Chapter II. Figure 6.2(a) shows the close-view spectrum of an undersampled sinusoid modulated at 2.025GHz. The sampling frequency for this sinusoid is 2.7 GHz ($(4/3)f_{RF}$). Figure 6.2(b) is the spectrum of the signal obtained from the chip, after decimation to a baseband frequency of 20MHz. The DFE dissipates 14mW of power with a maximum noise figure of 2.2dB in this experiment. Quantization of output to 12-bit words limits the SNR performance at the MODEM input to a maximum of 72dB. Although the chip was designed to process RF carriers up to 2.7GHz, the speed limitations of the FPGA and LVDS receivers on the PCB limited the testing to 2.025GHz. Power consumed for processing signals at different RF carriers and various LTE modem frequencies are shown in Figure 6.3. As expected, the power consumption increases with increasing RF carrier, since the extent of decimation increases. Also, the power levels increase with higher baseband frequency because components close to the modem operate at higher rates. The dashed line on the graph

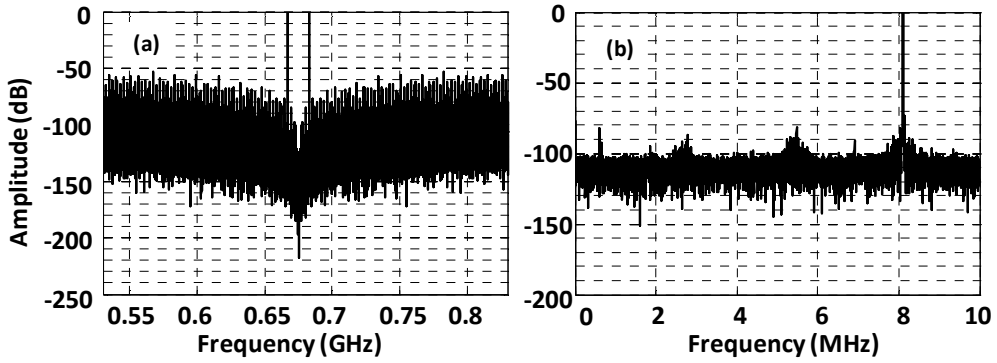

Figure 6.2: (a) FFT of 16 MHz bandwidth input modulated at 2.025 GHz after undersampling. (b) FFT of 20 MHz sampled output sinusoid from chip.

represents extrapolated power at f_{RF} = 2.7 GHz. We were unable to verify this measurement due to throughput restrictions on the PCB. From the extrapolated results we found that the maximum expected power consumption for the LTE sampling frequency of 30.72MHz and RF carrier of 2.7GHz is less than 20mW.

A summary of the measured results is presented in Table 6.1. The Table compares the performances of the proposed DFE with that of the low sample-rate DFE in [HZS08]. The work in [HZS08] proposes a conventional receiver with analog mixers and lowpass filters followed by an ADC operating at a fixed sample rate of 104Ms/s. The ADC is followed by a DFE that decimates and filters the baseband data to bring it to the modem sample rate. This DFE implementation

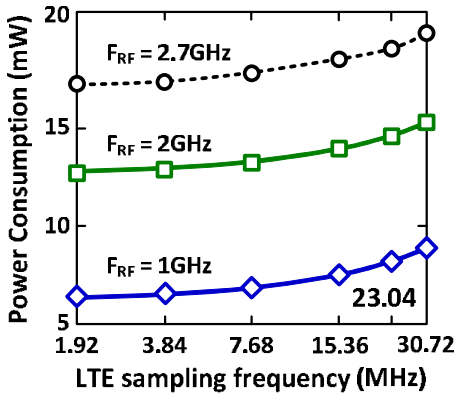

TABLE 6.1: CHIP PERFORMANCE SUMMARY

Reference	[9]	Our work
Technology	130 nm	65 nm
V_{DD}	1.5 V	1.0 V
Area	N/A	0.4 mm^2
Max Fs/f_{RF}	104 MHz / 2.1 GHz	2.7 / 2.025 GHz*
Max BW	5 MHz	20 MHz
Max I(V_{DD})	21 mA	14 mA
Noise figure	9.2 dB (AFE)	3 dB

* Restricted to 2.025 GHz due to testing infrastructure limitations.

Figure 6.3: Power consumption for LTE sampling frequencies. Solid lines: measured, dashed line: estimated.

was therefore designed to process baseband signals and low decimation factors (4-96), as opposed to our implementation that processes RF signal and supports larger scale decimation (factors between 32-2048). The power consumption for both implementations is comparable. The total power consumption of the receiver in [HZS08] varied between 60-150mW depending upon its mode of operation. This means that the direct-sampling ADC has to be implemented at power levels between 50-100mW for this approach to be comparable to state-of-the-art receiver performance.

6.2 Transmitter Implementation

The RF transmitter front-end implementation included the digital front-end, the ΣΔ modulator and the RFDAC, the components being integrated in 65nm CMOS technology. A schematic of the test setup is shown in Figure 6.4. The current output from the RFDAC was tapped by the 100Ω differential load of an RF balun. The unbalanced impedance of 50Ω in the RF balun captured the final output in voltage form, which was sent to a spectrum analyzer for measurement. A capacitive matching network parallel to the RF balun ensured maximum power transfer to the balun. The capacitance was tuned to ensure that the impedance seen at the output of the chip approximately equals the impedance of the matching network and balun at various RF frequencies.

The balun output was observed on a spectrum analyzer across a 50Ω load. The die photo of the transmitter is shown in Figure 6.5(a), with the respective components marked in the boxes. The RFDAC area was dominated by the current sources, which were implemented using 1um

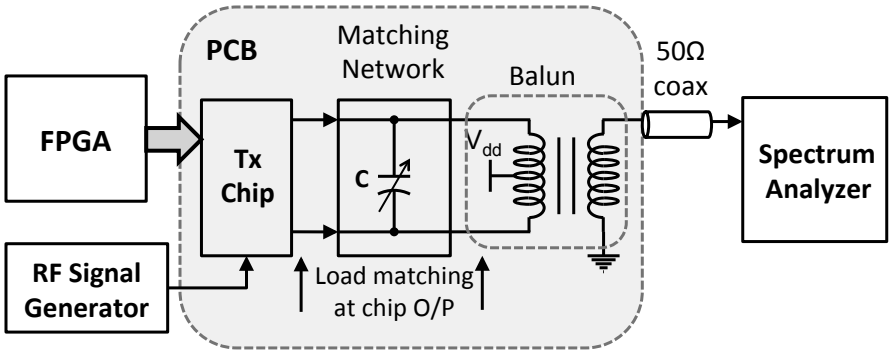

Figure 6.4: Setup used for testing the Tx chip.

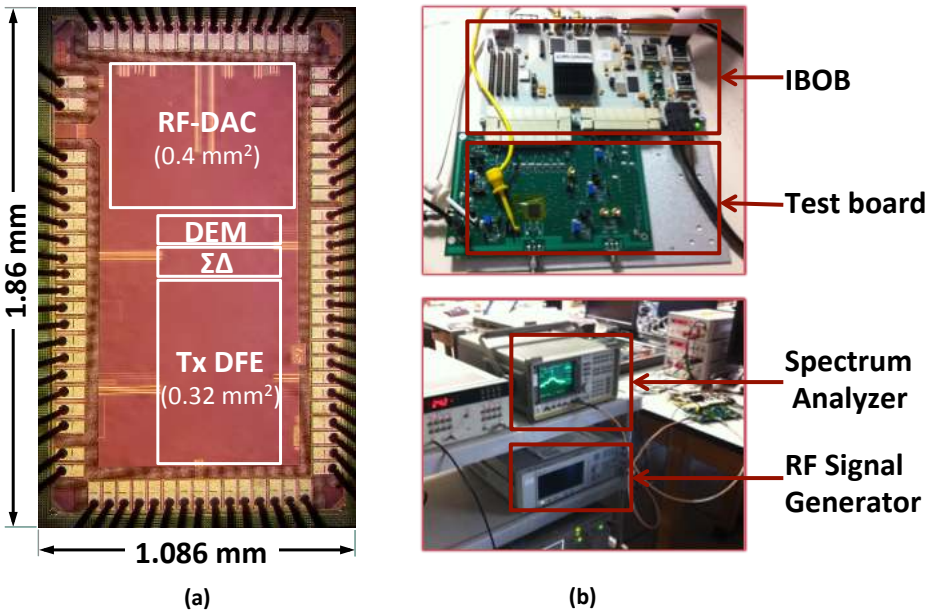

Figure 6.5: (a) Tx DFE die photo (b) Measurement setup.

long transistors. The RFDAC area can be reduced if regular transistors are used in the current sources, however, this comes at the expense of increased mismatch between the sources since the mismatches are inversely proportional to the area of the current sources. The core digital area is dominated by the interpolation filters in the DFE, while the ΣΔ modulator occupies a moderate area of 0.1mm^2.

6.2.1 Measured Results

Figure 6.5(b) shows the experimental setup used to test the RF chip. The top figure shows the setup of the FPGA interface from the IBOB platform to the PCB test board. The IBOB is a trimmed down version of the ROACH platform that was used to test the Rx DFE. The baseband modulated data is stored on the FPGA memory and then sent to the chip via LVDS signaling. The lower figure shows the connection between the chip output and the spectrum analyzer through SMA connectors. The LO signal, f_{RF}, is generated by an RF signal generator and transmitted to the PCB, also using SMA connectors.

For testing the chip, a digitally modulated, 5MHz, 64QAM SC-FDMA baseband signal with

a peak-to-average ratio of 8dB was sent to the chip input. The baseband signal was upconverted to 2GHz by the Tx modulator, and the result of the modulation is shown in the spectrum analyzer output of Figure 6.6. The ACPR performance was recorded at 42.5dB in the adjacent channel in Figure 6.6(a). The far out spectrum in Figure 6.6(b) shows the notches at frequency $f_{s1}/2$ = 333MHz, as is expected due to the time-interleaved DAC implementation. The far out spectrum also contains tones at intervals of 83MHz, which is a result of the digital noise leaking into the analog spectrum. These tones can be suppressed by ensuring greater isolation between the digital components and the analog output on the chip as well as the PCB. The measured error vector magnitude performance was at 2.8% for this experiment. The EVM measurement was obtained by capturing the RF output on a 20Gs/s sampling oscilloscope, followed by demodulation and decoding in MATLAB. The maximum output power recorded at the spectrum analyzer was -8.1dBm in the 5MHz transmitted channel. The output power can be tuned to a lower value by varying the bias voltages of the current sources in the RFDAC. This bias voltage was set by the output of a linear regulator on the PCB. The noise level at the Rx band 190MHz away (LTE band 1) was at -136dBm/Hz for the -8.1dBm output signal. For a $+20$dBm signal at the PA output, this translates to a noise level of -108dBm/Hz. This implies that the BPF following the RFDAC will need to suppress the noise level by 22dB in order to meet the target requirement of -130dBm/Hz in the Rx band. This can be achieved by the 6^{th}-order (3-pole) BPF filter with a Q of 12.5, as was discussed in Chapter IV.

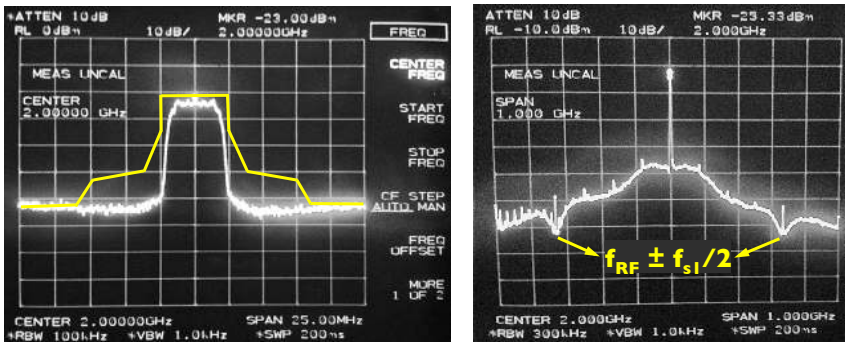

(a) (b)

Figure 6.6: (a) Close view spectrum of 5MHz bandwidth signal modulated at 2GHz (b) Far out spectrum spanning 1GHz.

The nonlinearity of the RFDAC proved to be a bottleneck in the chip implementation and reduced the ACPR performance from the expected 55dB at 5MHz to 42.5dB. The ACPR performance observed at 20MHz bandwidth was 36dB as opposed to the expected value of 48dB. A summary of the expected design performance and the actual performance obtained from the Tx chip is presented in Table 6.2. The output power of -8.1dBm was lower than the expected power of -7dBm due to transmission loss on the PCB and SMA cables. The chip met the timing specifications by modulating the baseband data to RF carriers upto 2.7GHz at a sample rate, f_{s1}, of 900Ms/s. The maximum baseband signal bandwidth during testing was 20MHz. The measured power consumption was at 85.5mW at an RF carrier of 2.4GHz and a baseband sample rate of 30.72MHz. This value was slightly higher than the expected power consumption of 75mW obtained from Synopsys and Cadence simulations during the design phase.

A comparison of the Tx chip performance with the work in [ESK07] and [PZH08] is presented in Table 6.3. The work in [ESK07] uses a straightforward I/Q modulator without any ΣΔ modulation, using a 10-bit RFDAC for modulation instead. The work in [PZH08] is closer to this work, using a 72-level RFDAC operating at a maximum of 5.4Gs/s and 3rd-order ΣΔ modulation. The main advantage of using the approach proposed in this work lies in the reduced power consumption and reduced complexity of the RFDAC. The power consumption of 85mW is about 1.7x lower than that in [PZH08] in comparable technology and RF frequency. The bandwidth and

Specifications	**Designed**	**Measured**
ΣΔ modulator	7th-order	7th-order
DAC structure	20-level DRFC	20-level DRFC
Max sample rate (f_{s1})	900 MS/s	900 MS/s
RF carrier (max)	2.7 GHz	2.7 GHz
ACPR	48 dB @ 20 MHz BW	42.5 dB @ 5 MHz BW
EVM	< 2%	2.8%
O/P channel power	−7 dBm (64-QAM)	−8.1 dBm (64-QAM)
Power (@ 2.4GHz)	75 mW	85.5 mW
Area	0.8 mm^2	0.8 mm^2
Technology	65 nm	65 nm
Flexible f_s & f_{RF}	Yes	Yes

Table 6.2: Summary of designed and measured performance of the RF modulator

Reference	[ESK07]	[PZH08]	This work
ΣΔ modulator	-	3rd-order	7th-order
DAC structure	10-bit DRFC	6-bit DRFC	5-bit DRFC
Sample rate (f_{s1})	307.2 MS/s	5.4 GS/s	900 MS/s
RF carrier (max)	1.9 GHz	2.7 GHz	2.7 GHz
Bandwidth	5 MHz	20 MHz	5 MHz
ACPR	55 dB	42.8 dB	42.5 dB
EVM	< 2%	2.1%	2.8%
O/P channel power	−9 dBm (64-QAM)	−8 dBm (64-QAM)	−8.1 dBm (64-QAM)
Power	157 mW (@ 1.9GHz)	150 mW (@ 2.4GHz)	85.5 mW (@ 2.4GHz)
Area	1mm^2	0.8mm^2	0.8mm^2
Technology	130nm	65nm	65nm
DEM	No	No	Yes
Flexible $f_{s1,}$ f_{s2}	No	Yes	Yes

Table 6.3: Modulator performance comparison.

ACPR performance of this work suffers compared to that of [PZH08] due to the nonlinearity of the RFDAC in our implementation. This can be improved in future versions by providing greater design margin by either increasing the resolution of the RFDAC or the sample rate. The reduced power consumption and DAC resolution in this work results from the use of a lowered sample rate of 900Ms/s ($f_{RF}/3$) and an optimized ΣΔ noise transfer function, as was discussed in Chapter IV.

CHAPTER 7

Conclusion & Future Direction

In this thesis we demonstrated the advantages and feasibility of radio front-ends implemented in a purely digital CMOS fabric. A digital front-end with a reconfigurable DSP core replaces the analog circuits that perform signal conditioning operations in conventional radio transceivers. The primary challenge was to design DSP circuits and data converters that can support the processing of high-speed RF signals centered at frequencies in the 2GHz range. We presented techniques that make such high-speed digital designs feasible at low levels of power consumption. A short summary of the research contributions is presented next.

7.1 Summary of Research Contributions

- We presented a low-power digital front-end core for a direct sampling receiver. Extensive architectural optimizations were used to ensure that the DFE supports high-throughput operations necessary to process RF signals. The DFE was made fully reconfigurable to handle variable modulation schemes, signal bandwidths, and RF carriers.

- We demonstrated a prototype of the Rx DFE core, implementation in 65nm CMOS, that can process RF carriers up to 2.025GHz. The DFE core consumes 14mW of power when processing an RF signal centered at 2.025GHz.

- On the transmitter side, we presented an architecture for a low-power signal conditioning chain that includes a Tx DFE core, a $\Sigma\Delta$ modulator, and a current-steering RFDAC. This Tx chain was made fully reconfigurable to support multiple modulation schemes, signal bandwidths, and RF carriers.

- We used an iterative approach to fully optimize the chain of operations in the transmitter to lower the power consumption and DAC resolution. The approach was based on co-optimizing the bandpass filter frequency response along with the ΣΔ noise-shaping characteristics. The objective was to converge to a noise-transfer-function (NTF) in the ΣΔ loop filter that minimizes the sample rate (f_{s1}) and allows operation at moderate DAC resolution. The DAC architecture uses a time-interleaved structure, which helps reduce quantization noise in the transmit spectrum, and lowers the specifications of the bandpass filter following the DAC.

- The optimized transmitter chain was prototyped in 65nm CMOS, and was shown to process RF carriers up to 2.7GHz and signal bandwidths up to 5MHz. The Tx chain consumed 85.5mW of power to modulate a 64QAM, 5MHz bandwidth baseband signal to a carrier frequency of 2.4GHz.

7.2 Future Directions

The concept of digitally-intensive transceivers is expected to gain ground in the future as CMOS technology advances and digital processing becomes cheaper. On the receiver side, the objective will be to implement a high-speed, highly linear ADC that can directly digitize the RF signal, following which the digitized signal can be conditioned by the Rx DFE core. As noted in Chapter I, the challenge will be to realize the ADC at power levels that will make it attractive for portable applications. The low-power nature of the DFE will allow the ADC in a direct-sampling receiver to have sufficient power headroom. As technology improves, the maximum sample rate and RF carriers processed by the DFE will also improve, eventually targeting the 5GHz bands used by standards like 802.11a/n.

On the transmitter side, the challenge will lie in implementing RF modulators that can support FDD operations in Rx bands close to the transmitted channel. In the present Tx modulator implementation, the closest Rx band in FDD mode was 190MHz away from the transmitted signal. The LTE standard, however, has certain bands where the Tx and Rx duplex bands are separated by only 30MHz (LTE band 2). Such requirements significantly increase the extent of noise suppression

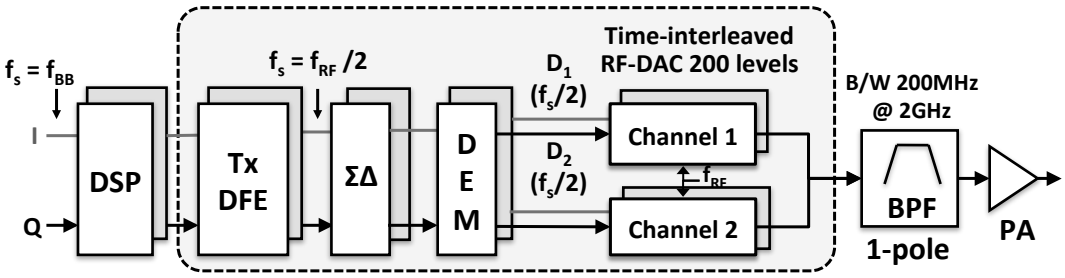

Figure 7.1: Proposed RF modulator for FDD operation in Rx bands at distance of 30MHz.

needed for FDD operation. Assuming 50dB duplexer attenuation, this translates to a maximum noise level of −130dBm/Hz in the 30MHz band away from the Tx band. For any filtering to be useful at such a small distance from the Tx band, the Q of the BPF will have to be very high, making it practically infeasible to integrate the BPF with the Tx modulator. Hence the noise suppression in the Rx bands close to the transmitted signal will have to be done in its entirety by the Tx modulator without any assistance from the BPF.

The architecture in Figure 7.1 can ensure the target noise level of −130dBm/Hz or −90dBr/MHz at a distance of 30MHz from the Tx channel. This architecture uses the same ΣΔ modulator as was used in the present Tx chip, but with a higher DAC resolution of 200 levels. The sampling frequency of the SDM and the DAC has also been raised to $f_{RF}/2$ from $f_{RF}/3$. The noise level in the 30MHz band away from the transmitted signal is −91dBr/MHz, lower than the required −90dBr/MHz, as shown in Figure 7.2(b). This solution uses a 2nd-order (1-pole) bandpass filter after the RFDAC to meet the noise specifications.

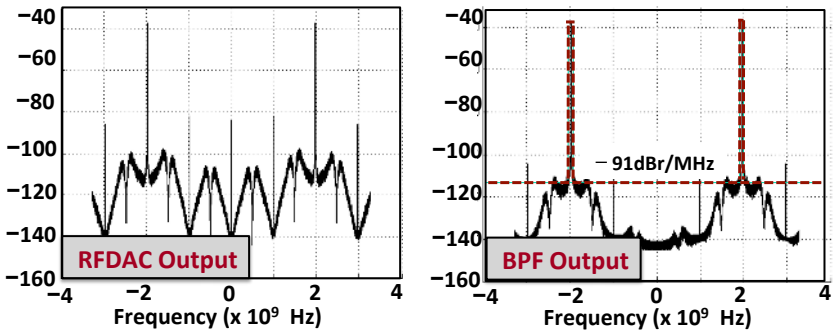

Figure 7.2: (a) Effect of 2-channel time-interleaved DAC filtering on upconverted RF signal. (b) Spectrum of RF modulated signal after 2nd-order bandpass filtering.

The sample rate of $f_{RF}/2$ for this modulator will increase the frequency of operation from 900Ms/s to 1.35Gs/s for a 2.7GHz RF carrier. This will increase the power consumed in the digital core of the modulator to above 100mW if implemented in 65nm CMOS. With more advanced CMOS fabric, the power levels will lower, thus improving the overall efficiency of the modulator. With the sample rate increasing to $f_{RF}/2$, we have the added advantage of having a more relaxed bandpass filter after the RFDAC. The implementation shown in Figure 7.1 only needs a 2^{nd}-order (1-pole) BPF with a bandwidth of 200MHz for a 2GHz RF carrier. This corresponds to a Q of 10, and the required Q will progressively decrease if we are able to raise sample rates even further, without any increase required in the DAC resolution. Hence with cheaper digital processing, it is expected that the digital content in RF front-end will steadily increase, yielding high-speed and low-power digitally-intensive transceivers.

REFERENCES

[Abi07] A.A. Abidi. "The Path to the Software-Defined Radio Receiver." *IEEE Journal of Solid-State Circuits*, **42**(5):954–966, May 2007.

[AGA04] M. Albiol, J. Luis Gonzalez, and E. Alarcon. "Mismatch and Dynamic Modeling of Current Sources in Current-Steering CMOS D/A Converters: An Extended Design Procedure." *IEEE Transactions on Circuits and Systems − I*, **51**(1):159–169, January 2004.

[Bai95] R.T. Baird. "Linearity enhancement of multibit ?? A/D and D/A converters using data weighted averaging." *IEEE Trans. on Circuits and Systems II*, **42**(12):753–762, December 1995.

[BAM09] N. Beilleau, H. Aboushady, F. Montaudon, and A. Cathelin. "A 1.3V 26mW 3.2GS/s undersampled LC bandpass ΣΔ ADC for a SDR ISM-band receiver in 130nm CMOS." *IEEE Radio Frequency Integrated Circuits Symposium*, pp. 383–386, June 2009.

[BMC06] R. Bagheri, A. Mirzaei, S. Chehrazi, M. Heidari, M. Lee, M. Mikhemar, M. Tang, and A. Abidi. "An 800MHz to 5GHz Software-Defined Radio Receiver in 90nm CMOS." *Proceedings of IEEE International Solid-State Circuits Conference*, pp. 480–481, February 2006.

[BMK11] Z. Boos, A. Menkhoff, F. Kuttner, M. Schimper, J. Moreira, H. Geltinger, T. Gossmann, P. Pfann, A. BelitzeR, and T. Bauernfeind. "A Fully Digital Multimode Polar Transmitter Employing 17b RF DAC in 3G Mode." *Proceedings of IEEE International Solid-State Circuits Conference*, pp. 376–378, February 2011.

[Bre03] G. Breed. "Bit Error Rate: Fundamental Concepts and Measurement Issues." *High Frequency Electronics*, pp. 46–48, 2003.

[CR83] R. E. Crochiere and L. R. Rabine. *Multirate Digital Signal Processing*. Englewood Cliffs, NJ: Prentice Hall, 1983.

[DSM04] E.H. Dagher, P.A. Stubberud, W.K. Masenten, M. Conta, and T. V. Dinh. "A 2-GHz Analog-to-Digital Delta-Sigma Modulator for CDMA Receivers With 79-dB Signal-to-Noise Ratio in 1.23-MHz Bandwidth." *IEEE Journal of Solid-state Circuits*, **39**(11):1819–1828, November 2004.

[ESK07] P. Eloranta, P. Seppinen, S. Kallioinen, T. Saarela, and A. Parssinen. "A Multimode Transmitter in 0.13um CMOS Using Direct-Digital RF Modulator." *IEEE Journal of Solid-State Circuits*, **42**(12):2774–2884, December 2007.

[Far88] C.W. Farrow. "A Continuously Variable Digital Delay Element." *IEEE Symposium on Circuits and Systems*, pp. 2641–2645, 1988.

[FFS09] A. Frappe, A. Flament, B. Stefannelli, A. Kaiser, and A. Kathelin. "An All-Digital RF Signal Generator Using High-Speed ΔΣ Modulators." *IEEE Journal of Solid-State Circuits*, **44**(10):2722–2732, October 2009.

[Gal10] I. Galton. "Why Dynamic-Element-Matching DACs Work." *IEEE Trans. on Circuits and Systems II*, **57**(8):69–74, February 2010.

[GNS09] V. Giannini, P. Nuzzo, C. Soens, K. Vengattaramane, M. Steyaert, J. Ryckaert, M. Goffioul, B. Debaillie, J. Van Driessche, J. Craninckx, and M. Ingels. "A 2mm^2 0.1-to-5GHz SDR Receiver in 45nm Digital CMOS." *Proceedings of IEEE International Solid-State Circuits Conference*, pp. 408–409, February 2009.

[Hay05] Simon Haykin. "Cognitive radio: brain-empowered wireless communications." *IEEE Journal on Selected Areas in Communications*, **23**(2):201–220, 2005.

[HLA09a] M.E. Heidari, M. Lee, and A.A. Abidi. "All-Digital Outphasing Modulator for a Software-Defined Transmitter." *IEEE Journal of Solid-state Circuits*, **44**(4):1260–1271, April 2009.

[HLA09b] M.E. Heidari, M. Lee, and A.A. Abidi. "An All-Digital RF Signal Generator Using High-Speed $\Delta\Sigma$ Modulators." *IEEE Journal of Solid-state Circuits*, **44**(10):2722–2732, October 2009.

[HMS05] G. Hueber, L. Maurer, G. Strasser, K. Chabrak, R. Stuhlberger, and R. Hagelauer. "Concept of a SDR compliant receive digital-front-end for cellular terminals." *Proceedings of the IEEE IASTED International Conference*, pp. 467–472, 2005.

[HSP08] C.-M. Hsu, M. Z. Straayer, and M. H. Perrott. "A low-noise, wide-BW 3.6 GHz Digital $\Delta\Sigma$ Fractional-N Frequency Synthesizer With a Noise-Shaping Time-to-Digital Converter and Quantization Noise Cancellation." *IEEE Journal of Solid-State Circuits*, **43**(12):2776–2886, December 2008.

[HZS08] G. Hueber, J. Zipper, R. Stuhlberger, and A. Holm. "An Adaptive Multi-Mode RF Front-End for Cellular Terminals." *IEEE Radio Frequency Integrated Circuits Symposium*, pp. 25–28, 2008.

[Ins08] Texas Instruments. "TMS320C6727, TMS320C6727B, TMS320C6726, TMS320C6726B, TMS320C6722, TMS320C6722B, TMS320C6720 Digital Signal Processors." *http://www.ti.com/lit/er/sprz232f/sprz232f.pdf*, October 2008.

[JS07] A. Jerng and C. G. Sodini. "A Wideband $\Delta\Sigma$? Digital-RF Modulator for High Data Rate Transmitters." *IEEE Journal of Solid-state Circuits*, **42**(8):1710–1722, August 2007.

[KRS08] K.Waheed, R.Staszewski, and S.Reze. "Curse of digital polar transmission: Precise delay alignment in amplitude and phase modulation paths." *Proceedings of the International Symposium on Circuits and Systems*, pp. 3142–3145, May 2008.

[KYW98] K.-Y. Khoo, Z. Yu, and A. N. Willson. "Efficient High-Speed CIC Decimation Filter." *Proceedings of Eleventh Annual IEEE International ASIC Conference*, pp. 251–254, 1998.

[MCB09] A. Mirzaei, S. Chehrazi, R. Bagheri, and A. A. Abidi. "A Second-Order Antialiasing Prefilter for a Software-Defined Radio Receiver." *IEEE Transactions on Circuits and Systems − I*, **56**(7):1513–1524, July 2009.

[MCR07] D. Markovic, C. Chang, B. Richards, H. So, B. Nikolic, and R.W. Brodersen. "ASIC Design and Verification in an FPGA Environment." *Proceedings of the IEEE Custom Integrated Circuits Conference (CICC)*, pp. 737–740, 2007.

[MHM05] K. Muhammad, Y.-C. Ho, T. Mayhugh, C.-M. Hung, T. Jung, I. Elahi, C. Lin, I. Deng, C. Fernando, J. Wallberg, S. Vemulapalli, S. Larson, T. Murphy, D. Leipold, P. Cruise, J. Jaehnig, M.-C. Lee, R. B. Staszewski, R. Staszewski, and K. Maggio. "A Discrete Time Quad-band GSM/GPRS Receiver in a 90nm Digital CMOS Process." *Proceedings of Custom Integrated Circuits Conference*, pp. 809–812, September 2005.

[Mit95] J. Mitola. "The software radio architecture." *IEEE Communication Magazine*, **33**(5):26–38, May 1995.

[MRM04] M. D. McKinley, K. A. Remley, M. Myslinski, J. S. Kenney, D. Schreurs, and B. Nauwelaers. "EVM Calculation for Broadband Modulated Signals." *64^{th} ARFTG Conference Digest*, 2004.

[MS04] K. Muhammad and Robert B. Staszewski. "Direct RF sampling mixer with recursive filtering in charge domain." *Proceedings of the International Symposium on Circuits and Systems*, pp. 577–580, February 2004.

[NCM11] R. Nanda, H. Chen, and D. Markovic. "A Low-Power Digital Front-end Direct-sampling Receiver for Flexible Radios." *Proceedings of the IEEE Asian Solid State Circuits Conference (ASSCC)*, pp. 377–380, November 2011.

[Nor93] S.R. Norsworthy. "Optimal Nonrecursive Noise Shaping Filters for Oversampling Data Converters." *IEEE Symposium on Circuits and Systems*, pp. 1353–1356, 1993.

[NST96] S. R. Norsworthy, R. Schreier, and G. C. Temes. *Delta-Sigma Data Converters, Theory, Design, and Simulation*. Wiley-IEEE Press, 1996.

[OGM05] M. Ortmanns, F. Gerfers, and Y. Manoli. "A continuous-time ΣΔ modulator with reduced sensitivity to clock jitter through SCR feedback." *IEEE Transactions on Circuits and Systems − I*, **52**(5):875–884, May 2005.

[OS89] A.V. Oppenheim and R.W. Schafer. *Discrete-Time Signal Processing*. Englewood Cliffs, NJ: Prentice Hall, 1989.

[PM96] J.G. Proakis and D.G. Manolakis. *Digital signal processing (3rd ed.): Principles, Algorithms, and Applications*. Prentice Hall, 1996.

[PZH08] A. Pozsgay, T. Zounes, R. Hossain, M. Boulemnakher, V. Knopik, and S. Grange. "A Fully Digital 65nm CMOS Transmitter for the 2.4 to 2.7GHz WiFi/WiMAX Bands using 5.4GHz ΔΣ RF DACs." *Proceedings of IEEE International Solid-State Circuits Conference*, pp. 360–361, February 2008.

[RBV09] J. Ryckaert, J. Borremans, B. Verbruggen, C. Armiento L. Bos, J. Craninckx, and G. Van der Plas. "A 2.4 GHz low-power sixth-order RF bandpass delta-sigma converter in CMOS." *IEEE Journal of Solid-State Circuits*, **44**(5):2873, May 2009.

[RCN04] Jan M. Rabaey, Anantha Chandrakasan, and Borivoje Nikolic. *Digital integrated circuits- A design perspective*. Prentice Hall, 2004.

[RDB08] F. Rivet, Y. Deval, J. Begueret, D. Dallet, P. Cathelin, and D. Belot. "A Disruptive Receiver Architecture Dedicated to Software-Defined Radio." *IEEE Tran. On Circuits and Systems −II*, **55**(4):344–348, 2008.

[ROA] "ROACH." *http://casper.berkeley.edu/wiki/ROACH*.

[SLE08] R. B. Staszewski, D. Leipold, O. Eliezer, M. Entezari, K. Muhammad, I. Bashir, C.-M. Hung, J. Wallberg, R. Staszewski, P. Cruise, S. Rezeq, S. Vemulapalli, K. Waheed, N. Barton, M.-C. Lee, C. Fernando, K. Maggio, T. Jung, I. Elahi, S. Larson, T. Murphy, G. Feygin, I. Deng, T. Mayhugh, Y.-C. Ho, K.-M. Low, C. Lin, J. Jaehnig, J. Kerr, J. Mehta, S. Glock, T. Almholt, and S. Bhatara. "A 24mm^2 Quad-Band Single-Chip GSM Radio with Transmitter Calibration in 90nm Digital CMOS." *Proceedings of IEEE International Solid-State Circuits Conference*, pp. 208–209, February 2008.

[SP12] P. Shettigar and S. Pavan. "A 15mW 3.6GS/s CT-?Æ ADC with 36MHz Bandwidth and 83dB DR in 90nm CMOS." *Proceedings of IEEE International Solid-State Circuits Conference*, pp. 156–157, February 2012.

[SPK92] K. Steiglitz, T. W. Parks, and J. F. Kaiser. "METEOR: A constraint- based FIR filter design program." *IEEE Trans. Signal Processing*, **40**(8):1901–1909, August 1992.

[SR73] R.W. Schafer and L.R. Rabiner. "A Digital Signal Processing Approach to Interpolation." *Proceedings of the IEEE*, **61**(6):692–702, June 1973.

[SRM03] T. Sowlati, D. Rozenblit, E. MacCarthy, M. Damgaard, R. Pullela, D. Koh, and D. Ripley. "Quad-band GSM/GPRS/EDGE polar loop transmitter." *Proceedings of IEEE International Solid-State Circuits Conference*, **47**:186–187, February 2003.

[ST04] R. Schreier and G. C. Temes. *Understanding Delta-Sigma Data Converters*. Wiley-IEEE Press, 2004.

[SWR05] R. B. Staszewski, J. Wallberg, S. Rezeq, C.M. Hung, O. Eliezer, S. Vemulapalli, C. Fernando, K. Maggio, R. Staszewski, N. Barton, M.-C. Lee, P. Cruise, M. Entezari, K. Muhammad, and D. Leipold. "All-digital PLL and GSM/EDGE transmitter in 90 nm CMOS." *Proceedings of IEEE International Solid-State Circuits Conference*, **48**:256–257, February 2005.

[TS07] B. K. Thandri and J. Silva-Martinez. "A 63dB 75-mW bandpass RF ADC at 950 MHz using 3.8-GHz clock in 0.25-um SiGe BiCMOS technology." *IEEE Journal of Solid-state Circuits*, **42**(2):269–279, February 2007.

[Vai90] P.P. Vaidyanathan. “Multirate digital filters, filter banks, polyphase networks, and applications: A tutorial.” *Proceedings of the IEEE*, **78**(1):56–93, June 1990.

[WN09] Renaldi Winoto and Borivoje Nikolic. “A Highly Reconfigurable 400-1700MHz Receiver Using a Down-Converting Sigma-Delta A/D with 59-dB SNR and 57-dB SFDR over 4-MHz Bandwidth.” *Proceedings of IEEE Symposium on VLSI Circuits*, June 2009.