MSc Thesis

A Wideband Frequency-Agile Phase Modulator for Digital Polar Transmitters

Michael Polushkin

Delft University of Technology

TU Delft Delft University of Technology

Challenge the future

A Wideband Frequency-Agile Phase Modulator for Digital Polar Transmitters

by

Michael Polushkin

to obtain the degree of Master of Science
at the Delft University of Technology,
faculty of Electrical Engineering, Mathematics and Computer Science,
to be defended publicly on January 25, 2017 at 14:00.

Student number:	4011554	
Project duration:	October 1, 2014 – February 1, 2017	
Thesis committee:	Prof. dr. ing. L. C. N. de Vreede,	TU Delft, supervisor
	Dr. ir. S. M. Alavi,	TU Delft
	Dr. ir. M. A. P. Pertijs,	TU Delft
	Ir. M. R. Mehrpoo,	TU Delft
	Ir. M. Hashemi,	TU Delft
	Dr. ir. J. Qureshi,	Ampleon
	Dr. ir. M. Acar,	Ampleon

This thesis is confidential and cannot be made public until January 25, 2018.

TU Delft

Abstract

Over the past two decades, the use of wireless communication has increased dramatically, reaching a total global data usage of several exabytes per month, as of 2015. The transmitters (TX) used in wireless communication systems typically operate at efficiencies in the order of 20–30%, which puts a large strain both on the wireless network operators, due to the high energy costs, and the environment, due to carbon emissions. Since the power amplifier (PA) is by far the largest energy consumer in any transmitter system, the development of highly efficient PAs has become a popular research topic in recent years.

One way in which the PA efficiency can be improved is by using the polar architecture, in which the signal to be amplified is split into envelope and phase components, which are presented separately at the input of the PA. This facilitates the use of non-linear but highly efficient amplifier classes, which improves the system efficiency. Further efficiency gains can be achieved by combining this approach with digital techniques as well as traditional efficiency enhancement techniques such as the Doherty topology.

One of the major challenges associated with the polar architecture is the bandwidth expansion that occurs when converting from a Cartesian representation of the signal to envelope and phase. In particular, very few phase modulators have been reported in literature that support modulation bandwidths large enough for advanced communication standards such as WiFi, WiMAX and LTE.

This thesis aims to tackle this problem by presenting a novel phase modulator concept which is both wideband enough to support 80 MHz QAM signals, and frequency-agile so as to support the numerous communication standards in the low-GHz range. The phase modulator is based on the constant-envelope direct upconversion architecture, which allows for very wideband operation. At the core of the modulator is a RF-DAC which makes use of harmonic rejection to suppress the 3^{rd} and 5^{th} harmonics at the RF-DAC output. This significantly relaxes the requirements of the output filter to the point that 2^{nd}-order RC filtering is sufficient to achieve the desired performance. In this way, the use of inductors can be avoided, which reduces the required chip area, aids in integration with a high-power PA and makes it easier to support a large range of operating frequencies. The output of the filter is then amplified to CMOS levels by a limiting amplifier in order to drive the digital polar PA. In order to benefit from technology scaling, the design is kept as digital as possible.

This concept has been implemented in TSMC's 40 nm CMOS technology, and although the performance of the phase modulator has not yet been experimentally verified, simulation results show that in a TX system with an ideal envelope path, the modulator can achieve −36 dB EVM and −47 dBc ACPR for an 80 MHz 64-QAM signal at 2.4 GHz. At 900 MHz, when the bandwidth is scaled proportionally to 30 MHz, the phase modulator achieves −34 dB EVM and −47 dBc ACPR. The power consumption for these two cases is 33.3 mW and 15.6 mW, respectively, while the occupied chip area is 0.17 mm^2.

Acknowledgements

This thesis would never have been possible without the support I received from others. For this support, the following people and groups deserve my sincerest gratitude:

The Dutch Technology Foundation STW – for providing the much-needed funding for the SEEDCOM project. These funds made it possible to conduct this project with a fantastic team of researchers, as well as to manufacture the final result.

Ampleon (formerly part of NXP Semiconductors) – for both the funding they provided to the project but also for their insight and expertise during the design process. Additionally, I would like to thank them for paying me a monthly stipend during my first year, which helped to lighten my personal financial burden.

Prof. Leonardus de Vreede – for giving me the opportunity to work on this project, and for teaching me most of what I know about RF circuit design. It was his vision and expertise that led to this project getting off the ground, and thanks to his numerous insights he was able to guide it to its conclusion. Also, I would like to thank him for listening to me complain about Cadence.

Dr. Seyed Morteza Alavi – not only for the work he did on the system-level design, but also for supporting me, personally, more times than I can count. From answering my numerous questions, to proof-reading this thesis, to encouraging my wild ideas, all the while doing this in the most kind-hearted manner, Morteza has helped me tremendously. Also, I would like to thank him for listening to me complain about Cadence.

Milad Mehrpoo MSc – to say that Milad helped me a great deal would be quite an understatement. Even though he was not involved in the project originally, he still offered to help out, not only by checking my work, but also by taking over the design and layout of the limiting amplifier as well as several other tasks. He has spent countless hours over the past two years trying to help me, to the point that he unintentionally became my daily supervisor, offering both insight and encouragement whenever it was needed. Also, I would like to thank him for listening to me complain about Cadence.

Mohsen Hashemi MSc – for helping me with all my issues related to signal processing. Even though he had his own *two* chips to tape out, he was alway willing to spend some time to help me make sense of my simulation results. It was only towards the end of the project that I had learned enough from him to be able to do the processing of both the system level and post-layout simulations on my own. Also, I would like to thank him for listening to me complain about Cadence.

Yiyu Shen Msc – for his work on both the DPA and almost all of the digital logic on the chip, including the LUTs and interpolation filters preceding my phase modulator. He has shown

true dedication to the project, sacrificing most of his weekends to get everything done, and at the time of writing he continues to do so in order to obtain the needed measurement results. This includes measurement results of the phase modulator, which, unfortunately, I did not have time to do myself. Also, I would like to thank him for listening to me complain about Cadence.

Lei Zhou Msc – for all of her work on the DPA concept and implementation, as well as her innovative off-chip matching network, which would allow this TX system to operate at an even higher efficiency. Without her contribution, my phase modulator would have nothing to drive. Also, I would like to thank her for listening to me complain about Cadence.

Atef Aknoukh – for his continued support throughout the implementation phase of this project. Given the sheer number of external connections on this chip, his work on the ESD ring must have been a great effort. Likewise, due to the large total area, getting the chip DRC-clean and ready for tape-out was no easy task. As for me, personally, he helped me to get started with layout, and was always there if I had any questions. Also, I would like to thank him for listening to me complain about Cadence.

My parents, Evgeny Polushkin and Olga Polushkina – for continuing to support me, both morally and financially, during my thesis (and of course long before that, as well). While they were not particularly happy to see this project last so long, they were still very understanding, and helped calm me down when I needed it. Also, I would like to thank them for listening to me complain about Cadence.

My loving girlfriend, Ioana Georgescu – both for being my voice of reason and for providing me with a steady stream of affection for the past several years. She has had to deal with a very grumpy boyfriend on far more occasions than I would like to admit, but she never gave up on me. Throughout the course of this project, she has supported me on intellectual, personal and financial levels, whenever I needed it. For this, I am forever grateful[1]. Also, I would like to thank her for listening to me complain about Cadence.

[1] or at least for the foreseeable future

Contents

Chapter 1

Introduction

The past two decades have seen a dramatic increase in the use of radio frequency (RF) communication, predominantly due to the popularization of mobile phones and wireless networking, to the point that it has now become indispensable in our everyday life. In 2015, a total of 7.9 billion mobile devices were in use worldwide, and that number is predicted to grow to 11.6 billion by 2020, as was published in a recent report by Cisco Systems Inc. [1]. Furthermore, with the market being driven to higher and higher connection speeds, and mobile phones now being used for purposes such as multimedia streaming and gaming, the amount of data being transferred has skyrocketed. The same report shows the total global data usage for 2015 was 3.7 exabytes per month (1 exabyte = 1 billion gigabytes) and is predicted to grow by more than 700 % in the next five years.

In a wireless system, data is transmitted over the air from a transmitter to a receiver. A block diagram of a traditional wireless system is shown in Figure 1.1. In a typical transmitter (TX) system, the digital baseband data to be transmitted is converted into an analog signal using a digital-to-analog converter (DAC), which is then passed through a low-pass filter to suppress spectral sampling replicas. This band-limited signal is then upconverted by mixing it with the local oscillator (LO) or carrier signal. The modulated RF signal is passed on to the power amplifier (PA) which amplifies it such that it can be reliably received at the other end of the wireless channel. The signal is then passed on to the antenna, allowing it to radiate into the air.

The receiver (RX) functionality is accomplished in a complementary fashion. First, a low-noise amplifier (LNA) amplifies the weak signal received by the antenna. This signal then gets mixed with the LO so that the desired frequency band appears at baseband. A low-pass filter serves to eliminate other signals that might be present in the received spectrum. The resulting signal is then sampled by the analog-to-digital converter (ADC) and passed on the the digital

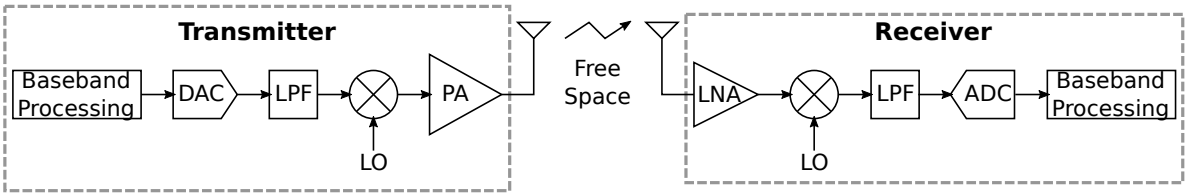

Figure 1.1: A basic RF system

processing to interpret the received information.

Of the various system blocks described above, the PA consumes by far the most power. First of all, the required power for transmission is often one or several watts (depending on the application), whereas the other components can operate in the milliwatt or even microwatt range. Secondly, PAs have to be designed to meet stringent linearity requirements, both in-band, so that the receiver is able to correctly identify the instantaneous signal values, and out-of-band, to not interfere with transmissions on neighboring frequency channels. This has traditionally required the use of linear amplifier topologies such as class-A, class-AB or class-B. These amplifier classes have maximum theoretical efficiencies of 50 %, 50 %-78.5 % and 78.5 %, respectively [2]. In a practical amplifier, the maximum efficiency is lower due to the non-idealities of real components. Furthermore, the maximum efficiency can only be achieved at the maximum output power. Since modern modulation schemes often modulate the amplitude, most of the time, the amplifier operates in power back-off, which lowers the efficiency even further. This is particularly true for advanced communication standards, such as LTE or WiMAX, which have a high peak-to-average power ratio (PAPR). Finally, if the power consumption of the preceding blocks in the TX chain is taken into account, the system efficiency in a typical system can be as low as 20–30%. [3] In the case of mobile phones, the power consumption of the transmitter directly relates to the battery life. In base stations, on the other hand, the transmitter power consumption represents a large part of the monthly operating costs [4]. It has been estimated that, globally, base stations and the backhaul networks that interconnect them consume roughly 60 billion kWh yearly, which corresponds to approximately 0.33 % of the global electricity usage [5]. As such, the development of highly efficient, yet linear, transmitters has become a very popular research topic.

1.1 Cartesian and Polar Transmitters

Numerous communication standards exist in the field of mobile telephony and wireless networking. Most of these communication standards modulate both the amplitude and the phase of the RF carrier to encode information. This RF signal can be expressed in terms of these quantities as

$$s(t) = \rho(t)\cos(\omega_c t + \phi(t)) \tag{1.1}$$

where ω_c is the angular frequency of the carrier, and $\rho(t)$ and $\phi(t)$ are the amplitude and phase, respectively. It is also possible to express $s(t)$ in terms of its in-phase (I) and quadrature (Q) components:

$$s(t) = I(t)\cos(\omega_c t) + Q(t)\sin(\omega_c t) \tag{1.2}$$

Equations 1.1 and 1.2 are known as the polar and Cartesian representations of $s(t)$, respectively. These two representations lead to two groups of transmitter topologies: polar and Cartesian transmitters.

Cartesian or I/Q transmitters have been the de-facto standard in the industry for over a decade [6], [7]. These transmitters modulate the amplitude of an in-phase and a quadrature carrier independently, and the two modulated signals are then summed at the input of the PA to produce an RF signal with arbitrary amplitude and phase. This is typically done using the architecture shown in Figure 1.2a. Because the summation is linear, the bandwidths of the I/Q components do not exceed the bandwidth of the output signal. As such, the Cartesian transmitter is able to achieve high modulation bandwidths. However, as explained above, this architecture requires the use of linear PAs, which are inherently inefficient.

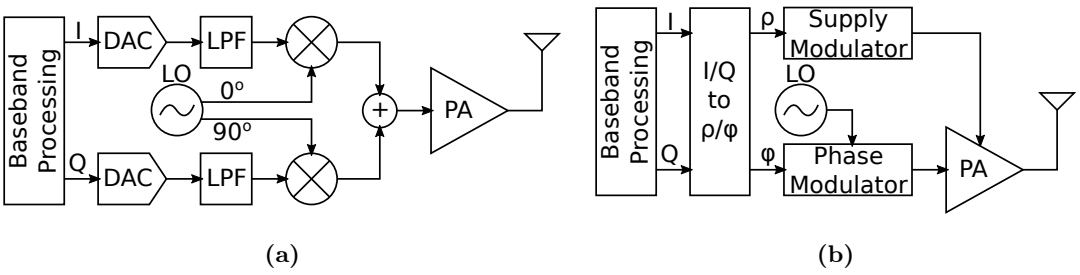

Figure 1.2: Typical **(a)** Cartesian and **(b)** polar transmitter architectures

Polar transmitters, on the other hand, operate on envelope (ρ) and phase (ϕ) information, rather than I/Q, and perform these modulations separately. These two signals are then combined within the PA to generate the output signal. Commonly, this is executed by driving the PA with a constant-envelope phase-modulated input signal, whereas the envelope information is supplied to the PA through supply modulation. This architecture is illustrated in Figure 1.2b. This method was first proposed by Khan in [8] and was termed the envelope elimination and restoration (EER) PA architecture, in which a modulated RF signal was split into its envelope and phase components and supplied to the PA in polar fashion. Nowadays, however, it is much more practical to generate the envelope and phase components directly in the digital domain. This was first proposed in [9], in which the authors presented a polar transmitter for GSM/EDGE systems.

The main advantage of the polar topology is that it can make use of a highly efficient nonlinear PA. However, since the relationship between I/Q and ρ/ϕ is nonlinear, polar transmitters suffer from bandwidth expansion. Thus, the amplitude and phase paths need to support significantly higher bandwidths than the combined signal to meet spectral requirements [10]. Another drawback of this approach is that the timing of the amplitude and phase paths needs to be very accurately aligned to avoid distortion. Because circuit implementations of supply modulators and phase modulators usually have very different time constants and propagation delays, this has proven to be a difficult task. Finally, the supply modulator, which is often implemented as a switch-mode power supply, can introduce a significant amount of noise and spurious tones in the system, further degrading the performance.

1.2 Efficiency Enhancement Techniques

Over the years, several efficiency enhancement techniques have been proposed for TX systems. One of the most popular ones is the Doherty topology [11], which is illustrated in Figure 1.3. As discussed above, any signal of which the envelope is modulated exhibits a non-zero PAPR, which severely degrades the efficiency of linear PAs. In a Doherty configuration, this effect is mitigated by employing two PAs: the main and the peak amplifiers. In a typical system the main PA is sized such that it reaches its maximum voltage swing at the average signal power so that, at this power level, this PA operates at its maximum efficiency (78.5 % for an ideal class-B amplifier). For power levels up to the average power, the peak PA is turned off. If a higher signal power is required, the main PA cannot provide it alone and thus the peak PA is turned on to provide the difference. The main PA is able to maintain operation at its maximum efficiency point for higher power levels through a technique known as active load pulling. As the peak

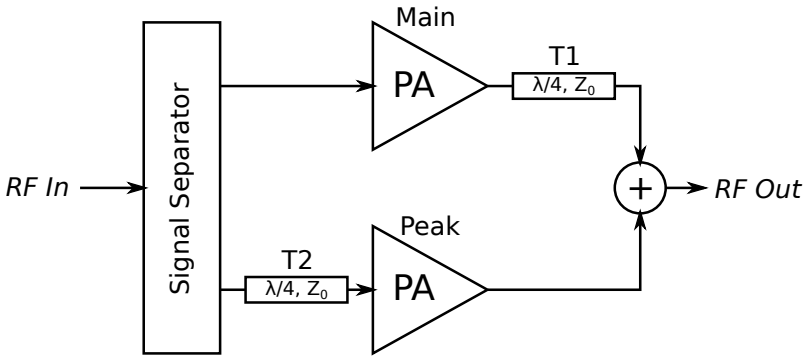

Figure 1.3: The Doherty architecture

PA turns on, its output will "pull" the load impedance (not shown in the figure) as seen by a phase-related signal. For increasing power levels of the peak PA, the effective load impedance will increase. A $\lambda/4$ transmission line (T1) inserted between the main amplifier and the load serves as an impedance inverter, so that for increasing power levels the main amplifier sees a decreasing load impedance. This allows the main PA to continuously operate at its maximum voltage swing, thus maintaining its peak efficiency. [2]

To compensate for the phase delay introduced by T1, another $\lambda/4$ line is placed before the peak PA. This ensures the correct phase relationship of the main and peak output signals at the load. The modulated RF signal needs to be split into two components to drive the main and the peak PA separately. The signal splitter can be implemented as a simple power divider if, for instance, the main PA is biased in class-B and the peak PA is biased in class-C. However, other possibilities exist, including doing the signal separation entirely in the digital domain, optionally incorporating the delay of T2 as part of the signal processing.

An alternative method of improving the efficiency of TX systems is the outphasing topology, first proposed by Chireix in [12]. A typical implementation of this topology is shown in Figure 1.4a. Like the Doherty architecture, this technique employs two PAs to generate the RF output signal. However, in this case, the RF output is produced by summing two constant-envelope phase-modulated signals: s_1 and s_2. This concept is illustrated by means of a phasor diagram in Figure 1.4b. As the figure shows, branch signals s_1 and s_2 are signals of unit magnitude, with a phase of $\phi + \theta$ and $\phi - \theta$, respectively. Here, ϕ is the desired phase of the RF output signal, whereas θ is the outphasing angle required to achieve the desired amplitude when s_1 and s_2 are summed. The fact that s_1 and s_2 are constant-envelope signals allows for the use of non-linear digital PA topologies, such as class-E or class-F, which are inherently more efficient then their linear counterparts. The input signals to the PAs are somewhat more difficult to generate than for the simple Doherty case. While it would be possible to split a modulated RF signal into the two signal components required for outphasing, this is generally considered impractical. Instead, the components are typically generated in the digital baseband and phase modulators are employed to drive the PAs.

Despite the efficiency advantages provided by the switching PAs in outphasing transmitters, this configuration has a number of disadvantages. One such disadvantage is the bandwidth expansion that the branch signals exhibit, similar to the polar architecture. Thus, the phase modulators need to support a much higher bandwidth than the original I/Q input signal. Another disadvantage is the poor accuracy of the of the output signal in deep power back-off, stemming

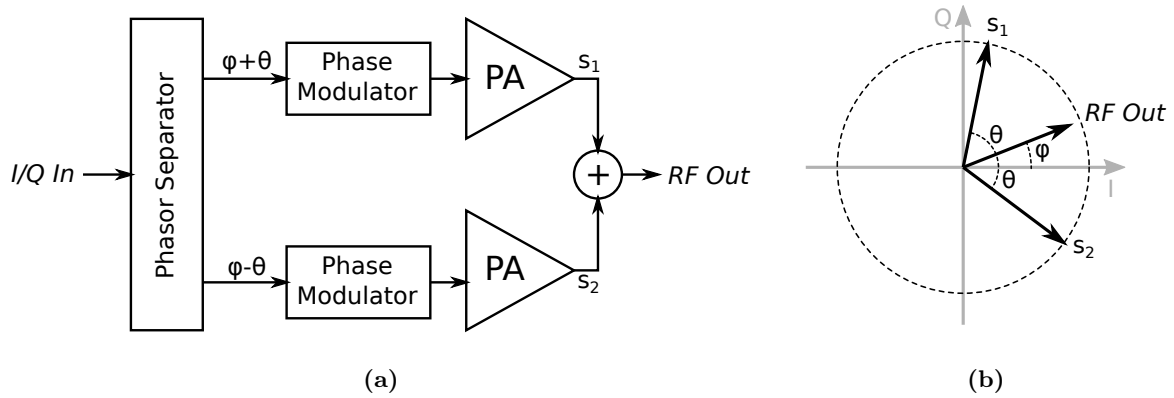

Figure 1.4: **(a)** A typical outphasing architecture and **(b)** a phasor diagram of the relevant signals

from attempting to produce a small vector by summing two anti-phase large vectors. A small mismatch between the branch vectors can, thus, have a large relative effect on the resulting output, limiting the dynamic range of the transmitter. Furthermore, when the branch signals are summed with a power combiner, their phase difference will introduce a reactive component to the load seen by each amplifier. This effect gets exacerbated in deep power back-off, which degrades the efficiency of the switching PAs. These effects have been analyzed and several solutions to some of these issues have been proposed in [13]. Nevertheless, applying this concept to wideband, high-PAPR signals, remains a challenge.

1.3 The SEEDCOM Project

This thesis work is part of the SEEDCOM project, funded by STW and Ampleon (formerly NXP), which was started to develop an advanced highly-efficient TX chain which is suitable for use in base stations. The key goal of the project is to design a TX solution which:

- has a high efficiency, both at peak power and in power back-off
- maintains a high spectral purity
- has a high RF bandwidth (i.e. can operate at a large range of carrier frequencies) to support the numerous standards in the low-GHz range
- has a high video bandwidth (or signal bandwidth) to accommodate the ever-growing demand for high-speed data transfer
- is as digital as possible, in order to take advantage of the superior digital performance and power consumption in deep sub-micron CMOS technology nodes as the analog performance continues to degrade [14].
- is fully integrated, in order to reduce both costs and form factor, allowing for cost-effective implementation of small-cell base stations

Although in the initial stages of the project the design efforts were focused on high performance outphasing concepts ([13], [15]), due to the limitations of the outphasing architecture, as discussed above, other concepts had to be considered. The chosen approach builds upon

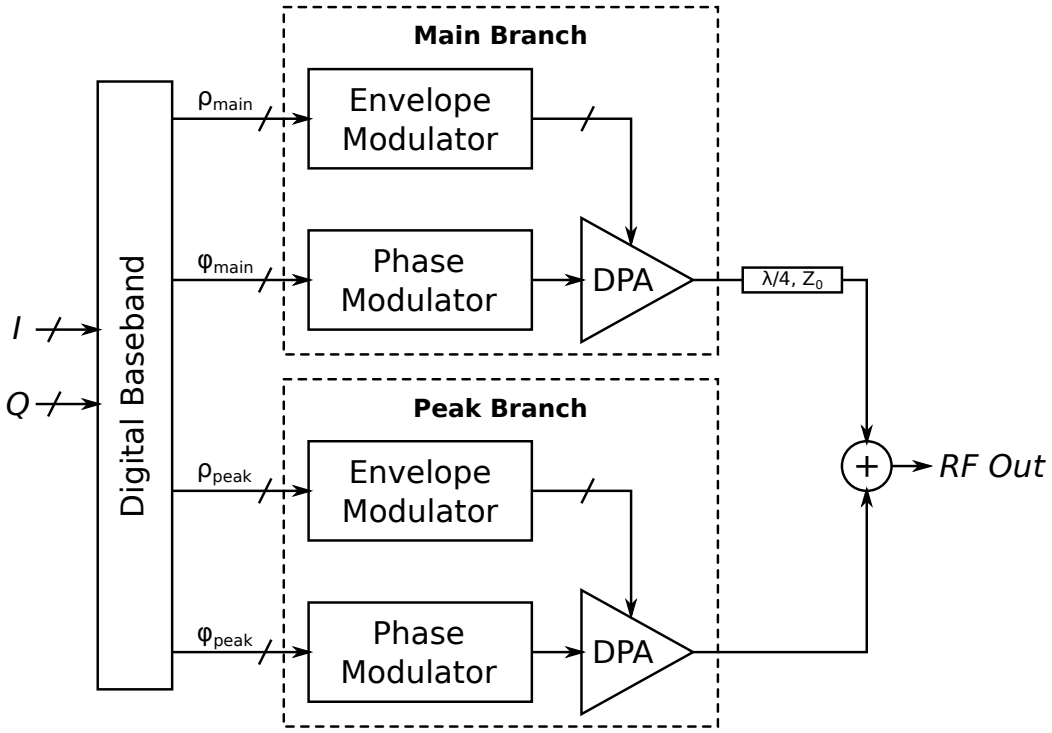

Figure 1.5: The modified Doherty architecture based on digital polar power amplifiers

the popular Doherty architecture, but addresses one of the major limiting factors for achieving high efficiency, namely, the need for linear PAs. In a typical Doherty system, the employed PAs typically operate in class-B, which limit the peak efficiency to 78.5 %. By replacing these amplifiers with PAs based on the polar architecture, it is possible to reap the benefits of the Doherty architecture while having a theoretical peak efficiency of 100 %.

A strongly simplified overview of this system is shown in Figure 1.5. The main and peak branches each consist of a phase modulator and an envelope modulator driving a digital power amplifier (DPA). Although traditional polar systems rely on supply modulation to achieve envelope modulation, the DPA consists of numerous small PA cells which can be turned on or off individually. This DPA operates in load-insensitive class-E ([16]), which is able to provide high efficiency, while the digital segmentation allows for precise envelope control via a digital codeword. In this way, the difficulties with designing a wideband high-power supply modulator are avoided. Furthermore, by relying heavily on digital techniques, the performance of this concept are expected to improve as finer technology nodes are introduced.

The digital baseband signal processing present in this system accepts I/Q data as input and computes the required envelope (ρ) and phase (ϕ) signals for the two branches. Note the absence of the $\lambda/4$ transmission line at the input of the peak branch. Its functionality is also incorporated as part of the digital baseband, which is accomplished simply with a 90° phase shift in ϕ_{peak} and ρ_{peak}.

One of the major design challenges of this system is the phase modulator. This circuit block takes a digital codeword representing the desired phase as input and outputs a phase-modulated RF signal. This output signal is used to drive the input of the DPA. In this case, since the

DPA relies on switching, rather then analog amplification, the phase-modulated RF signal has to be a square wave with rail-to-rail swing. The phase modulator is a critical part of any polar system (and outphasing system) and due to the bandwidth expansion experienced by the phase signal, designing a phase modulator which maintains good performance for wideband signals is a challenging task. As such, the design and implementation of a new phase modulator concept which is compatible with the goals of the SEEDCOM project is the topic of this thesis work.

1.4 Phase Modulator Design Challenges

The design of a phase modulator for a wideband multi-band/multi-standard polar system brings with it several challenges. First and foremost, the aforementioned bandwidth expansion when converting from I/Q to polar coordinates requires the phase modulator to be able to handle a modulation bandwidth of several times the bandwidth of the signal being transmitted. This is the fundamental reason why Cartesian systems are typically preferred for wideband applications. Nevertheless, with the continuous technology scaling of nano-scale CMOS processes, the speed of MOS devices has improved dramatically. This facilitates developing phase modulators which support bandwidths of several hundreds of MHz.

One of the goals of the SEEDCOM project is to develop a system which can support multi-band/multi-mode communication in the low GHz range. This requires the carrier frequency to be configurable over a large range of frequencies. Although eventually it is desirable to cover the entire frequency range up to approximately 6 GHz, for the purposes of this iteration of the project it was chosen to limit the maximum operating frequency to 2.4 GHz for easier implementation. For the lower frequency limit, 900 MHz was selected in order to support, for instance, the GSM-900 band.

For any RF transmitter, it is critical to maintain a sufficient in-band signal-to-noise ratio and a high degree of out-of-band spectral purity. The former is typically expressed in terms of error vector magnitude (EVM), the latter is quantified as adjacent channel power ratio (ACPR)[1]for the close-in spectrum and the level of the noise floor for the far-out spectrum. Any non-idealities in the system, including those in the phase modulator, will contribute to errors in the system and will adversely effect the performance in terms of these metrics. Unfortunately, due to the various nonlinear processes involved in polar transmission systems, it is difficult to quantify what effect phase modulator non-idealities will have on the performance of the overall system. As a consequence, for the majority of this thesis work the phase modulator performance is evaluated as part of a complete polar system, including the digital baseband signal processing and an envelope path which is assumed to be ideal.

As stated above, this TX system, and thus also the phase modulator, is meant to support multiple different standards, which can use a number of different modulation types. Evaluating the phase modulator performance for all these standards is a time consuming process, in particular in the later stages of the design, when post-layout simulations can take several days. Therefor, it was decided to focus on only a single type of modulation, namely quadrature amplitude modulation (QAM). It can be argued that evaluating the performance for orthogonal frequency-division multiplexing (OFDM) would be more relevant, as this type of modulation is used in many modern communication standards such as 802.11 (WiFi) and LTE. However, not only is a signal utilizing OFDM more complex to generate, but its symbol time is also many times longer than that of a QAM signal, which would require much longer simulation times. As

[1]Sometimes this performance measure is instead called adjacent channel leakage ratio (ACLR)

such, OFDM modulation was not considered for this design.

It should be noted that the various standards define different ways to quantify both EVM and ACPR. In the case of EVM, the definition used throughout this thesis work relates the root mean square (RMS) value of the measured error vector to the RMS value of the ideal input vector as follows:

$$\mathrm{EVM}_{RMS} = \frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(S_{meas}(i) - S_{ideal}(i)\right)^2}}{\sqrt{\frac{1}{N}\sum_{i=1}^{N} S_{ideal}(i)^2}} \tag{1.3}$$

where $S_{ideal}(i)$ and $S_{meas}(i)$ are the ideal (input) and measured (output) signal values, respectively, at the i^{th} constellation point and N is the total number of points being evaluated. ACPR measures the ratio between the signal power in the main transmit band and the signal power in the adjacent channels. It can be defined as follows:

$$\mathrm{ACPR}_L = \frac{P_{adj,L}}{P_{main}} \tag{1.4}$$

$$\mathrm{ACPR}_H = \frac{P_{adj,H}}{P_{main}} \tag{1.5}$$

where P_{main} is the total signal power in the main transmit channel, whereas $P_{adj,L}$ and $P_{adj,H}$ are the signal powers in the lower and higher adjacent channels, respectively. However, the spacing between these channels varies from standard to standard. In this thesis work, whenever ACPR evaluated, the channel spacing is assumed to be equal to half of the channel bandwidth.

1.5 Thesis Objectives

In light of the considerations described above, the following objectives were set for this thesis work:

- Develop a novel phase modulator concept for use in polar transmitter systems
- The phase modulator must have a high enough modulation bandwidth to support QAM signals with a channel bandwidth of at least 80 MHz
- When evaluating the phase modulator performance with an ideal envelope path, the polar system should meet the following requirements:
  - EVM < -35 dB
  - ACPR < -45 dBc
  - Noise floor < -60 dBc
- The phase modulator should support operating frequencies between 900 MHz and 2.4 GHz
- The phase modulator design should be as digital as possible to take advantage of technology scaling
- The design should be cost-effective, i.e. occupy a small area and have a low power consumption

- The use of inductors in this design should be avoided, as these occupy a large area, make it more difficult to support a large range of operating frequencies, and can cause high-power signals of the DPA to inductively couple to the phase modulator degrading its performance
- This phase modulator should be implemented in TSMC's 1-poly-7-metal 40 nm CMOS technology (4X0Y1Z0R1U) together with the remainder of the system described in Section 1.3 to form the demonstrator for this project

1.6 Thesis Outline

This thesis is organized as follows:

Chapter 2 presents a study of the state-of-the-art phase modulator topologies presented in recent literature. The benefits and drawbacks of the various topologies are compared on a qualitative level. A distinction is made between closed-loop and open-loop phase modulator topologies. Whereas closed-loop topologies generally have a lower power consumption, literature shows that open-loop topologies allow for vastly superior modulation bandwidths. Of the different open-loop topologies, the constant-envelope direct upconversion architecture is selected as the most promising topology to build upon in this design.

In Chapter 3 a more detailed overview of the full TX system is presented and some system-level considerations are discussed namely: bandwidth expansion, phase resolution and the choice of interpolation filter for the baseband signal. As a result of the bandwidth expansion when converting from Cartesian to polar coordinates, it is found that the phase modulator needs to support a modulation bandwidth of at least six times the signal bandwidth to meet the target ACPR requirement. A resolution of 9 bits is selected for both the in-phase and quadrature-phase components of the input phase signal to provide sufficient margin for the noise floor specification. Finally, is shown that in a digitally-intensive system, the default zero-order hold (ZOH) response will result in spectral replicas at multiple of the sampling frequency. To suppress these replicas, the use of an interpolation filter is required to up-sample the data to as high a sample-rate as possible. For this purpose a second-order hold (SOH) filter is selected.

Chapter 4 deals with the conceptual design of the phase modulator. The constant-envelope direct upconversion architecture is discussed in more detail, and the importance of filtering the output of the DACs in this architecture is illustrated. It is shown that without the use of inductors, it is difficult to sufficiently suppress higher harmonics present in the signal, which leads to significant phase error at the output of the phase modulator. To combat this issue, the harmonic rejection RF-DAC (HR-RF-DAC) is introduced.

Chapter 5 details the most relevant design decisions that were made to arrive at the final implementation of the phase modulator concept. The design process is explained for the three main components of the modulator: the HR-RF-DAC, the limiting amplifier and the RC filter. These components are discussed both at the schematic level and the layout level. Finally, the top-level layout of the phase modulator is presented.

Chapter 6 presents the simulation results of the phase modulator design. Most of the simulations are performed at the post-layout level. The performance of the phase modulator is evaluated in several ways. First, the static phase error versus input phase is evaluated, both at the post-layout level and in a schematic-level Monte-Carlo simulation. Then, the step response is analyzed in order to estimate the modulation bandwidth. Finally, the performance of the phase modulator as part of a polar system with an ideal envelope path is evaluated for a QAM input signal. This is done at a carrier frequency of 2.4 GHz with 80 MHz channel bandwidth for

both ZOH and SOH interpolation. Furthermore, the results of a simulation at 900 MHz with a proportionally scaled channel bandwidth (30 MHz) are also shown to demonstrate the phase modulator's frequency-agile operation.

Finally, Chapter 7 concludes this thesis and provides a discussion of the results and the overall design process. Some directions are suggested for future work.

Chapter 2

Phase Modulator Concepts

To choose a suitable concept for the phase modulator design, it is important to consider the state-of-the-art designs presented in literature. This chapter provides an overview of the different phase modulator concepts and compares them on a qualitative level. At the end, a concept is chosen as a base to build upon in this thesis work.

Perhaps the easiest way to achieve phase modulation is to apply the modulation data directly to the input port of a voltage-controlled oscillator (VCO). Although this approach allows for wideband modulation, it suffers from numerous problems such as frequency drift, VCO transfer function non-linearity, high close-in phase noise and large variability due to process, voltage and temperature (PVT) variations. All these issues render this phase modulation approach unusable for practical communication systems. Over the years, several phase modulator concepts have emerged to overcome these problems. These concepts can be broadly classified into two categories: closed-loop and open-loop phase modulators [17].

2.1 Closed-Loop Phase Modulators

Closed-loop phase modulators are based around a phase-locked loop (PLL), which is typically already present in RF systems to synthesize the LO waveform. A block diagram of a conventional PLL is shown in Figure 2.1. The operation principle of a PLL involves placing a VCO in a feedback configuration in order to lock the phase of the VCO output voltage to that of a reference clock. To achieve phase lock, first a phase-frequency detector (PFD) measures the phase (and frequency) error of the VCO output with respect to the reference. This error is fed to a charge pump (CP) which adjusts the input voltage of the VCO. This voltage is filtered by a low-pass filter (LPF), which serves to stabilize the loop, as well as filter out disturbances in the loop caused by, for instance, thermal noise.

The system, as described so far, outputs a waveform at the same frequency as the reference clock. In order to achieve higher frequencies, a frequency divider can be placed in the feedback path. Since communication standards define several frequency channels, the division ratio needs to be adjustable. As such, a multi-modulus divider (MMD) is required. Such a divider, however, can only divide the frequency by an integer factor. To allow for more fine-grained frequency control, the modulus select of the MMD can be driven by a sigma-delta modulator, such that the average division ratio over time tends toward the desired fractional value.

The PLL has several sources of error apart from thermal noise. These include feed-through of the reference signal, as well as fractional frequency spurs due to the frequency division. Fur-

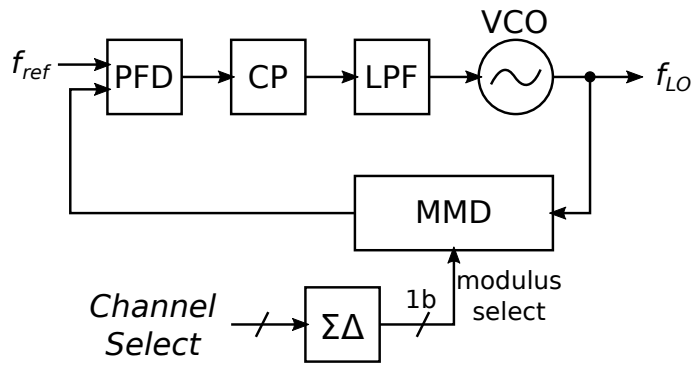

Figure 2.1: A conventional phase-locked loop without modulation

thermore, the sigma-delta modulator introduces quantization noise, which is shaped to have low noise power at low frequencies and high noise power at high frequencies. If these errors were to appear in the output spectrum of the VCO, it would be detrimental to the performance of the RF system. For this reason, the loop filter needs to have as low a cut-off frequency as possible, so that it has good suppression at the frequencies at which these errors appear [18].

A further source of error that presents itself when a PLL is integrated on the same chip as a high-power amplifier is VCO pulling. This occurs when the modulated RF signal couples to the VCO causing the oscillation frequency to shift. VCO pulling can be counteracted by ensuring good isolation between the VCO and the PA and by operating the PLL at a higher frequency than the carrier frequency. Ideally, these two frequencies should not be harmonically related. A fractional frequency divider can then be used at the PLL output to produce the desired carrier signal [19].

2.1.1 Basic Phase Modulation

In order to achieve phase modulation, the phase data can be converted to frequency data and superimposed on the channel select signal [20], [21]. This technique is illustrated in Figure 2.2. As can be seen, only a very minor addition of circuitry is needed for this technique. This make this type of phase modulator very cheap in terms of power consumption. However, this technique is only suitable for very narrowband modulation, since the loop filter removes any high-frequency content.

Various techniques have been devised to increase the available modulation bandwidth. One approach is based on the idea that if the errors in the PLL can be suppressed by other means than low-pass filtering, then the bandwidth of the loop filter can be increased while maintaining the required spectral purity. One example of this approach is phase quantization noise (PQN) cancellation [22]. This technique is quite effective because PQN is normally the dominant source of error for larger loop bandwidths. Another example is the multiphase fractional-N PLL [23], where one of several phases can be used in the feedback path to further reduce quantization noise and suppress fractional spurs.

An alternative approach to increase the modulation bandwidth is digital pre-emphasis [24], in which the PM data is amplified in the digital domain for higher frequencies such that the low-pass response of the loop filter is canceled out. However, for very wideband modulation an extremely high dynamic range is required to accurately represent the signal, both in the digital

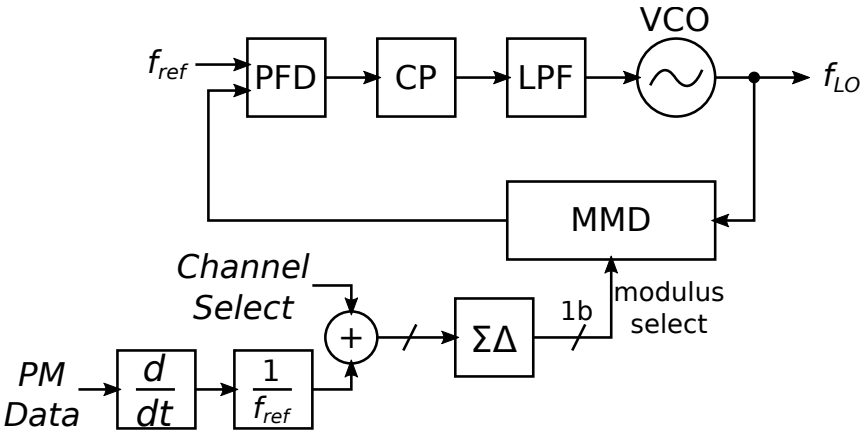

Figure 2.2: A basic closed-loop phase modulator

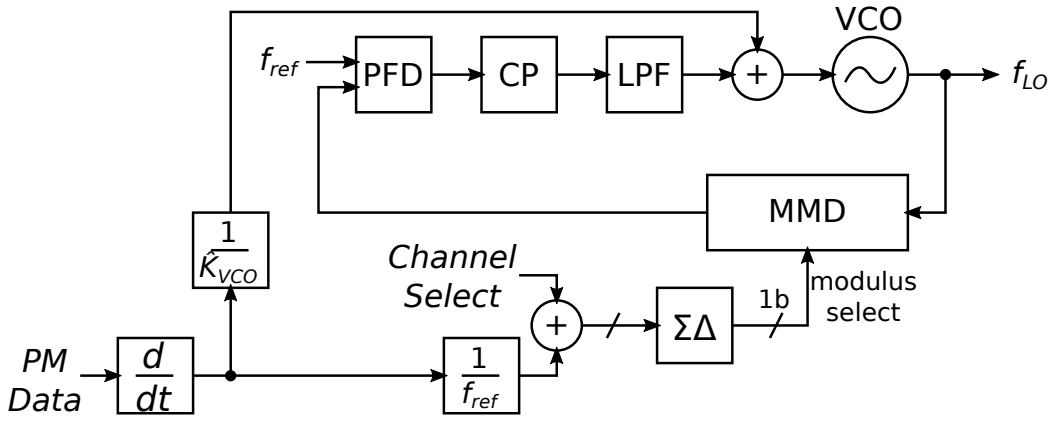

Figure 2.3: A two-point closed-loop phase modulator

and analog domains.

Despite the various bandwidth-extension techniques discussed above, even the most advanced designs using the basic phase modulation method have not exceeded 3 MHz modulation bandwidth [17].

2.1.2 Two-Point Phase Modulation

To overcome this apparent bandwidth limitation, the two-point modulation method is commonly employed as was first proposed in [25] and implemented in [26]. A block diagram of a typical two-point phase modulator is shown in Figure 2.3. This architecture relies on the fact that any node in the PLL has either a high-pass or a low-pass response. Applying a modulation signal to two nodes simultaneously results in a wideband frequency response, assuming proper signal conditioning. The most common nodes used for this method are the MMD and VCO input ports for the low-pass and high-pass paths, respectively.

One of the major difficulties in with this architecture is the gain and phase mismatch between the two paths, which can result in performance degradation. Phase mismatch stems from the

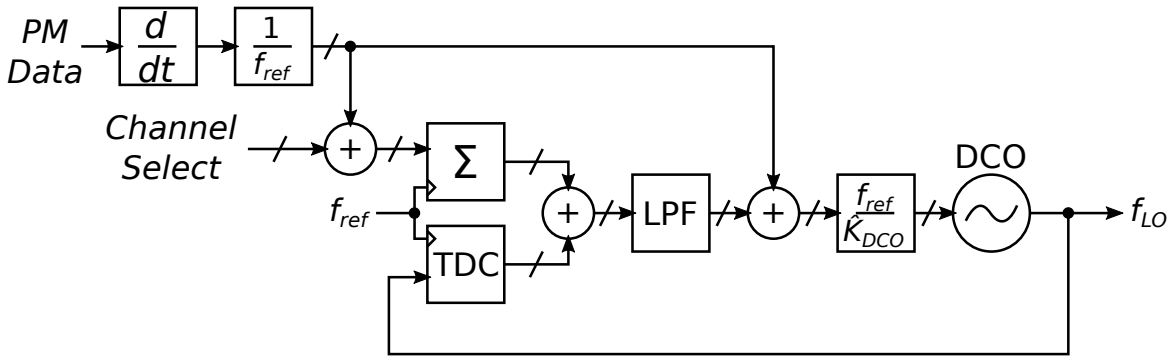

Figure 2.4: An all-digital PLL with two-point modulation

heterogeneous circuitry used for the two paths. Whereas the the modulation of the MMD happens in the digital domain, the VCO requires an analog signal to be applied at its input port. However, it has been shown in [27] that for bandwidths of a few MHz phase mismatch can be suppressed to acceptable levels using digital techniques, assuming a sufficiently high sample-rate. Gain mismatch, on the other hand poses a more serious problem. Although the gain of the low-pass path can be very well controlled, the high-pass path requires accurate estimation of the VCO gain K_{VCO}, which is both highly sensitive to PVT variation and nonlinear. One solution to this problem was proposed in [28], where a VCO calibration scheme was proposed using a 5-bit switched capacitor bank. This resulted in coarse digital control, so that the VCO could be designed for a small K_{VCO}, reducing the sensitivity to PVT. This method, however, does not improve the linearity. An alternative solution, proposed in [29], is to enclose the VCO in a feedback loop with a frequency-to-voltage converter, in order to linearize the VCO transfer function. Though this method has shown to be effective for EDGE and WCDMA applications, the feedback loop around the VCO inevitably limits the achievable bandwidth.

2.1.3 All-Digital PLL

With the move towards deep-submicron CMOS technology, the digital performance of CMOS devices improved dramatically, at the cost of sacrificing analog performance. This prompted the development of the all-digital phase-locked loop (ADPLL), which replaced the analog components with their digital counterparts [30], [31]. A block diagram of an ADPLL with the two-point modulation technique discussed above is shown in Figure 2.4. The VCO of the traditional PLL is replaced with a digitally-controlled oscillator (DCO) and the operation of the PFD is accomplished by a time-to-digital converter (TDC). The frequency control word is integrated in an accumulator and converted to phase such that it can be readily compared to the TDC output. Since all of the signals in the ADPLL are in the digital domain, the behavior of each node in the loop is much more predictable, and digitally-intensive calibration techniques can more easily be applied.

Similar to the traditional PLL, the modulation bandwidth of the ADPLL is limited by the gain and phase mismatch of the signal paths used in two-point modulation. In [32] and [33], two background gain estimation algorithms were proposed which worked by continuously measuring the instantaneous phase error and adjusting the estimated gain factor accordingly. Additionally, [33] proposed a method for measuring the phase error between the two paths, and predistorting

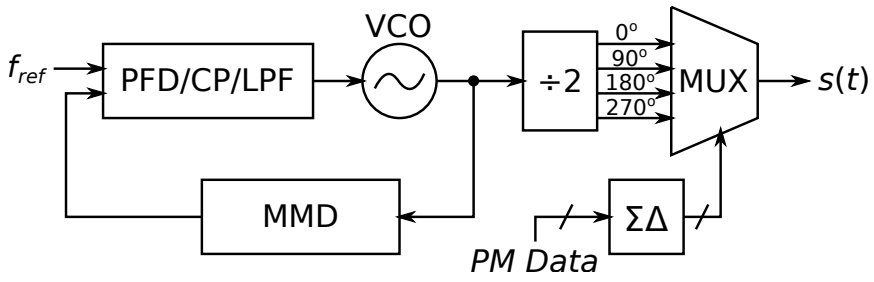

Figure 2.5: An open-loop phase modulator with ΣΔ-modulated phase multiplexer

one of them in the digital domain in order to correct for it.

It should be noted that the techniques discussed in the previous sections that were used to improve the performance of analog PLLs are equally applicable to the ADPLL. For example, [34] proposed a PQN cancellation method in the digital domain, allowing the ADPLL loop bandwidth to be extended without a significant degradation in phase noise performance.

2.2 Open-Loop Phase Modulators

Despite the various performance enhancing techniques for PLLs that have been developed in recent years, to the best of the author's knowledge, no closed loop phase modulators have been reported that can achieve high enough modulation bandwidths for very wideband applications such as WLAN, LTE and WiMAX. To make wideband modulation possible, other concepts have been proposed which use the PLL to generate one or several static phases of the LO which are then modulated outside the loop by separate circuitry. These concepts are collectively known as open-loop phase modulators [17].

2.2.1 ΣΔ-Modulated Phase Multiplexer

One technique for modulating the phase of the carrier outside the loop is to pass the output of a PLL through a frequency divider to generate four quadrature phases and switch between these phases using a multiplexer (mux) to generate the phase-modulated output. With this approach, the phase modulator is not limited by the small bandwidth of the PLL, nor does it suffer from the gain and phase mismatch that plagues the two-point modulation scheme. Instead, the modulation bandwidth is only limited by the speed of the mux, which can be very fast in deep sub-micron CMOS technologies. The obvious drawback of this approach is the high PQN, since only one of four discrete phases can be selected. To alleviate this issue, the multiplexer can be controlled by a ΣΔ modulator, which has the effect of pushing the quantization noise to higher frequencies. The resulting system architecture, proposed in [35], is shown in Figure 2.5.

Unfortunately,to achieve the desired spectral purity, the output signal needs to be band-pass filtered, in order to eliminate the high out-of-band PQN caused by ΣΔ modulation. While for a PLL, the loop filter at baseband was sufficient, for this open-loop architecture a separate filter at the carrier frequency is required. As a direct consequence, the design in [35] achieved only approximately 8 MHz modulation bandwidth.

Another drawback of this design is that the PLL has to synthesize double the desired carrier frequency. Apart from making it more challenging to design, this also increases the power

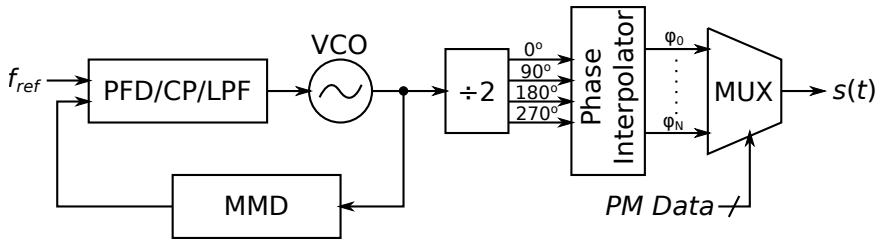

Figure 2.6: An open-loop phase modulator based on a multiplexed phase interpolator

consumption of the modulator. However, alternative methods can be used to generate quadrature phases. These methods include using a quadrature VCO [36] or an inverter-based ring oscillator [37]. In fact, the design in [37] generated 31 different phases rather than 4, which significantly reduces the quantization noise. Unfortunately, as ring oscillators suffer from poor phase noise performance, they are rarely used for modern communication systems [38].

2.2.2 Multiplexed Phase Interpolator

One way to suppress the high PQN of the technique described above is to increase the number of phases that can be selected by the multiplexer. This can be achieved by using a phase interpolator. A phase interpolator is a circuit block which can generate one or more output phases by performing a weighted summation of the two closest input phases. An example of phase-iterpolation-based architecture is shown in Figure 2.6.

Although the phase interpolator can be designed to have any number of output phases and switching between them can be very fast, practical designs typically limit themselves to only a few outputs. One reason for this is the high power consumption that results from having to drive so many signals at the LO frequency. Another reason is the added layout complexity of such an approach. As a result, this architecture is rarely used on its own. For instance, the design in [39] used a phase interpolator to generate only 8 phases, and the modulator presented in [40] generated 64. This translates to a phase resolution of only 3 bits and 6 bits, respectively. Since in both cases a higher resolution was required, the multiplexed phase interpolator was used only for the most significant bits (MSB), while the least significant bits (LSB) of the modulation were accomplished using other techniques.

2.2.3 Delay-Line-Based Phase Modulator

An alternative method for phase modulation is to use inverter-based delay elements do delay the LO signal by the desired phase angle. This can be done either statically, creating a tapped delay line and selecting one of the phases with a mux, or dynamically, by controlling the delay of a single delay element with the modulation data. For example, the tapped delay line was used in [41] to create a low resolution, high speed phase modulator. Much like the multiplexed phase interpolator discussed above, high resolution designs based on this approach are impractical. In this case, it is due to the limited switching speed of the inverters. The digitally controlled delay, on the other hand, is limited in range but can provide very fine control, and is one of the techniques used to increase the resolution of the design in [40] beyond what the phase interpolator would allow. These two approaches can be combined such that the former is used for the coarse

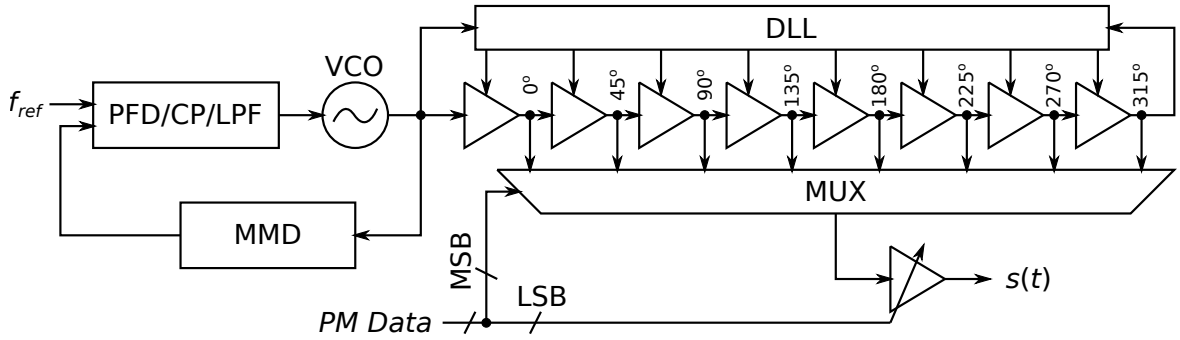

Figure 2.7: An open-loop phase modulator with a coarse tapped delay line and a fine digitally-controlled delay

MSB selection while the latter accomplishes the fine LSB modulation. This design was proposed in [42] and is illustrated in Figure 2.7.

In both cases, the delay needs to be accurately controlled to achieve accurate phase modulation. Since inverter delay is heavily dependent on PVT variations, some form of calibration is required. For a tapped delay line, a typical approach is to enclose the line in a delay-locked loop (DLL), as shown in the figure. A DLL controls the delay of each element in order to align the edges of the VCO output and the output of the last delay element. The drawback of this design is that although the DLL regulates the delay of the whole line, the delay of the individual elements is controlled with the same signal. This means that any mismatch between the elements will still result in phase errors.

For the digitally controlled delay, the input signal can be predistorted to account for the PVT variations. For this to be effective, the delay element needs to be designed for a slightly larger range and resolution than in the ideal case. It should be noted that predistortion of the input signal can also be used to compensate for the for the mismatch between the elements of the tapped delay line discussed above.

Despite the issues associated with mismatch, it has been shown that with proper calibration this architecture shows great promise for wideband applications. The design in [42] used two such phase modulators to drive two class-D PAs in outphasing configuration and was able to transmit WLAN signals with both 20 MHz and 40 MHz channel bandwidths.

2.2.4 Constant-Envelope Direct Upconversion

Another open-loop phase modulator concept is based on the traditional direct upconversion (Cartesian) transmitter designs. While such transmitters are inferior to polar topologies in terms of power efficiency, they are known to be very wideband. If phase modulation data is converted to a constant-envelope I/Q signal by means of a phase-to-IQ converter, the Cartesian architecture can serve as a phase modulator for a polar PA. If a switching PA is used, additionally a limiting amplifier is required at the output to convert the phase-modulated signal into a square wave. This is illustrated in Figure 2.8. In this way, a highly efficient non-linear polar PA can be employed while still exploiting the high bandwidth of the direct upconversion structure. Using this architecture, the design proposed in [43] was able to meet the requirements for 20 MHz WLAN signals.

Compared to other open-loop architectures, this technique in general has a higher power

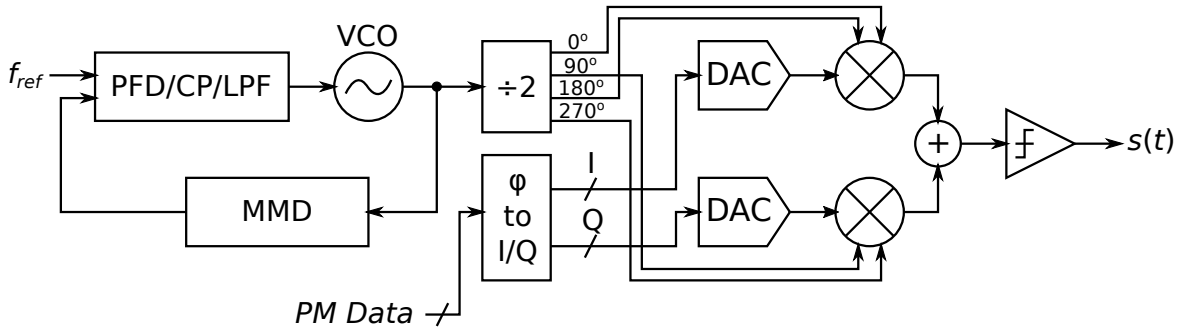

Figure 2.8: An open-loop phase modulator based on constant-envelope direct upconversion

consumption, due to the relatively power-hungry DACs and the additional digital circuitry associated with the phase-to-I/Q conversion. Nevertheless, low-power implementations have been reported in finer technology nodes. For instance, the phase modulator described in [44], implemented in 65 nm CMOS, consumed only 9 mW of power despite having a high enough bandwidth to support 20 MHz WLAN signals. Such a low power consumption is mainly due to a simplification of the phase-to-I/Q converter. Typically, such a phase-to-I/Q converter needs to run at a very high clock frequency in order to push spectral replicas (resulting from the zero-order hold (ZOH) operation of the DACs) outside of the band of interest. Since the relationship between phase and I/Q is non-linear, complex digital circuitry is required, which drives up the power consumption. In this design, phase-to-I/Q converter ran at only 200 MS/s. Its output was then linearly interpolated up to 1 GS/s to push away the spectral replicas. A relatively simple digital circuit was then used to normalize the I/Q data such that $I + Q$ remained constant. Because the signal was no longer constant-envelope, the limiting amplifier at the output exhibited AM-PM distortion. Nevertheless, this effect was small enough that it could be compensated for by predistorting the input phase signal.

More recently, it has been demonstrated that this architecture can be used for even more wideband signals. The design presented in [45], implemented in 45 nm CMOS silicon-on-insulator technology, used two such phase modulators in outphasing configuration. It managed to maintain good performance for 256-QAM OFDM signals with bandwidths up to 133 MHz in the 10 GHz band.

2.3 Conclusion

This chapter has presented several phase modulator concepts which have been developed over the course of the past two decades. Traditional concepts based on modulation of the PLL control word, have shown to be very well suited for ultra-low-power narrowband applications, since very little additional circuitry is required. Several bandwidth extension techniques have been proposed, such as PQN cancellation and digital pre-emphasis, in order to increase the available modulation bandwidth. Nevertheless, the inherently narrowband nature of the PLL has proven difficult to overcome. Two-point modulation has been proposed to increase the modulation bandwidth beyond the loop bandwidth of the PLL, but the effectiveness of this method is limited due to the gain and phase mismatch between the two paths. Despite the numerous efforts to correct for this mismatch, neither analog nor digital PLL-based phase modulator implementations

have been reported which support modulation bandwidths beyond a few MHz.

To allow for more wideband modulation, open-loop phase modulator concepts have been developed. These concepts typically rely on a PLL to provide quadrature LO phases and then modulate these outside of the PLL loop. A simple way to achieve this is through the use of a multiplexer controlled by a $\Sigma\Delta$ modulator. While this allows for a higher bandwidth than what is achievable with a closed-loop phase modulator, the bandwidth is still limited by the high quantization noise at higher frequencies that results from $\Sigma\Delta$ modulation. Furthermore, $\Sigma\Delta$ modulators need a high oversampling ratio to be effective, limiting the effective baseband sample rate. Phase interpolation can be used to achieve finer phase quantization, but the resolution of the phase interpolator cannot be pushed indefinitely, since the power consumption and layout complexity associated with driving so manly signals at the LO frequecy would make the design impractical.

Both delay-line-based phase modulators and modulators based on constant-envelope direct upconversion have shown to be promising candidates for wideband phase modulation. In both cases, it has been shown that these architectures can meet the requirements of modern wideband communication standards. In the case of the delay-line-based modulators, however, strong dependence on PVT variations make on-chip calibration algorithms a requirement. Furthermore, this architecture is not amenable to frequency-agile operation. This is due to the fact for high frequencies a fine resolution is needed, whereas lower frequencies need a large delay range. Combining these two requirements would result in a very bulky and power-hungry design.

Phase modulators based on constant-envelope direct upconversion do not suffer from this issue, since in theory, all of the circuitry after the PLL is completely frequency-independent. The main drawback of this architecture is the relatively high power consumption. However, when used with a high-power amplifier, as is the case in base-stations, this can have very little impact on system efficiency. It is for this reason, that the constant-envelope direct upconversion architecture has been chosen as a base to build upon in this thesis work.

Chapter 3

System Level Considerations

In order to better understand what is required of the phase modulator, it is important to consider its performance in the context of the whole TX system. This chapter provides an overview of the most important system level considerations for the phase modulator design. First, an overview of the TX system is presented in Section 3.1. The remainder of the chapter deals with the three main sources of error at the system level, namely, finite modulation bandwidth, quantization noise and errors associated with input signal interpolation. Section 3.2 discusses the issue of bandwidth expansion due to the conversion from Cartesian to polar coordinates and its effect on performance. The effects of finite phase resolution are evaluated in Section 3.3. Section 3.4 shows the advantages and disadvantages of interpolating the input signal to a higher sample rate using several methods. Finally, some conclusions on the system level are drawn in Section 3.5.

3.1 TX System Overview

While a simplified overview of the digital Doherty TX system developed during this project has already been discussed in brief in Section 1.3, it is illustrated in much more detail in Figure 3.1. Included in the figure is a detailed overview of the digital signal processing (DSP) that precedes the envelope and phase modulators. The input to the system is a 600 MS/s baseband I/Q signal as well as an upconverting input RF clock at four times the desired carrier frequency. The RF input is fed to a divide-by-four circuit to produce eight different phases at the carrier frequency which are all separated by 45°. The main reason for using eight phases will be discussed in Section 4.3. The I/Q immediately gets converted to envelope (ρ) and phase (ϕ) signals using a COordinate Rotation DIgital Computer (CORDIC) [46].

As with any Doherty system, the input signal needs to be split into two different paths to feed the main and peak amplifier branches. In a convensional Doherty transmitter the two branches are driven by identical input signals apart from a 90° phase shift. The amplifiers in the two branches are biased differently such that the main branch is always operational, whereas the peak branch only turns on when a high power level is required. In this digital Doherty system, on the other hand, the main and peak branches are identical. The difference between the two lies in the input signals that are supplied to them. In the case of the ρ signal, it is split into ρ_{main} and ρ_{peak} using a look-up table (LUT). The ϕ signal, on the other hand, is fed to both branches unaltered. In order to get the desired behavior, a 90° phase shift between the two branches is still required, both in the envelope and in the phase paths. For the phase, this shift is achieved by shifting the phase of the LO rather than that of the input signal. This is easily accomplished

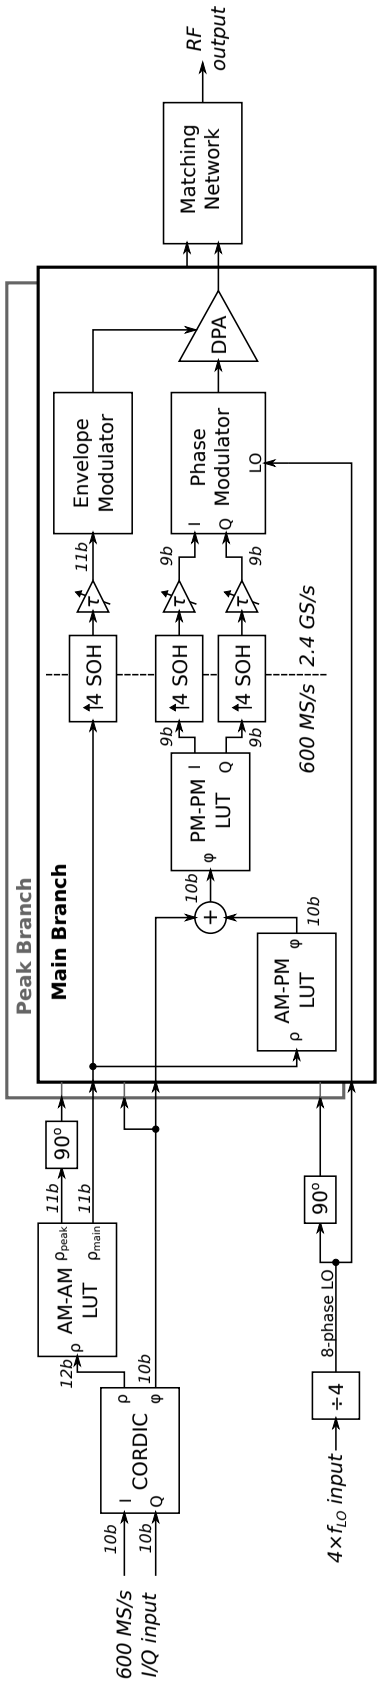

Figure 3.1: An overview of the digital Doherty TX system including digital processing of the input signals

by changing the order of the different LO phases which are supplied to the peak branch. For the envelope path, the 90° shift is accomplished using a delay element.

The LUT mentioned above serves an additional purpose apart from separating the envelope signal. It can also be used for static AM-to-AM digital pre-distortion (DPD). Within each branch, another LUT handles the AM-to-PM DPD by adjusting the phase signal based on the appropriate envelope input.

As explained in Section 2.3, the phase modulator developed in this thesis work is based on the constant-envelope direct upconversion architecture. As such, its input phase signal needs to be converted to a constant-envelope I/Q signal. This too is done using a LUT, which additionally serves the purpose of PM-to-PM DPD. Because this system employs a polar PA and the phase modulator outputs a square wave, which is by nature constant-envelope, this system has no PM-to-AM distortion. Therefor, PM-to-AM DPD was not required.

All of the digital processing blocks described so far are clocked at 600 MS/s. As will be discussed Section 3.4, the resulting signals need to be upsampled and filtered to suppress unwanted frequency spurs in the output spectrum. In this system, both the envelope and the phase (represented in the I/Q domain) are upsampled by a factor of four to the carrier frequency and filtered using a second-order hold (SOH) interpolation filter.

These upsampled envelope and phase signals then get fed to their respective modulators. The envelope modulator primarily consists of a thermometer encoder to drive the unary-weighted cells of the DPA. The phase modulator takes the various LO phases and modulates them by the input phase signal. The output of the phase modulator is a square wave at the LO frequency with a phase that is equal to the input signal. The outputs of the envelope and phase modulators get combined at the DPA which constructs the RF signal of either the main or peak branch. Finally, the outputs of the two branches get combined using a matching network to generate the RF output. The matching network also provides a 90° phase shift at the output of the main branch, which is needed for Doherty operation.

One of the major difficulties of the polar topologies is the alignment of the envelope and phase paths. In a traditional system, where the envelope modulation is achieved by modulating the supply voltage of the PA, this alignment is made difficult due to the very different time constants associated with the envelope and phase paths. In this system, however, since the envelope modulation is accomplished in the digital domain, the alignment of envelope and phase is made easier. Nevertheless, some mismatch in propagation delay is still expected. To compensate for this, programmable delay elements are inserted after the interpolation filters, such that the envelope and phase paths can be aligned as part of the calibration procedure.

As explained above, the main and peak branches are functionally identical. The difference between the two lies only in the supplied input signals and how the outputs are handled by the matching network. From the perspective of the phase modulator, each branch can be seen as a separate TX system, which significantly simplifies the analysis and evaluation of the system performance. For this reason, throughout the rest of this thesis work only a single branch is considered and the Doherty operation is disregarded for simplicity.

3.2 Bandwidth Expansion

One of the major issues associated with polar transmission is the bandwidth expansion due to the conversion from Cartesian to polar coordinates. The relationship between I/Q (Cartesian)

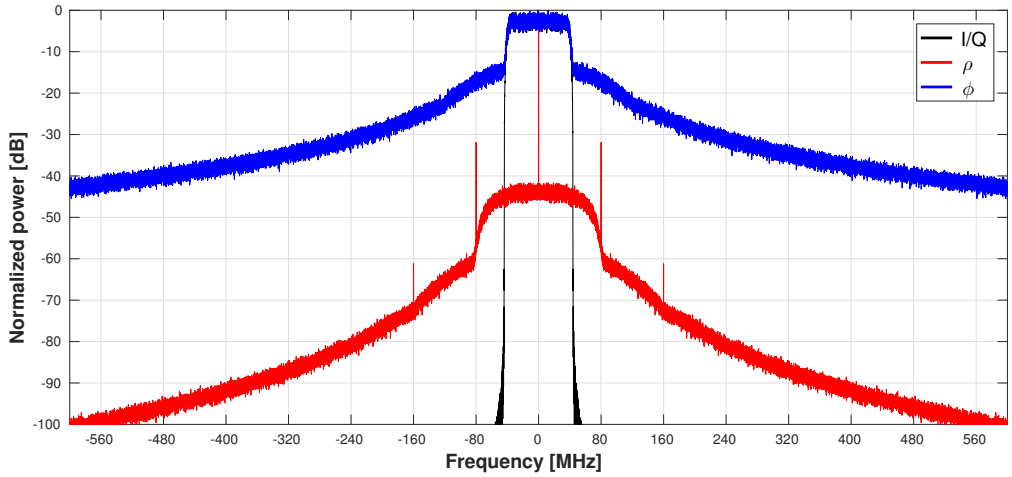

Figure 3.2: Power spectra of an 80 MHz 256-QAM signal and its envelope and phase components

and ρ/ϕ (polar) representations of the signal are given by

$$\rho = \sqrt{I^2 + Q^2} \tag{3.1}$$

$$\phi = \arctan\left(\frac{Q}{I}\right) \tag{3.2}$$

As can be seen from these equations, this relationship is highly nonlinear, which is the root cause of the bandwidth expansion.

To demonstrate this, an 80 MHz 256-QAM signal was generated and separated into its envelope and phase components. The resulting signal spectra[1] can be seen in Figure 3.2. As the figure shows, although the I/Q signal is contained within an 80 MHz bandwidth, the same cannot be said for its polar components. Both the envelope and phase signals occupy a bandwidth which tends towards infinity. This clearly shows that the envelope and phase modulators need to support a much larger bandwidth than would be required in a Cartesian transmitter in order to faithfully reconstruct the input signal. To avoid confusion, a distinction is made in this thesis work between signal bandwidth and modulation bandwidth. Signal bandwidth refers to the bandwidth of the I/Q signal which is supplied as input to the TX system. Modulation bandwidth, on the other hand, refers to the bandwidth that the envelope and phase modulators need to support.

To determine what modulation bandwidth is required of the phase modulator, the phase signal was filtered with a set of second-order low-pass Butterworth filters with varying cutoff frequencies. The resulting phase spectra can be seen in Figure 3.3. Note that the filter bandwidths reported in the figure are the double sideband (DSB) bandwidths. Since this thesis work deals only with the design of the phase modulator, any non-idealities in the envelope path are not considered.

The filtered phase signals were combined with the ideal envelope signal to produce the RF output. The resulting power spectra and constellation diagrams can be seen in Figures 3.4a and 3.4b, respectively. Note that is that the limited bandwidth of the phase path has only a

[1]The power spectrum of the phase signal is computed using $\exp(j\phi)$ rather than ϕ because ϕ itself has no physical representation. Nevertheless, throughout this thesis work it is denoted as the spectrum of ϕ to keep to the convention followed in literature.

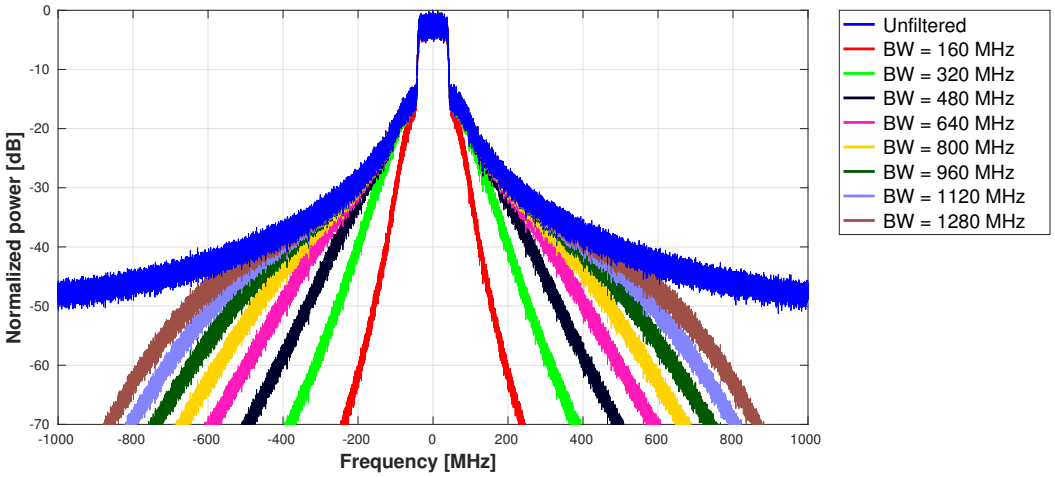

Figure 3.3: Power spectra of the phase when filtered with a set of second-order low-pass Butterworth filters with varying cutoff frequencies

minor effect on the in-band performance. Even a modulation bandwidth of only twice the signal bandwidth, can meet the EVM requirement, albeit by a very narrow margin. The effect on the out-of-band performance, however, is much more severe. As can be seen in Figure 3.4a, the RF output exhibits significant spectral regrowth due to the limited bandwidth of the phase signal. A modulation bandwidth of at least six times the signal bandwidth is required to meet the ACPR requirement. An even higher bandwidth is needed to provide enough margin so that the desired spectral purity can be maintained even when other sources of error are considered.

3.3 Phase Quantization

A critical issue to consider in any digital system is the finite resolution of the various signals. This introduces quantization noise in the system, which is often the limiting factor for the out-of-band noise floor. This section investigates the effects of finite resolution of the phase path on the system performance. Apart from the finite resolution of the I/Q input signal, there are two DSP operations in the phase path for which the desired resolution should also be determined. The first is the conversion from I/Q to ϕ (and ρ) which is performed by the CORDIC, while the second is the conversion from ϕ back to constant-envelope I and Q, which is achieved with a LUT. This is illustrated in Figure 3.5, in which the points at which re-quantization occurs are made explicit.

A system level simulation was run with an 80 MHz 256-QAM signal in which the input phase signal (as presented at the output of the CORDIC) was quantized to several different resolutions. In this simulation, the finite resolution of the I/Q input signal was not taken into account. The resulting signal spectra and constellation diagrams can be seen in Figures 3.6a and 3.6b, respectively. As shown in Figure 3.6a, each additional bit of resolution improves the out-of-band noise floor by approximately 6 dB. Additionally, for the lower resolutions, the in-band noise floor (which also determines the EVM) is significantly worse than the out-of-band noise floor. This is due to the non-linear relationship between phase and I/Q. This translates into poor EVM performance for the 4-bit and 5-bit cases, as illustrated in Figure 3.6b.

One of the performance targets for this design is an out-of-band noise floor of better than

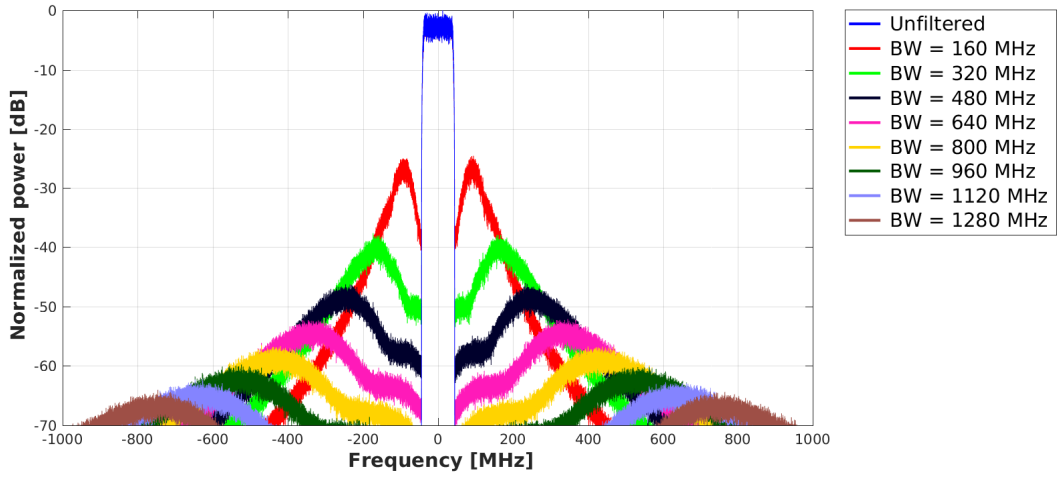

(a)

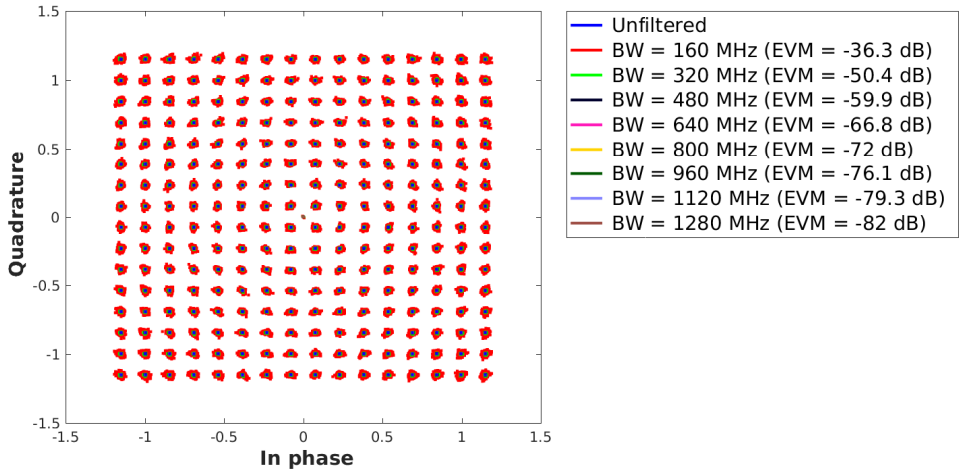

(b)

Figure 3.4: (a) Power spectra and **(b)** constellation diagrams of the RF output with band-limited phase input

−60 dBc. As can be seen from Figure 3.6a, a phase resolution of 8 bits would be sufficient to meet this target. However, this analysis did not take into account the quantization noise associated with the envelope path or any other non-idealities that could affect the noise floor. To provide a margin for error of at least 10 dB, a 10-bit resolution for the phase signal was chosen for this design.

So far, only the quantization of the phase signal itself has been considered. However, as explained in Section 3.1, the phase modulator concept that was chosen for this design requires its input to be supplied as an I/Q signal. Since this too is a digital signal, the ϕ-to-I/Q conversion introduces an additional quantization operation. In order to select a suitable resolution for the I and Q signals, the same simulation was run with a 10-bit phase input and the I and Q after conversion were quantized to several different resolutions. The resulting signal spectra and constellation diagrams can be seen in Figures 3.7a and 3.7b, respectively.

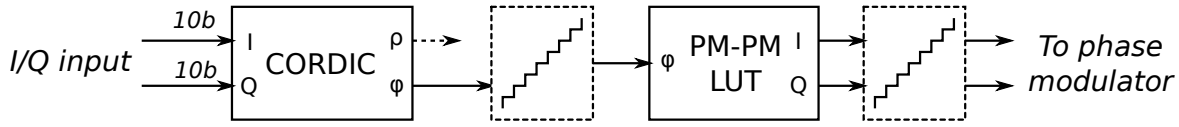

Figure 3.5: Block diagram of the DSP chain of the phase path, showing where re-quantization of the input signal occurs

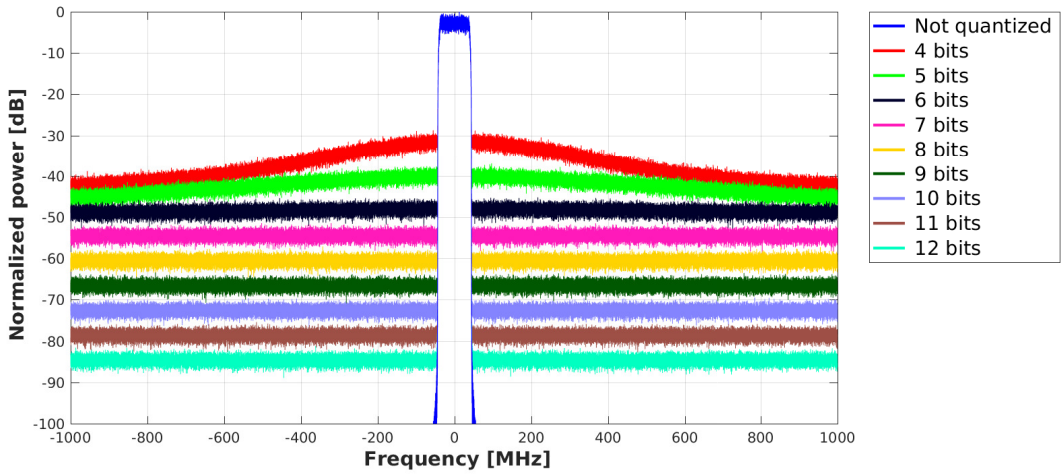

(a)

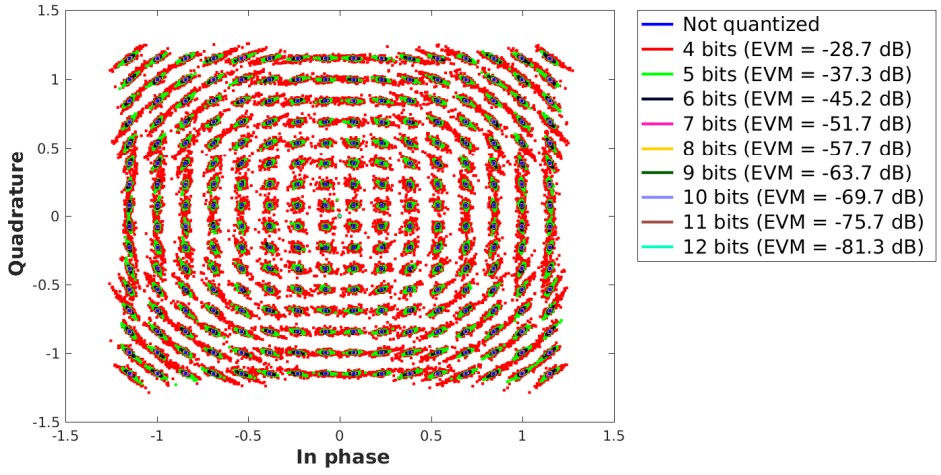

(b)

Figure 3.6: **(a)**Power spectra and **(b)** constellation diagrams of the RF output with finite phase resolution

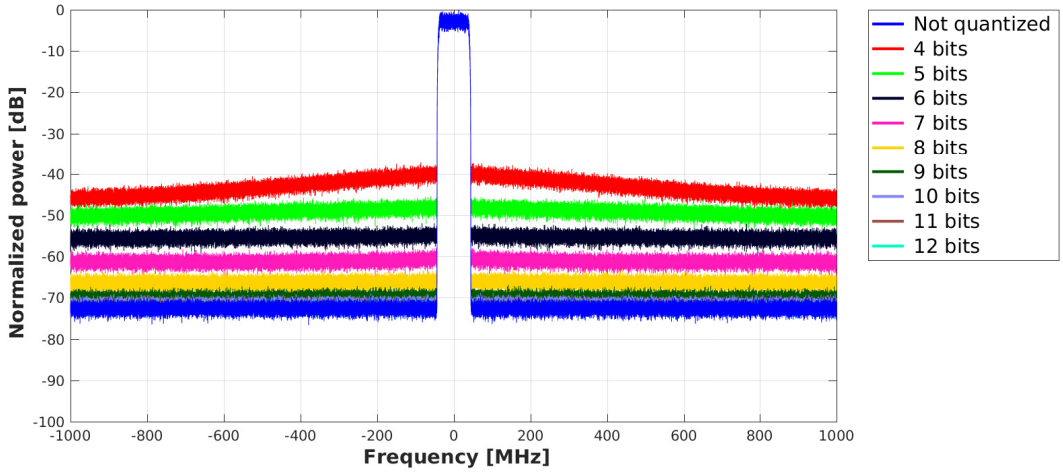

(a)

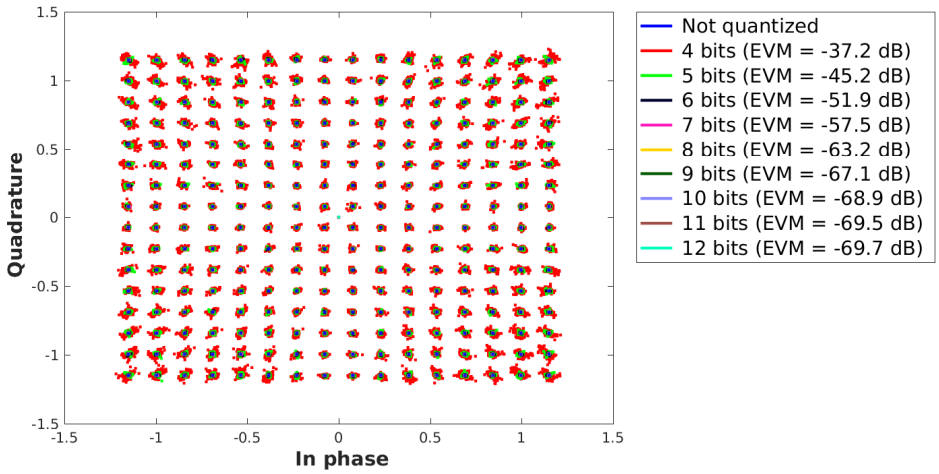

(b)

Figure 3.7: (a) Power spectra and **(b)** constellation diagrams of the RF output with 10-bit phase input and various resolutions in the I/Q domain after ϕ-to-I/Q conversion

Much like in the previous case, it is the out-of-band performance, rather than the EVM, that determines what resolution is required. As can be seen from Figure 3.7a, a resolution of at least 7 bits is required to meet the -60 dBc noise floor requirement. For higher resolutions, the spectrum quickly converges to that of the input signal. Initially, a resolution of 10 bits was selected for the I and Q signal, in order to not degrade the out-of-band performance any further. However, this is directly related to the resolution of the DACs in the phase modulator. In the later stages of the design, it was found that this decision resulted to a very large area for the phase modulator. This, in turn, gave rise to high power consumption (due to longer high speed signal lines) and a high parasitic capacitance. It was later decided to reduce the I/Q resolution to 9 bits, which largely alleviated these issues. As can be seen in Figure 3.7a, this sacrifices only about 2 dB in the noise floor when compared to the 10-bit case.

3.4 Input Signal Interpolation

One of the drawbacks of digital TX systems with respect to their analog counterparts is that the output spectrum of a digital system will exhibit spectral replicas at integer multiples of the sampling frequency. This occurs because in a digital system the signals are held at a constant value between the sampling points and the transition from one value to the next happens almost instantly. This is known as zero-order hold (ZOH) operation. As such, it is beneficial to operate the digital TX system at as high a sample rate as possible, to push the spectral replicas far outside of the band of interest. If they are far enough away, they will be attenuated by the inherent band-limited nature of the other system components such as, for instance, the antenna.

As mentioned in Section 3.1, the input to this design is an I/Q signal sampled at 600 MS/s. Therefor, spectral replicas are expected to appear at multiples of 600 MHz away from the carrier. The reason for this sample rate stems from the CORDIC, which performs the I/Q-to-ρ/ϕ conversion. Due to the complex digital circuitry contained within the CORDIC, it was deemed infeasible to operate it at a much higher sample rate than this. Thus, other means of suppressing the spectral replicas had to be considered.

A common solution to this problem is to upsample the data and filter it using an interpolation filter. Although many different types of interpolation filters exist, only k^{th}-order hold finite impulse response (FIR) interpolation filters were considered, due to their relative ease of implementation. The simplest of these filters is the ZOH interpolation filter. For an upsampling factor of L, its impulse response can be defined as

$$h_{ZOH}(n) = \begin{cases} 1 & \text{if } n \in [0, L-1] \\ 0 & \text{otherwise} \end{cases} \tag{3.3}$$

where n is an integer sample number. Note that if no explicit upsampling is performed, this is effectively the interpolation filter that is applied when viewing a discrete-time signal in the continuous-time domain.

Higher-order interpolation filters can be recursively derived from the ZOH impulse response as follows:

$$h_{kOH}(n) = \frac{1}{L} h_{(k-1)OH}(n) * h_{ZOH}(n) \tag{3.4}$$

where $h_{kOH}(n)$ and $h_{(k-1)OH}(n)$ are the impulse responses of the k^{th}-order hold and the $(k-1)^{\text{th}}$-order hold, respectively, and $*$ is the convolution operator.

In the frequency domain, the transfer function of a k^{th}-order hold interpolation filter is given by

$$H_{kOH}(f) = \text{sinc}^{k+1}\left(\frac{f - f_0}{f_s}\right) \tag{3.5}$$

where f_0 and f_s are the center frequency and sampling frequency, respectively.

In this design, an upsampling factor of four was chosen, such that the resulting signal is clocked at the LO frequency. Besides the unaltered ZOH response, first-order hold (FOH), second-order hold (SOH) and third-order hold (TOH) interpolation filters were considered. Their respective impulse responses can be seen in Figure 3.8a. In the figure, the responses are centered around $n = 0$ for easier comparison. The frequency domain responses, evaluated at baseband (i.e. $f_0 = 0$), are shown in Figure 3.8b.

As discussed above, the phase needs to be converted to an I/Q signal before it can be fed to the phase modulator. This gives rise to two distinct ways to interpolate the phase signal:

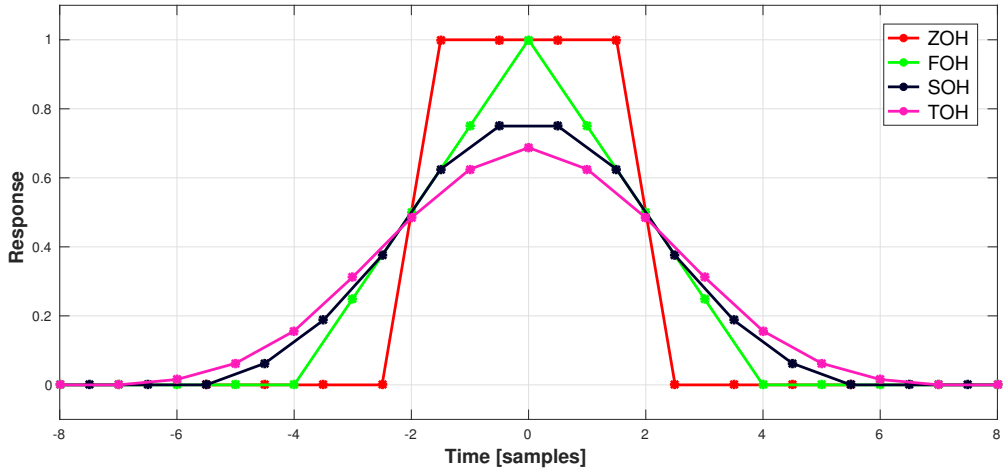

(a)

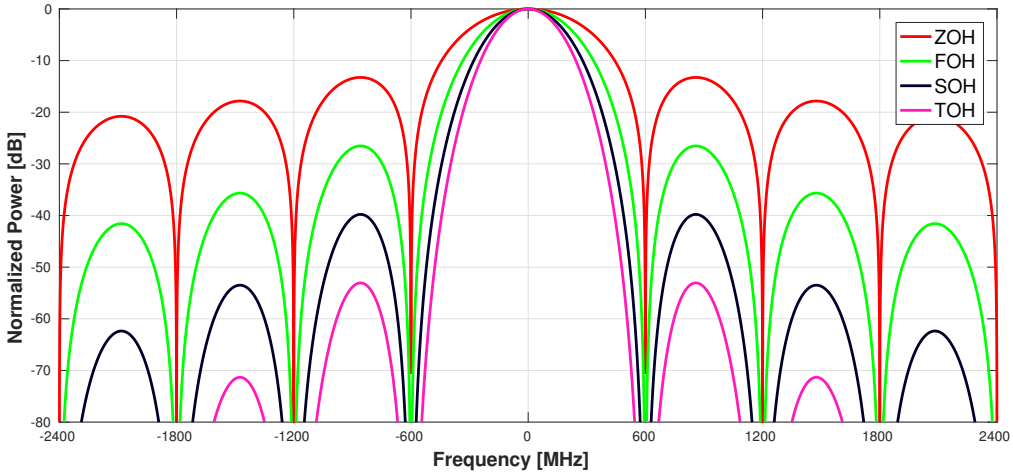

(b)

Figure 3.8: **(a)** Impulse responses and **(b)** frequency responses of ZOH, FOH, SOH and TOH interpolation filters

either before the conversion (in the phase domain), or after the conversion (in the I/Q domain). These two situations are presented graphically in Figure 3.9. Both cases were considered for this design. In terms of implementation, the former option has two major drawbacks. Firstly, the interpolation filter would need to operate on the unwrapped phase signal. This would most likely require the use of modulo arithmetic, as the true unwrapped phase signal is theoretically unbounded. Secondly, the LUT which is responsible for the ϕ-to-I/Q conversion would have to run at 2.4 GS/s, which would increase both the power consumption and the design complexity. By interpolating the phase in the I/Q domain, these issues can be avoided. However, by doing this, the resulting I/Q signal is no longer constant envelope, which may pose problems for the phase modulator.

To compare the effects of the various interpolation filters on the system performance, a simulation was run with an 80 MHz 256-QAM signal sampled at 600 MS/s. Both the envelope

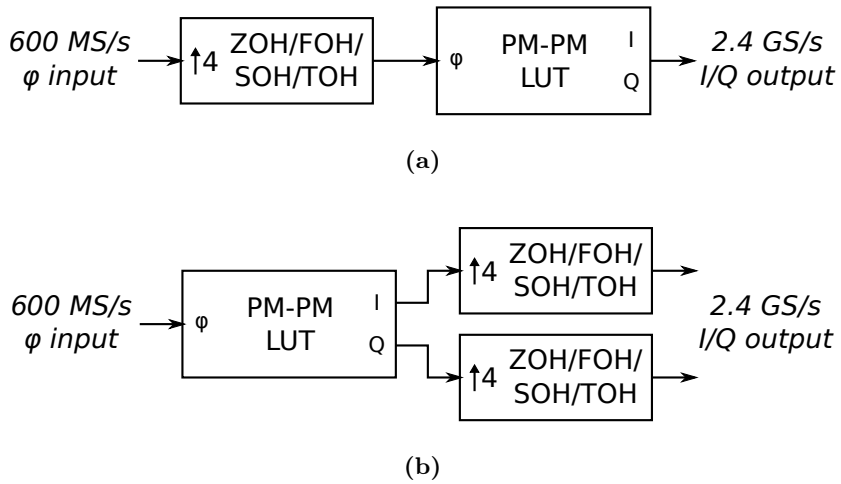

Figure 3.9: Two ways of interpolating the phase signal: **(a)** in the phase domain, **(b)** in the I/Q domain

and phase signals were upsampled to 2.4 GS/s and filtered with one of the selected interpolation filters. The resulting spectra and constellation diagrams in the case of phase interpolation in the phase domain can be seen in Figure 3.10. The same plots for phase interpolation in the I/Q domain can be found in Figure 3.11.

As expected, in both cases, the ZOH interpolation has a negligible effect on system performance, except for the spectral replicas at multiples of 600 MHz away from the carrier. Higher-order interpolation filters suppress these replicas but sacrifice both in-band and out-of-band performance. As shown in Figure 3.10a, if the phase is interpolated in the phase domain, FOH interpolation suppresses the replicas by approximately 20 dB. However, both the in-band and out-of band noise floors are significantly degraded. As can be seen in Figure 3.10b, the in-band degradation results in an EVM of approximately -41 dB. With SOH interpolation, spectral replicas are attenuated even further, as is the out-of-band noise floor. In-band, however, the EVM is degraded by approximately 3.5 dB. Finally, with TOH interpolation, the out-of-band noise floor is even lower and the spectral replicas are no longer visible. However, this comes at the cost of another 2.4 dB of EVM performance.

If the phase is interpolated in the I/Q domain, the effect of interpolation on the spectral replicas is similar, as shown in Figure 3.11a. However, the out-of-band noise floor seems to be limited to approximately -65 dB. As a result, although there is some benefit to using SOH over FOH in terms of suppression of the spectral replicas, the out-of-band performance gain of TOH interpolation is negligible. The EVM, shown in Figure 3.11b, follows a similar trend as in the previous case. It should be noted, however, that for FOH, SOH and TOH, the EVM is roughly 4 dB better than if the phase is interpolated in the phase domain.

As these results have shown, suppressing the spectral replicas by interpolation of the input signals comes at the cost of perfomance in other areas. For this design, it was chosen to do the phase interpolation in the I/Q domain, partly because it is easier and more power-efficient to implement, and partly due to the somewhat better in-band performance. Of the different interpolation filters, SOH was chosen because it attenuates the replicas almost down to the noise floor, while maintaining a better in-band performance than TOH. All in all, it can be seen that of the three different sources of error discussed in this chapter, this is by far the most severe, and it leaves little to no margin for achieving the target specifications. Yet, it was deemed a

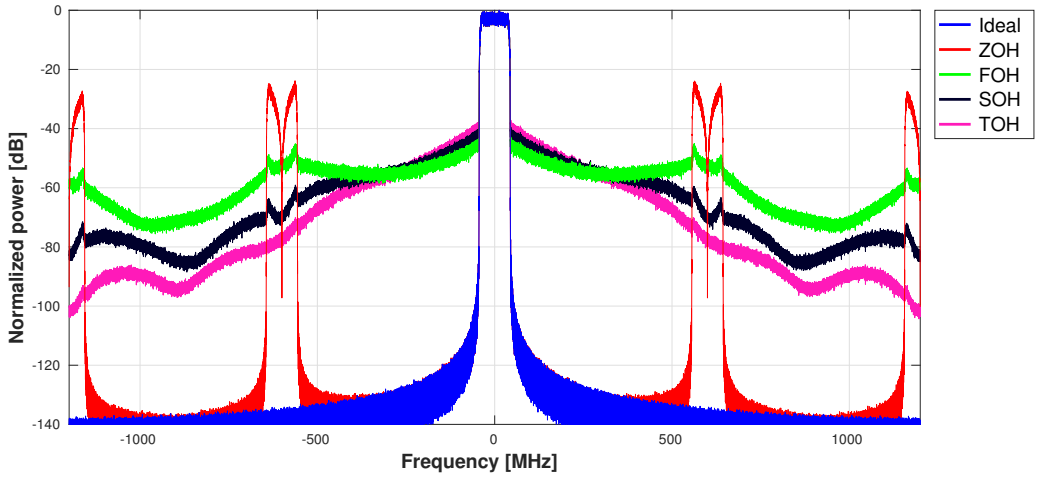

(a)

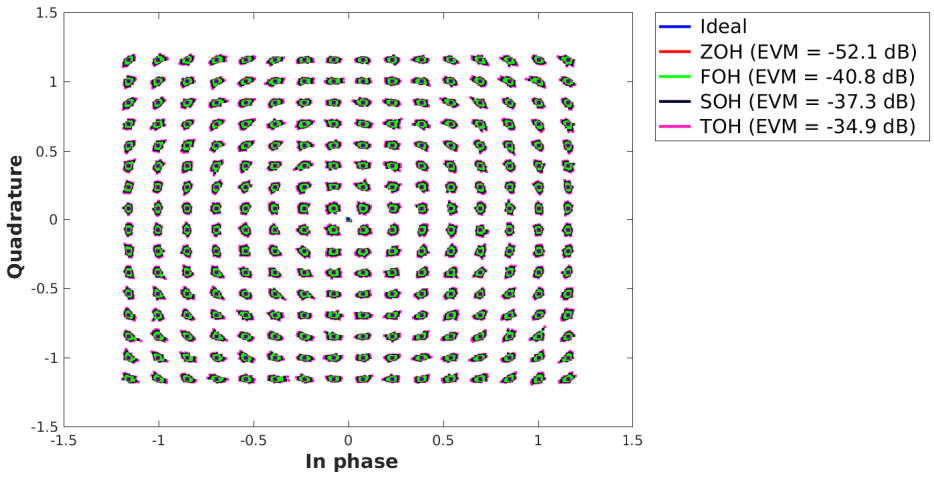

(b)

Figure 3.10: **(a)**Power spectra and **(b)** constellation diagrams of the RF output with phase interpolation in the phase domain

necessary measure to ensure that the spectral replicas can be suppressed. In order to make sure that the system performance could be evaluated without being limited by the input signal, an option was added to bypass the interpolator entirely. In this way, if the spectral replicas can be tolerated, the system performance can be measured with a much cleaner input signal.

3.5 Conclusion

This chapter presented a system-level overview of the transmitter system including the digital processing blocks that precede the envelope and phase modulators. The issue of bandwidth expansion in polar transmitters was examined, and it was found that the phase modulator needs to support a modulation bandwidth of at least six times the signal bandwidth in order to meet the ACPR requirement. The quantization noise caused by a finite phase resolution was also

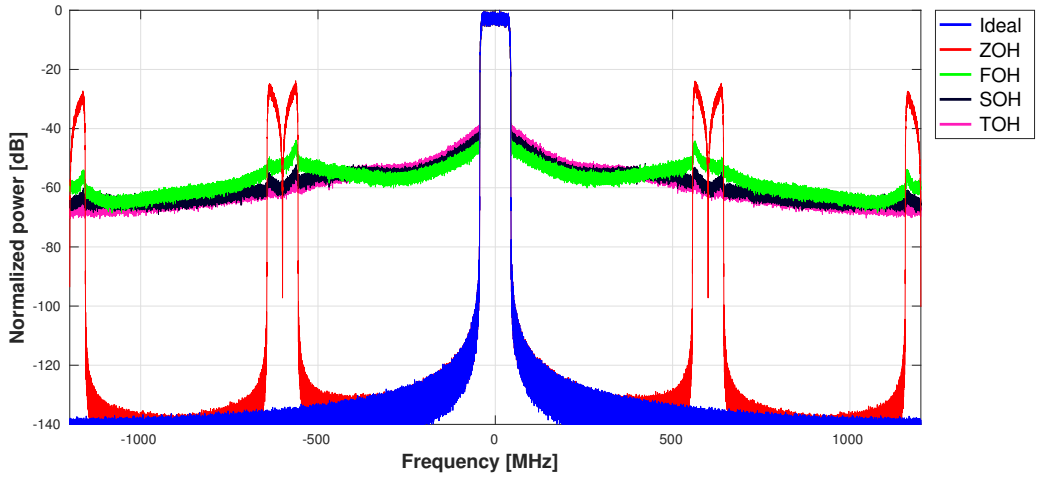

(a)

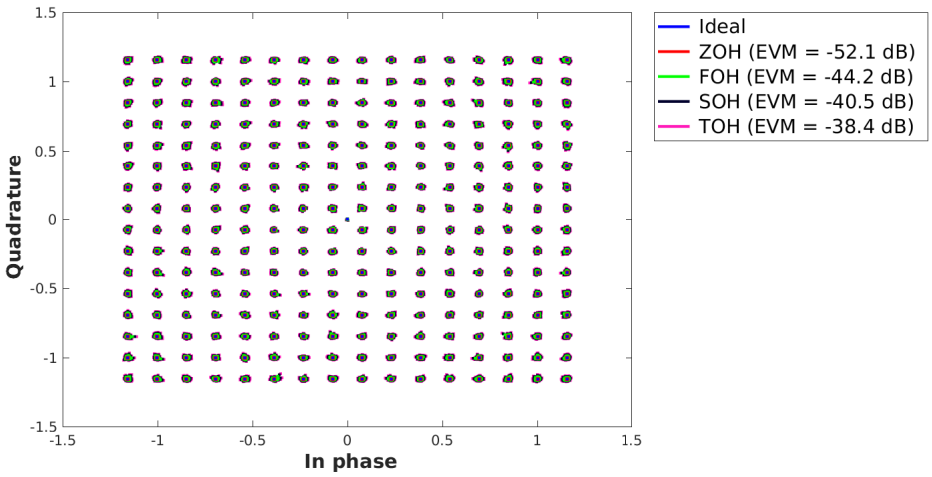

(b)

Figure 3.11: **(a)**Power spectra and **(b)** constellation diagrams of the RF output with phase interpolation in the I/Q domain

considered. This primarily affects the achievable out-of-band noise floor. A resolution of 10 bits was chosen for the phase signal which is then converted to a constant envelope I/Q signal with 9-bit resolution. This provides approximately a 10 dB margin to the target noise floor.

Finally, it was shown that a digital TX system exhibits spectral replicas at integer multiples of the sampling rate. These replicas need to be either pushed far enough outside the band of interest, or they must be attenuated by other means, to not interfere with other communication channels. The former was not an option for this design due to the speed limitation of the CORDIC, so SOH interpolation was employed to suppress the replicas. In the phase path, this interpolation was done after ϕ-to-I/Q conversion to simplify the design, save power and to achieve better in-band performance. Unfortunately this interpolation scheme limits the achievable system performance providing almost no margin to meet the target specifications. An option to bypass the interpolators was added to be able to evaluate the performance of the system when it is not

limited by the input signal, but in this case the spectral replicas need to be tolerated.

Chapter 4

Conceptual Design

As was explained in Chapter 2, the constant-envelope direct upconversion architecture (Figure 2.8) was considered the most suitable starting point for the phase modulator developed in this thesis work. Chapter 3 presented the system-level perspective and set some requirements on the attainable modulation bandwidth and resolution of the phase path. It also showed the digital processing blocks that precede the modulator and how they affect the performance. This chapter deals with the internals of the phase modulator itself, outlining the steps taken to arrive at the final conceptual design.

First, the starting point for this design, a modified version of the constant-envelope direct upconversion architecture, is presented in Section 4.1. Section 4.2 discusses the importance of low-pass filtering in this architecture and shows that restricting the design to only using RC filtering limits the achievable performance. Section 4.3 proposes a solution to this issue: using a harmonic rejection architecture to significantly relax the filtering requirements. Finally, the chapter is summarized and several conclusions are drawn in Section 4.4.

4.1 Phase Modulator Architecture Overview

Although the general architecture of the chosen phase modulator concept was covered in Section 2.2.4, it is instructive to consider how it can best be applied in the context of this TX system. The modified phase modulator architecture which served as the starting point for this design is shown in Figure 4.1. The ϕ-to-I/Q converter, which was covered in Chapter 3, is included in the figure for completeness.

At the core of the phase modulator is an I/Q radio frequency digital-to-analog converter (RF-DAC), which consolidates the mixing operation, the digital-to-analog conversion and the summation of the I and Q paths into a single circuit block [47]. The DACs in this block are based on the well-known current steering architecture. This is a common choice for high-speed designs since its speed is only limited by the pole of the output impedance [48]. These DACs are fully differential and operate in current mode. As such, the summation can easily be achieved, simply by connecting the differential output nodes together. As can be seen in Figure 4.1, the mixing operation occurs in the digital domain, before the digital-to-analog conversion. This is common practice for RF-DAC designs ([49]–[51]) and will be explained in more detail in Chapter 5.

Since the RF-DAC employs digital mixing, its output will be the sum of two square waves of varying amplitudes. This means that apart from the fundamental (LO) frequency the signal

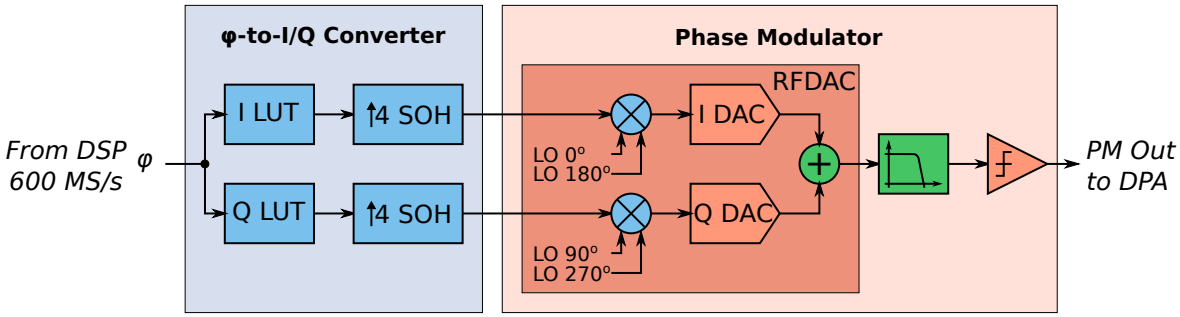

Figure 4.1: Modified constant-envelope direct upconversion architecture — the starting point for this design

will contain significant odd-order harmonic content. As will be discussed in more detail in the following section, it is vital that these higher harmonics are filtered out. Therefor, the output of the RF-DAC is fed into a low-pass filter (LPF).

The final stage of the phase modulator is a limiting amplifier which amplifies the filtered signal until it clips at the supply voltage. Provided the amplifier has enough gain, the output will approximate a square wave, the phase of which is specified by the input signal. This square wave can then be used to drive the phase input of the digital PA.

4.2 The Importance of Filtering

As explained above, it is essential that the output of the RF-DAC is low-pass filtered before it is fed to the limiting amplifier. The reason for this stems from the fact that, ideally, a limiting amplifier approximates the sign function:

$$\operatorname{sgn}(x) := \begin{cases} -1 & \text{if } x < 0 \\ 0 & \text{if } x = 0 \\ 1 & \text{if } x > 0 \end{cases} \tag{4.1}$$

Thus, it can be said that a limiting amplifier operates on the zero-crossings of the input signal. The output of the RF-DAC, however, is a summation of two square waves, in which only the fundamental component is the signal of interest. Due to the higher harmonics present in this signal, its zero-crossings will be unrelated to the zero-crossings of the fundamental, which will result in a phase error at the output of the limiting amplifier.

To demonstrate this, the time-domain waveforms of the various signals were analyzed for an input phase of $\phi_{in} = 30°$. Figure 4.2a shows the waveforms of the I and Q components as they appear at the outputs of their respective DACs. The I and Q waveforms are square waves with a frequency equal to f_{LO} and their amplitudes are given by $\cos(\phi_{in})$ and $\sin(\phi_{in})$, respectively. The Q waveform leads the I waveform by a quarter of the LO period, which constitutes a 90° phase shift.

The sum of these two waveforms is shown as the green curve in Figure 4.2b, while the black curve is its fundamental component, which is the desired signal. This clearly shows that if the output of the RF-DAC is fed directly to the limiting amplifier without any filtering, it would produce erroneous results. In fact, if no filtering is applied, then for any ϕ_{in} the output of the

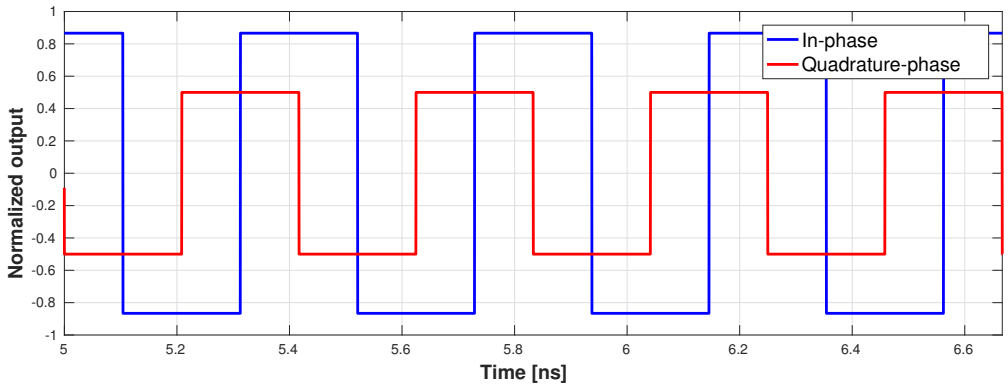

(a)

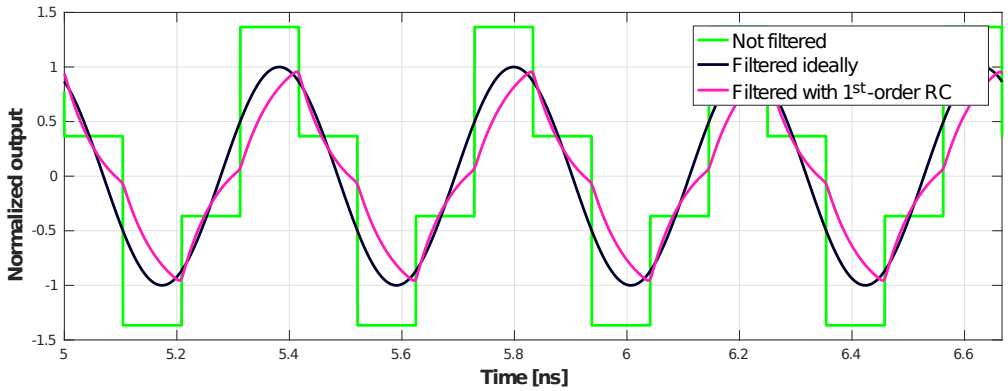

(b)

Figure 4.2: Time-domain waveforms of **(a)** the I and Q components and **(b)** their unfiltered and filtered sum for $\phi_{in} = 30°$

limiting amplifier can only assume one of four discrete phases, namely: 0°, 90°, 180° or 270°, since these are the only phases at which the zero-crossings can occur.

It should now be clear that some form of filtering is required to achieve adequate phase modulator performance. Typically, in low-GHz RF designs, on-chip filtering is accomplished by passive RC or LC filters. LC filters are generally considered to be superior because they allow for complex poles in the filter transfer function. This allows for finer control of the filter characteristics. However, as discussed in Section 1.5, one of the goals of this thesis work is to arrive at a phase modulator design which does not include inductors. Thus, LC filters were not an option for this design.

Complex poles could also be achieved by employing active filters, which consist of amplifiers in feedback configuration. The filter characteristics are then set by the RC feedback network. However, the amplifiers must have a very high gain to ensure good performance, which is difficult to achieve at such high frequencies. Therefor, active filters were not considered for this design.

Thus, it was decided to do the filtering with RC filters only. Only 1^{st}-, 2^{nd}- and 3^{rd}-order filtering was considered, as higher order filters were deemed impractical. To visualize what can be expected from RC filtering, the RF-DAC output in the example above was filtered with a 1^{st}-order RC filter with a cut-off at the LO frequency. The result was the pink curve in Figure 4.2b.

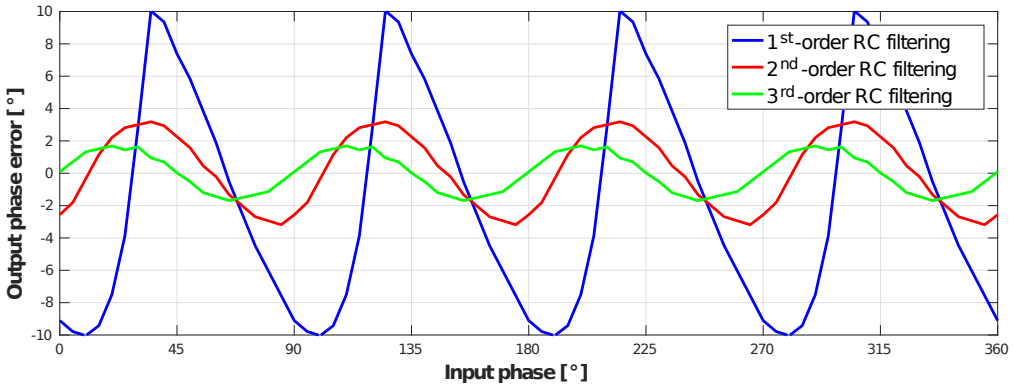

Figure 4.3: Static phase error at the output of the limiter for 1st-, 2nd- and 3rd-order RC filtering

As the figure shows, although the filtered waveform more closely approximates the desired signal than the unfiltered RF-DAC output, sharp transitions are still clearly visible at the time points where the RF-DAC output changes value. This leads to a phase error at the output of the limiting amplifier.

To determine the extent of this phase error, a transient simulation was run with a close-to-ideal RFDAC the output of which was filtered by a cascade of RC filter stages. Each stage had a cut-off at the LO frequency. This was considered a decent compromise between poor attenuation of harmonics if the cut-off frequency is too high and too strong attenuation of the fundamental if the cut-off frequency is too low. Buffers consisting of ideal voltage-controlled voltage sources were inserted between the stages to eliminate any loading effects. Non-idealities in the RF-DAC were kept to a minimum and were only introduced to achieve reliable convergence in the simulation. To emulate the behavior of an ideal limiter, the sign function was applied to the output of the filter.

In the simulation, the input phase ϕ_{in} was swept from 0° to 360° in 5° steps. For each ϕ_{in}, the output was allowed to settle to steady-state and the waveforms were analyzed. The phase error at the output of the limiter was computed by comparing the zero-crossings of the output to those of an ideal square wave with a phase equal to ϕ_{in}. Note that any causal filter also introduces a constant phase shift. This phase shift is not considered to be part of the phase error, since it can easily be corrected for by a programmable delay in the envelope path before the envelope and phase paths are recombined at the input of the DPA. In this analysis, this phase shift was computed in such a way that minimum and maximum values of the phase error were of equal magnitude (i.e. so that the error is centered around 0°).

The simulation results can be seen in Figure 4.3. As the figure shows, if only 1st-order RC filtering is used, the phase error is as large as ±10°. For 2nd- and 3rd-order filtering the maximum phase error is ±3.2° and ±1.8°, respectively. As can be seen, the phase error as a function of ϕ_{in} is periodic with a period of 90°. For 2nd- and 3rd-order filtering in particular, the phase error ϕ_ϵ can be modeled by a sine wave as

$$\phi_\epsilon(\phi_{in}) = A_\epsilon \sin(4\phi_{in} + \phi_{\epsilon,0}) \tag{4.2}$$

where A_ϵ is the amplitude of the error and $\phi_{\epsilon,0}$ represents a constant shift along the ϕ_{in} axis.

To investigate what phase error can be tolerated while still meeting the target requirements, a system-level simulation was run with an 80 MHz 256-QAM signal. The phase modulator had a distortion profile as defined in Equation 4.2. In this simulation, A_ϵ was swept from 0.5°

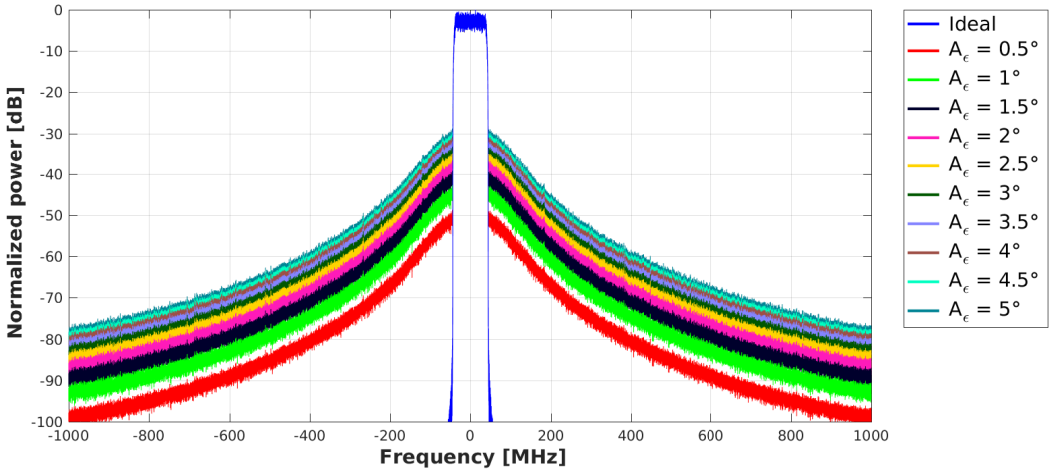

(a)

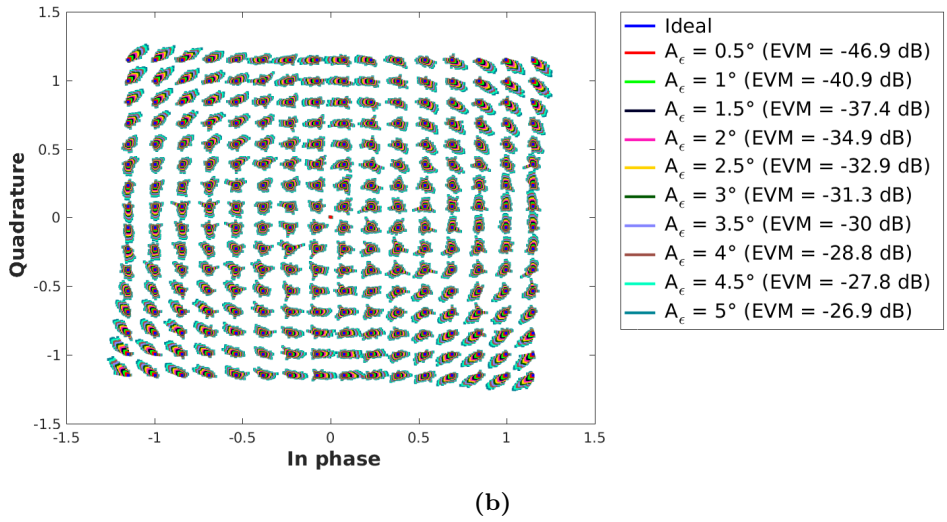

(b)

Figure 4.4: **(a)**Power spectra and **(b)** constellation diagrams of the RF output with a sinusoidal phase error profile with varying amplitudes

to 5° in 0.5° steps, whereas $\phi_{\epsilon,0}$ was set to 0°. Although various values of $\phi_{\epsilon,0}$ were initially considered, the simulation results were all within 0.5 dB of each other, so the effect of $\phi_{\epsilon,0}$ was considered to be negligible. The resulting power spectra and constellation diagrams of the RF output are shown in Figure 4.4. As the figure shows, A_ϵ must be less than 1° to satisfy the ACPR requirement, although less than 0.5° is desirable to have some margin. For the EVM, the performance degradation due to phase error is somewhat less severe, with $A_\epsilon = 1°$ resulting in a 6 dB margin.

These simulation results clearly show that even 3$^{\text{rd}}$-order RC filtering under near-ideal conditions would not be able to meet the targeted system requirements. In reality, however, the phase error would be even worse due to, for instance, DAC nonlinearity. Furthermore, in this simulation only the static phase error was considered, i.e. the phase error after settling. In a practical circuit, due to memory effects, a dynamic phase error would also be introduced, which

would degrade the system performance even further. Thus, another solution to this filtering issue had to be found.

4.3 The Harmonic Rejection Architecture

In light of the problems discussed above, another solution was sought to be able to sufficiently suppress the harmonics in the RF-DAC output while being restricted to RC filtering only. Upon further investigation, it was found that the dominant source of error stemmed from the 3^{rd} and the 5^{th} harmonics, as these were only moderately attenuated by the RC filters alone. To solve this problem, the harmonic rejection architecture was introduced. This was first proposed in [52] as a mixer architecture, but is equally applicable to RF-DACs. By employing three RF-DACs with proper gain scaling and their LO inputs shifted by 45° relative to one another, the 3^{rd} and the 5^{th} harmonics can be canceled out. This would significantly relax the filtering requirements and make RC filters a viable solution. A block diagram of this concept is shown in Figure 4.5a. It shows the three RF-DACs connected in parallel but supplied with different clocks. RF-DACs 1 and 3 are set to have a normalized gain of 1, whereas RF-DAC 2 has a gain of $\sqrt{2}$. For brevity, this architecture is termed the harmonic rejection RF-DAC (HR-RF-DAC), and is referred to as such throughout the rest of this thesis work.

The harmonic cancellation operation resulting from this architecture can be explained by examining the phasor diagrams at the fundamental frequency and the 3^{rd} and 5^{th} harmonics. These phasor diagrams can be seen in Figure 4.5b. If it is assumed that the output of RF-DAC 1 has a fundamental component of $1\angle 0°$, then the outputs of RF-DACs 2 and 3 will have fundamental components of $\sqrt{2}\angle 45°$ and $1\angle 90°$, respectively. The sum of these three phasors is $2\sqrt{2}\angle 45°$, thus, for the fundamental component the signals are added constructively.

For higher harmonics, a similar analysis is performed, but because the frequency of the n^{th} harmonic is n times the frequency of the fundamental, the phase shift between the phasors will be $n \cdot 45°$. As can be seen in Figure 4.5b, for the 3^{rd} and 5^{th} harmonics, this results in the outputs of the three RF-DACs adding destructively, thereby canceling out the harmonics.

The effect of harmonic rejection can also be seen in the time domain. An input phase of $\phi_{in} = 30°$ was taken as an example. The time-domain output of the HR-RF-DAC is depicted in Figure 4.6. Also shown in the figure are the waveforms after filtering, both with an ideal brickwall filter and with a 1^{st}-order RC filter. It can be seen that even without filtering, the HR-RF-DAC output resembles a sine wave much more closely that the output of a single RF-DAC (Figure 4.2b). Nevertheless, the zero-crossings can only occur at eight discrete time points during a period, so filtering is still required. Similarly, when this waveform is filtered with a 1^{st}-order RC filter with a cut-off at the LO frequency, the waveform approximates that of the ideal signal rather well, albeit with a slight phase shift.

To determine the filter order needed to satisfy the static error requirement, the transient simulation described in the previous section was run again, now with the HR-RF-DAC architecture. The resulting static phase error ϕ_ϵ as a function of ϕ_{in} is shown in Figure 4.7. Compared to the case without harmonic rejection (Figure 4.3), the situation is considerably improved. With only 1^{st}-order RC filtering the error is limited to $\pm 2.2°$. For 2^{nd}- and 3^{rd}-order filtering, the amplitudes of the error are 0.37° and 0.14°, respectively.

It should be noted that due to the introduction of four additional LO phases, the period of $\phi_\epsilon(\phi_{in})$ has decreased to 45°. The system-level simulation from the previous section was run with this new distortion profile for various error amplitudes. The results are illustrated in Figure 4.8. Comparing this to the case without harmonic rejection (Figure 4.4), it can be seen that for

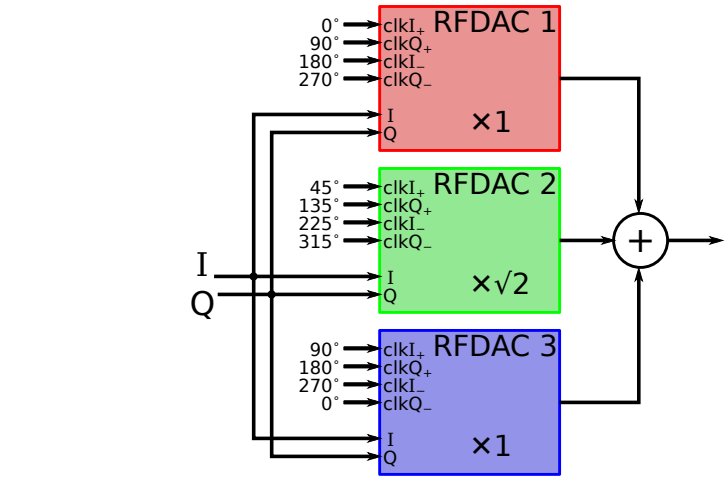

(a)

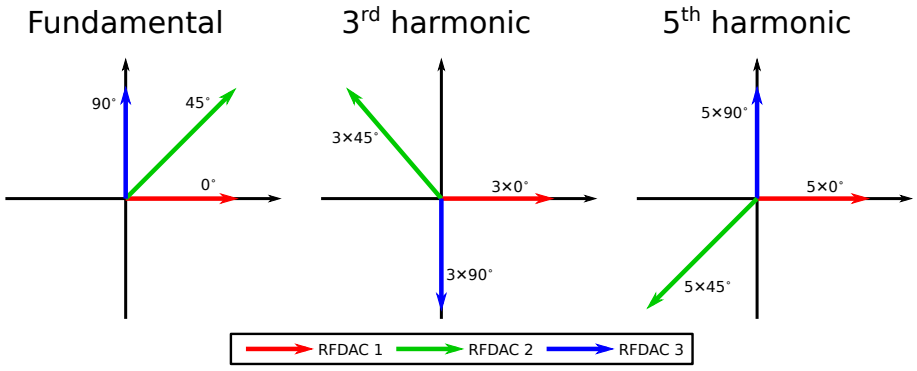

(b)

Figure 4.5: The harmonic rejection architecture showing **(a)** the block diagram and **(b)** the phasor diagrams for the fundamental frequency and the 3rd and 5th harmonics

the same A_ϵ both the EVM and the close-in spectrum show an improvement of approximately 3 dB. This comes at the cost of worse far-out performance. Nevertheless, the static phase error specification of $A_\epsilon < 0.5°$ remains the same, in order to leave enough margin for other error sources. As shown in Figure 4.7, 2nd-order RC filtering is sufficient to meet this specification, and was therefor chosen for this design.

4.4 Conclusion

This chapter discussed the importance of filtering the output of the RF-DAC in order to minimize the static phase error of a phase modulator based on the constant-envelope direct upconversion architecture with digital mixing. It was shown that if the design is restricted to using only RC filtering, even a 3rd-order filter is not enough to meet the EVM and ACPR specifications. However, by making use of the harmonic rejection architecture to cancel out the 3rd and the 5th

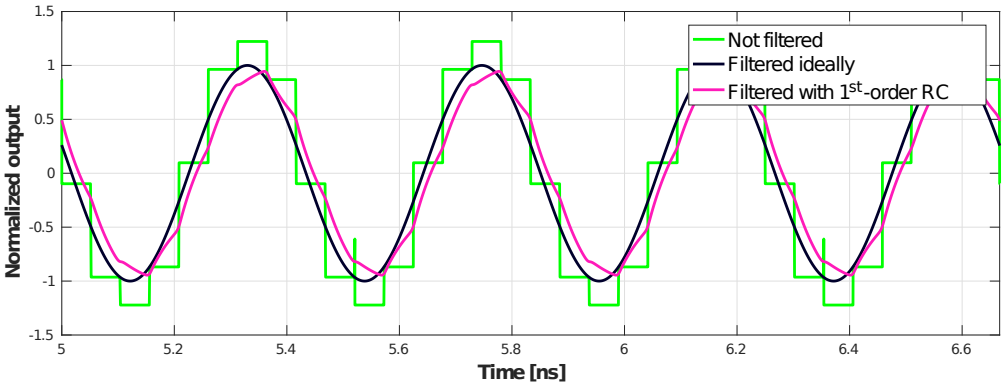

Figure 4.6: Time-domain waveforms at the output of the HR-RF-DAC before and after filtering for $\phi_{in} = 30°$

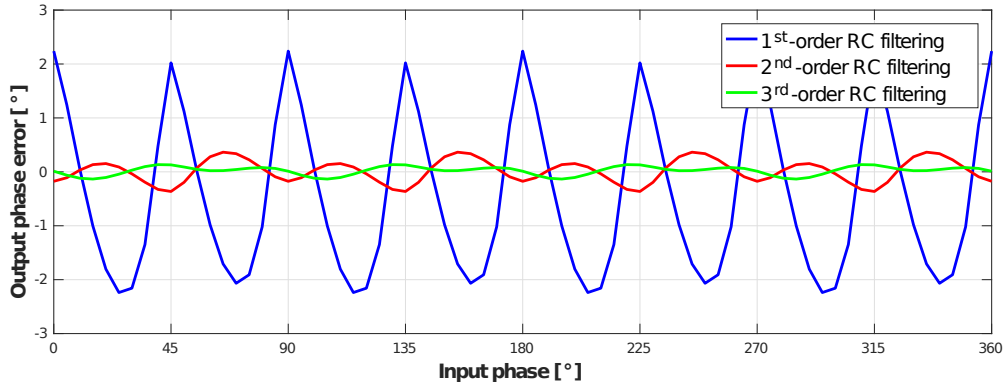

Figure 4.7: Static phase error at the output of the limiter for 1st-, 2nd- and 3rd-order RC filtering with harmonic rejection

harmonics, the filtering requirements can be significantly relaxed. This makes it possible to meet the requirements with only a 2nd-order RC filter. This architecture requires the use of three RF-DACs rather then one, with the clocks of each successive RF-DAC shifted by 45°. Additionally, the gain of the second RF-DAC has to be scaled by a factor of $\sqrt{2}$, relative to the other two RF-DACs. This is reflected in Figure 4.9, which shows the revised phase modulator concept chosen for this design. The term harmonic rejection RF-DAC (HR-RF-DAC) was introduced to describe the combination of these three RF-DACs as a single circuit block.

Use of the HR-RF-DAC architecture comes at the cost of a larger area and a higher power consumption due to the use of three RF-DACs rather then one. Furthermore, this architecture requires eight different LO phases instead of four, which increases the power consumption even further. However, if this phase modulator concept is used in conjunction with a high-power PA, as is the case in base-station applications, then these disadvantages will have a minimal effect on system efficiency and overall cost. Thus, these issues were tolerated to be able to meet the in-band and out-of-band spectral purity requirements.

For RC filter used in this design consists of two RC low-pass stages, each with a cut-off frequency at the LO frequency, namely 2.4 GHz. This ensures good suppression of higher harmonics with only moderate attenuation of the fundamental component. At first glance, this may

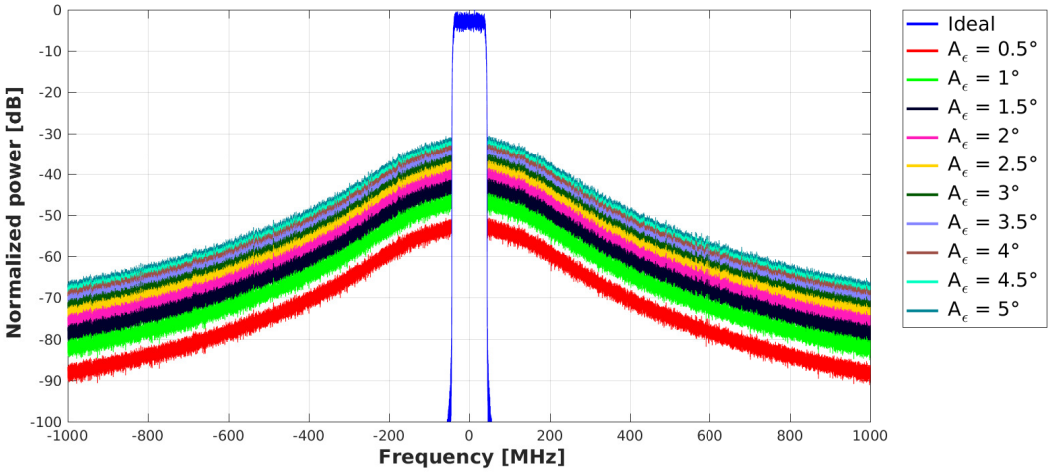

(a)

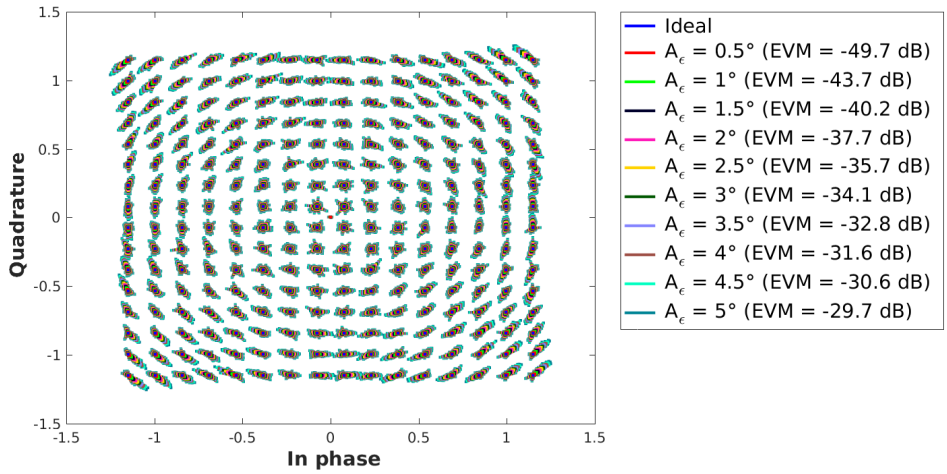

(b)

Figure 4.8: (a) Power spectra and **(b)** constellation diagrams of the RF output with a sinusoidal phase error profile with varying amplitudes, in the case of harmonic rejection

seem to conflict with the modulation bandwidth requirement of 480 MHz specified in Section 3.2. However, as the transfer function of a 2$^{\text{nd}}$-order RC filter exhibits a very gentle roll-off around the cut-off frequency, neither the attenuation nor group delay vary significantly in the 480 MHz band centered at the LO. Thus, the filter will not have a significant performance impact on the modulated signal.

It should be noted that, apart from the filter itself, this concept is frequency-independent. At the high side, the operational frequency will still be limited by parasitics and the speed of the digital logic. Lower operational frequencies, however, can easily be achieved simply by lowering the LO frequency and changing the cut-off frequency of the filter. To accommodate for this, the capacitors used in the filter can be implemented as programmable capacitor banks. In this design, the lowest cut-off frequency of the filter was chosen to be 900 MHz, to limit the area required for the capacitor banks.

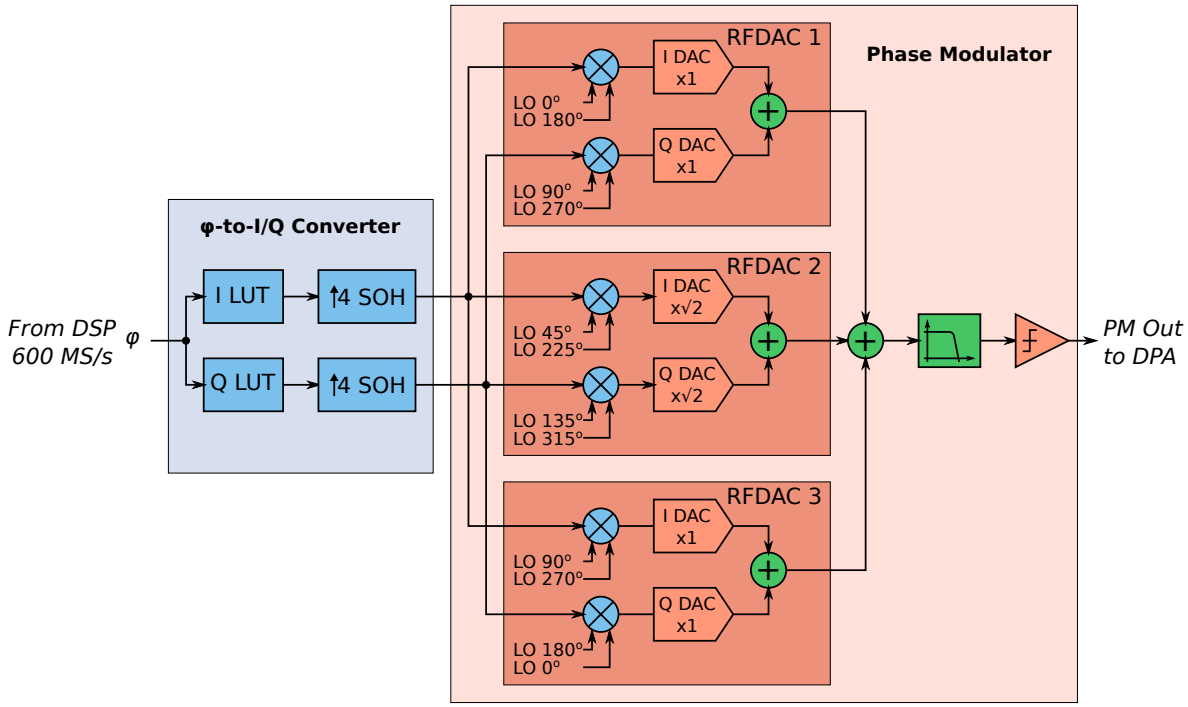

Figure 4.9: Revised phase modulator concept making use of the harmonic rejection RF-DAC architecture

Chapter 5

Detailed Design

The previous chapter presented the high-level phase modulator concept that was chosen for this design. This chapter describes the process of implementing this concept in 40 nm complementary metal-oxide-semiconductor (CMOS) 1-poly-7-metal TSMC technology with a 1.1 V supply.

First, a top-level overview of the system is presented in Section 5.1, showing the three major components of this design: the HR-RF-DAC, the filter and the limiting amplifier. The majority of the design effort is focused on the constituent mixing DACs of the HR-RF-DAC, and their design is detailed in Section 5.2. The designs of the limiting amplifier and the filter are described in Section 5.3 and Section 5.4, respectively. The layout of the individual blocks is shown in their respective section. The top-level layout combines these blocks to form the complete phase modulator implementation, and is discussed in Section 5.5. Finally, some conclusions are drawn in Section 5.6.

5.1 Top-Level Overview

The top-level schematic of the phase modulator developed during this thesis work is shown in Figure 5.1. As the figure shows, the inputs for this design are I and Q data words as well as eight phases of the LO clock, spaced 45° apart. As discussed in Chapter 3, the I and Q data is quantized to 9 bits and has a sample rate of 2.4 GS/s. The output of this phase modulator is a differential rail-to-rail square wave at the LO frequency with a modulated phase based on the I/Q inputs. In the figure, this signal is designated $v_{o,lim}$ and is the output of the limiting amplifier.

As was discussed in Chapter 4, the core of this design is the HR-RF-DAC, which comprises a total of six DACs: three for I and three for Q. Each DAC includes a mixer in the digital domain, which upconverts the I or Q data from baseband up to the LO frequency. To avoid confusion, this combination of mixer and DAC is termed a 'mixing DAC', whereas simply 'DAC' is used to refer to the only the part after the mixer. It should be noted that from the mixing operation onwards, the entire phase modulator design is fully differential. This serves to eliminate any even-order harmonic distortion and effectively doubles the dynamic range of all signals in the voltage domain when compared to a single-ended topology. For differential mixing, however, both the LO and the I/Q inputs need to be differential as well. In the case of the LO signal, this is achieved by supplying each mixer with two different phases of the LO, spaced 180° apart. As the I and Q data is supplied from digital logic which is single-ended in nature, its conversion to a differential form is less straightforward, and will be discussed in more detail in Section 5.2.6.

The output of the mixing operation is a 9-bit digital word which is supplied as input to a current steering DAC. As was mentioned in Section 4.1, the current steering topology was chosen both for its high speed capabilities as well as the relative ease with which the output of multiple DACs can be summed together. Figure 5.1 shows the ideal model of such a DAC. The model comprises two current sources, with the output currents of these sources given by

$$I_{out+} = mI_0 \tag{5.1}$$
$$I_{out-} = (N-m)I_0 \tag{5.2}$$

where m is the input code, N is the maximum value this code can take ($N = 511$ for a 9-bit DAC), and I_0 is the reference current of a single unit cell. Note that in the case of I-DAC 2 and Q-DAC 2, the output is scaled by a factor of $\sqrt{2}$, as is required for harmonic rejection. The outputs of all six DACs are summed by connecting them at the same two nodes, such that the DACs are connected in parallel. This produces the differential output current of the HR-RF-DAC (superimposed on a DC current component).

The load for the HR-RF-DAC is a differential RC filter. This filter is composed of two branches, each consisting of a resistor and capacitor connected in parallel, with one end connected to the supply. The capacitors are made programmable such that different cut-off frequencies can be selected for use with different LO frequencies. As discussed in Chapter 4, this filter serves to filter out higher harmonics resulting from non-linear mixing, which would otherwise degrade the phase modulator performance. Furthermore, the filter converts the current-mode output of the HR-RF-DAC to a differential voltage, $v_{o,DAC}$, which is supplied as input to the limiting amplifier. The limiting amplifier amplifies this voltage until it clips at the supply rails, producing the phase-modulated square wave, $v_{o,lim}$, which is the output of this phase modulator.

One thing to note is that the filter in this design is implemented as a single RC stage, even though it was concluded in Section 4.3 that two stages are required to meet the target specifications. During the design of the limiting amplifier, it was found that several gain stages were needed in order to achieve a high enough amplification. Each stage had a voltage gain of only approximately 3 to 4, and as a result, the signal at the output of the first stage was small enough that the amplification could be considered to be linear. Furthermore, the bandwidth of the first stage was only slightly higher than 2.4 GHz. It was decided to make use of these characteristics and thus the first stage of the limiter now additionally serves as the second stage of the filter. This not only saved additional design time (due to not needing to increase the bandwidth) but also isolated the filter stages, removing any loading effects. An additional capacitor was added to load the output of the first stage, setting the pole of the transfer function to 2.4 GHz. Much like in the case of the RC filter, this capacitor was made programmable, to maintain frequency-agile operation.

5.2 Mixing DAC Design

The mixing DAC in this design integrates a current-steering DAC and a mixer in a single circuit block. As was discussed in the previous section, six such mixing DACs make up the HR-RF-DAC, where two of the mixing DACs have a gain a factor of $\sqrt{2}$ larger than the other four. This section describes the design process of the DAC and mixer components as well as the associated digital circuitry, both at the schematic and layout levels.

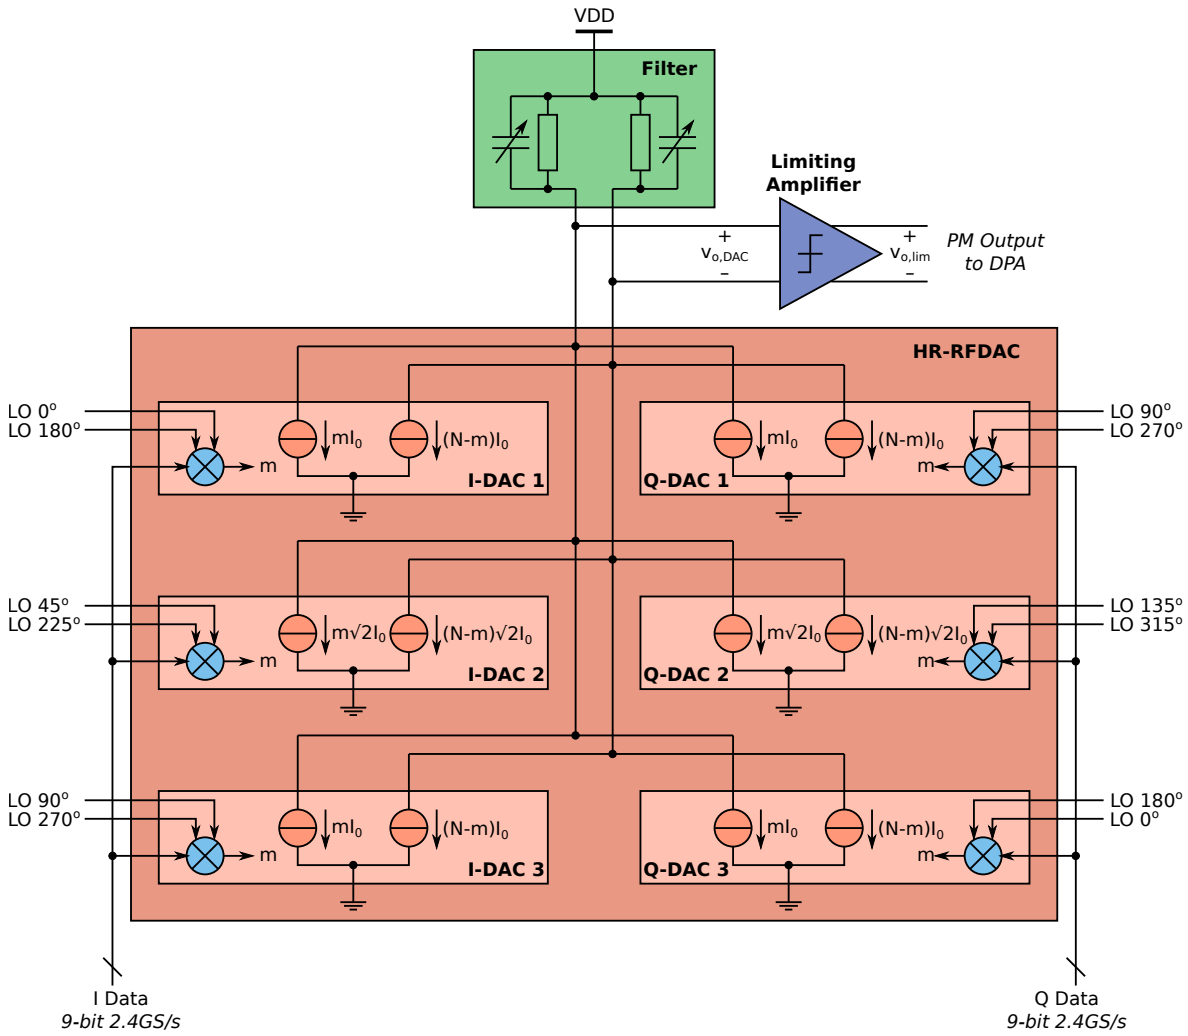

Figure 5.1: Top-level schematic of the phase modulator

5.2.1 Current-Steering DAC

The principle of operation of the current-steering topology is shown in Figure 5.2a. Such a DAC consists of a set of DC current sources each connected to a switch. The input code m is a digital word consisting of n bits $(m_0 \dots m_{n-1})$ where each bit controls the switch of the corresponding current source and connects the current source to either the positive or the negative output. The DAC output ports are connected through a load to the supply voltage. A number of different load configurations are possible, including coupled inductors or a transimpedance amplifier, but as discussed above, for this design, the load is an RC filter.

The current-steering topology is a popular choice for high-speed DAC designs because the speed of this DAC is only limited by the pole in the output impedance [48]. In this case, the output impedance is the parallel combination of the load, the output impedance of the DAC itself (which is finite due to non-idealities of the current sources) and any interconnect parasitics. One of the reasons for the high speed that can be achieved with this topology is the fact that the current sources are always on, and do not need to settle to a new value whenever the input code

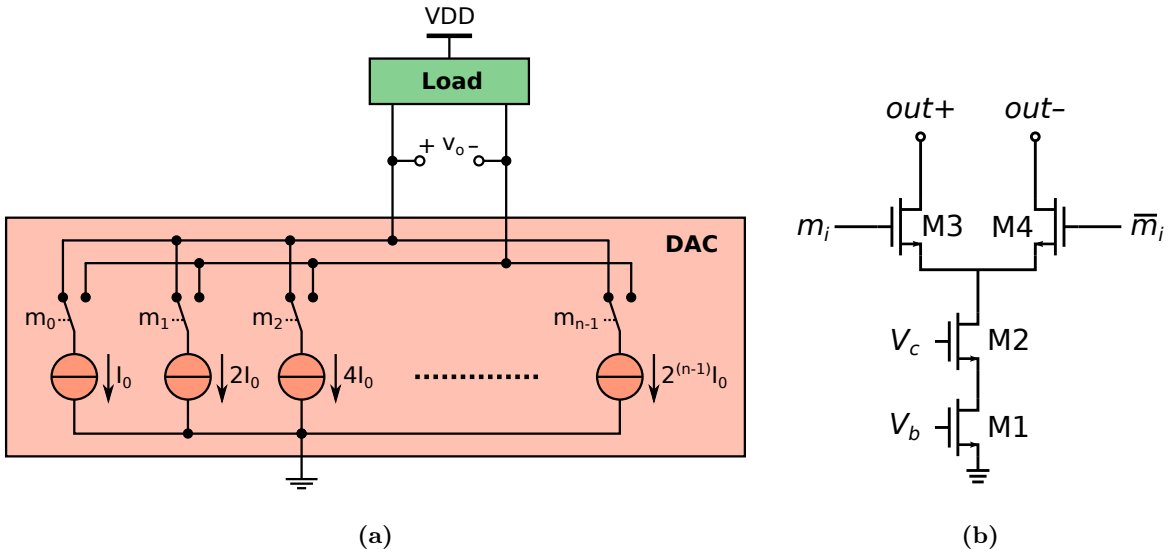

Figure 5.2: The current steering DAC topology showing **(a)** the principle of operation and **(b)** a typical implementation of a unit cell

of the DAC changes. The change in code only redirects the current to flow through a different path by using a digital switch. Such switches can be made very fast in modern deep submicron CMOS technologies, and usually do not limit the speed of the DAC.

The fact that the current sources are always on also means that the total current drawn from the supply remains constant. Assuming a linear load, this also means that the common-mode voltage of v_o stays constant. This common mode voltage can be set to a desired value by properly sizing the load resistors and the unit cell current I_0.

A typical implementation of a unit cell for a current-steering DAC is shown in Figure 5.2b. A unit cell contains a current source providing a DC current I_0 and a switch to steer this current through either the positive or negative output port. In the figure, the current source transistor is denoted as M1 and is biased in the saturation regime by bias voltage V_b. Although a single transistor can be a good enough current source in some applications, the analog performance of deep submicron CMOS devices suffers from poor output impedance, which would result in DAC non-linearity. For this reason, the current source is supplemented with a cascode transistor (M2) biased with V_c, which helps to increase the output impedance of the current source. This comes at the cost of a higher saturation voltage for the current source, which in turn reduces the available voltage headroom left for the signal. However, in modern CMOS technologies, it is often a necessary measure to achieve adequate DAC linearity. The switch is implemented as a digitally driven differential pair (M3 and M4). Note that the digital code needs to be available as a differential signal for this switch to function.

5.2.2 Digital Mixer

The baseband data needs to be upconverted to the LO frequency through mixing. Traditionally, this is done after the DAC using a double-balanced mixer configuration as shown in Figure 5.3. By providing a sinusoidal LO signal and with proper biasing, this circuit can perform a linear mixing operation. The phase modulator in [44] used this approach to eliminate higher harmonics in the output of the mixer, thereby relaxing the filtering requirements. The sinusoidal LO signal

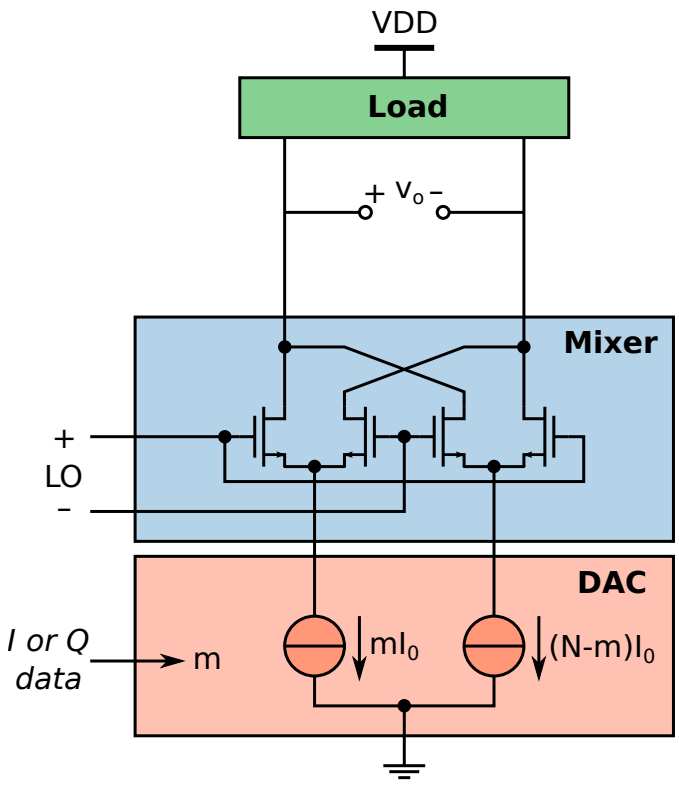

Figure 5.3: Typical mixing DAC implementation making use of a double-balanced mixer

was generated from a square wave by tuning the LO driver with an LC resonant circuit. This approach was dismissed for this thesis work, however, for a number of reasons. First of all, one of the goals of this work is to develop an inductor-less phase modulator concept. Using an LC tank for tuning would also make it difficult to achieve a wide range of operating frequencies. Furthermore, this design aims to be as digital as possible to benefit from technology scaling and using a linear mixer would only degrade the performance at finer technology nodes. Finally, due to their resistive nature, linear mixers are known to have poor noise performance. As such, in modern systems, switching mixers are generally preferred.

The mixer in Figure 5.3 can also be driven by a rail-to-rail square wave LO signal, such that the transistors act as switches. This eliminates the issues mentioned above but introduces higher-order harmonic content at the output of the mixer because the output signal itself will also be a square wave. As explained in Chapter 4, this is corrected for by the HR-RF-DAC architecture and a 2^{nd}-order RC filter.

Although this was the initial architecture considered for this phase modulator design, several drawbacks surfaced during the design process. First of all, the CMOS devices in the mixer core needed to be quite large to be able to handle the output current of the DAC, which resulted in a high parasitic capacitance at the output nodes. Furthermore, because the source voltage V_s of the mixer devices had to be at least as high as the saturation voltage of the current sources in the DAC (approximately 250 mV), the overdrive voltage V_{gt} (given by the difference between the gate-source voltage V_{gs} and the threshold voltage V_{th}) was significantly reduced. This, combined with the high gate-source capacitance C_{gs} made it difficult to achieve fast switching behavior.

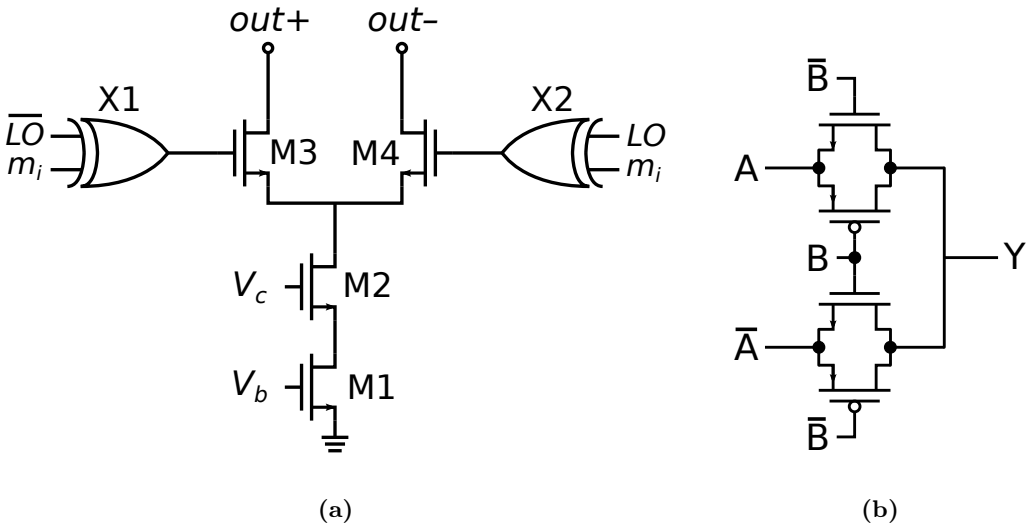

(a) **(b)**

Figure 5.4: Schematics of **(a)** the mixing DAC unit cell and **(b)** the XOR gate

Finally, due to the finite on-resistance of the devices, the mixer incurred a voltage drop which reduced the available voltage headroom.

To combat these issues, it was decided to remove the mixer from the transistor stack and instead implement it in the digital domain before the analog-to-digital conversion. This has the drawback of requiring the DAC to have a higher bandwidth, since the signal is now centered at the LO frequency rather than DC, but the current-steering DAC architecture is well-suited to handle that bandwidth. For this design, it was chosen to implement the mixing operation on a bit-by-bit basis inside the mixing DAC unit cell. The schematic of this unit cell is shown in Figure 5.4a. Rather than driving the switch transistors M3 and M4 directly with the appropriate data bit, as with a traditional current-steering architecture, the switches are driven by two XOR gates (X1 and X2). The inputs to these XOR gates consist of the data bit m_i and the positive and negative LO waveforms. For a constant m_i, the unit cell output current will be switched between the positive and negative output nodes at the LO frequency. When m_i changes state, the polarity of the output current waveforms will be reversed. In this way, digital mixing is achieved.

The schematic of the XOR gate used in this design is shown in Figure 5.4b. The XOR gate consists of two transmission gates, each consisting of an NMOS and a PMOS device. Input A was used for the data bit, whereas input B was used for the LO. To achieve the differential signal required for the switch transistors M3 and M4, the positive and negative LO inputs of X1 were reversed with respect to X2. Since each XOR gate only has to drive the gate of a single transistor, which as will be explained in Section 5.2.4 was set to minimum size, the load at the output of the XOR gates was very small. Because of this, all four transistors in the XOR gate could be made minimum size and still maintain the desired switching speed. Note that, although it is not made explicit in Figure 5.4a, both the LO and m_i inputs need to be provided as complementary signals for the XOR gate to function.

Although this mixing DAC architecture addresses the issues of the double-balanced mixer discussed above, it has a serious disadvantage that has not yet been considered. While a double-balanced mixer has excellent suppression of LO feed-through to the output, the same can not be

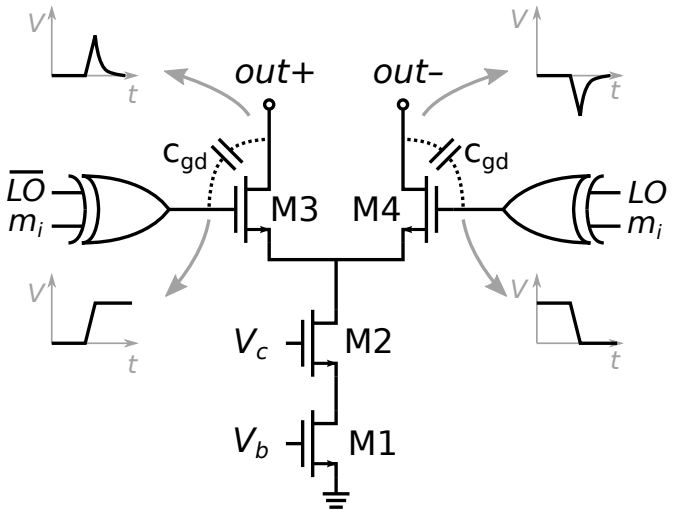

Figure 5.5: The effect of charge injection on the mixing DAC unit cell output

said of this topology. Due to the gate-drain capacitance C_{gd} of switching transistors M3 and M4, a sharp transition in the gate voltage will cause a disturbance at the drain of the device. This mechanism is known as charge injection and is illustrated in Figure 5.5. As shown in the figure, each rising or falling edge at the output of the XOR gates will cause some charge to couple to the output to the output nodes resulting in a voltage (and current) spike. Since the resulting error at the output of the unit cell is periodic with a frequency equal to the LO frequency, the DAC output will exhibit LO feed-through. In principle, the gate-source capacitance C_{gs} also causes charge to be injected at the source of transistors. However, since the sources of M3 and M4 are connected to the same node, the charge injected through one capacitor will be ejected through the other, thus canceling out this effect. The double-balanced mixer configuration does not suffer from LO feed-through due to the fact that each output node is connected to two transistors, each driven by opposite phases of the LO.

Preliminary schematic-level simulations with this mixing DAC architecture showed that the maximum static phase error of the phase modulator was approximately 1.5°, which is significantly higher than expected phase error discussed in Section 4.3. It was found that LO feed-through due to charge injection was the dominant source of this error. To address this issue dummy switch transistors M5 and M6 were employed, in the way illustrated in Figure 5.6. These transistors were made the same size as M3 and M4 so that C_{gd} was equal for all four devices. The gates of M5 and M6 were connected to the gates of M3 and M4, respectively, but the drains were cross-connected so as to cancel out the charge-injection. The sources of M5 and M6 were only connected to each other to not have any further effect on the circuit. With this addition, the maximum static phase error of the modulator was reduced to approximately 0.4° degrees, which is only slightly higher than in the ideal case presented in Section 4.3. Despite effectively doubling the load of the XOR gates, it still proved to be sufficiently small that the devices comprising the XOR gates could be left minimum size while maintaining a high switching speed.

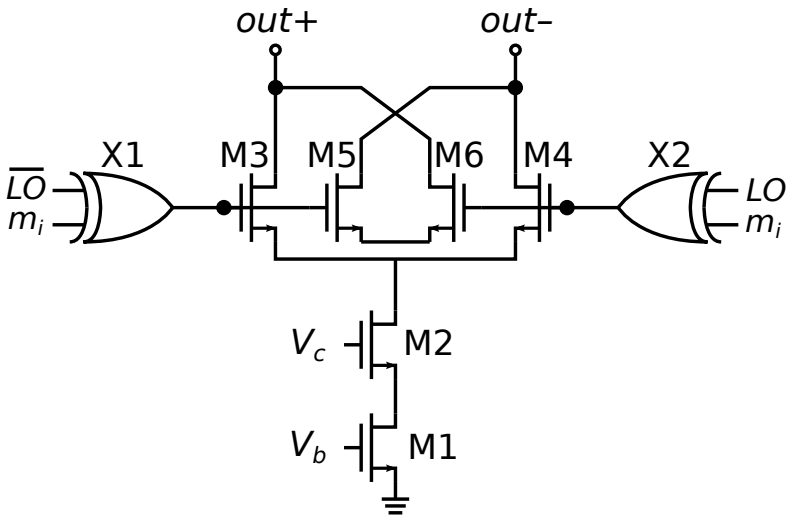

Figure 5.6: Mixing DAC unit cell with charge-injection cancellation

5.2.3 Unary/Binary Segmentation

An important factor in DAC performance is the choice between unary and binary weighted cells. In a binary weighted DAC design, each input bit is directly used to switch on or off a cell which is sized proportionally to the weight of that bit. In other words, the input bit m_i will control a cell whose output current (in the case of a current-steering DAC) will be a factor of 2^i times larger than the LSB current I_0. Figure 5.2a in Section 5.2.1 illustrates this approach. In the unary weighted, or thermometer coded, approach each cell has the size of the LSB. The n-bit binary input code is converted to thermometer code of $2^n - 1$ bits – one for each cell. If the input code is increased by one, then one additional unit cell will be turned on.

The advantages and disadvantages of both approaches were investigated in [53]. While the binary weighted approach is conceptually straightforward, any mismatch between the current sources used in cells can cause severe degradation of the differential non-linearity (DNL) performance. This is most around at the mid-code transition where the most significant cell is switched on and all other cells are switched off. Due to the statistical spread of the current sources, the step size associated with this transition can vary dramatically which harms the performance. To limit the DNL associated with this transition to within 0.5 LSB (a common specification to guarantee monotonicity) the current of the most significant cell must be within 0.5 LSB of the sum of all other current sources. In a given technology, the main way to control the mismatch is by changing the area A of the MOS devices. A first-order approximation of the relationship between A and the standard deviation σ of the current source is given as follows [54]:

$$\sigma^2 \propto \frac{1}{A} \tag{5.3}$$

Thus, a DAC with very precise matching requirements will occupy a large area.

The unary weighted approach does not suffer from this problem because at any point, if the input code is increased or decreased by one, only a single cell will be turned on or off, respectively. This means that the matching requirements of the current sources are significantly relaxed when compared to the binary weighted case. However, this comes at the cost of additional digital logic

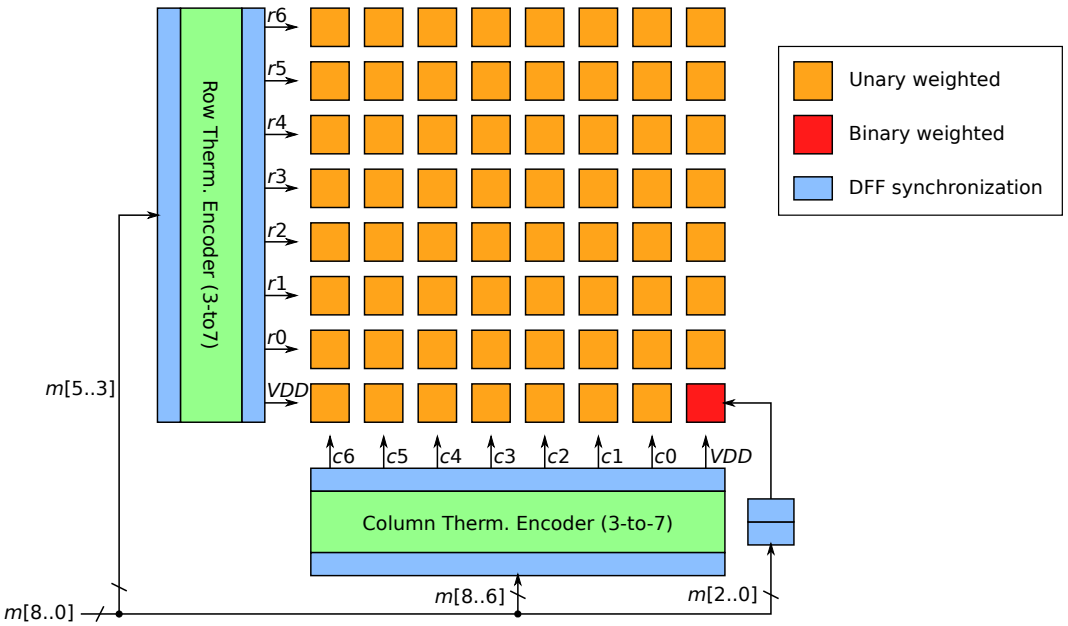

Figure 5.7: Segmented DAC structure with 6 bits unary and 3 bits binary

which in turn results in a higher power consumption and layout complexity. Furthermore, in some cases, the additional area required for the digital logic can be greater than the additional area to ensure good matching in the binary weighted case. In modern deep submicron CMOS technologies, however, the latter is not often the case.

The authors of [53] conclude that by segmenting the design into two sections where the MSB part is unary weighted and the LSB part is binary weighted, the advantages of both approaches can be leveraged, while keeping the disadvantages to a minimum. For this design, the 9-bit DAC was segmented such that the 6 MSBs were unary weighted and the 3 LSBs were binary weighted. This was considered a reasonable compromise between power consumption and layout complexity on the one hand and area on the other. For the unary weighted part this means that a total of 63 cell is needed, each with an output current equal to $8I_0$. For the remainder of this thesis work, these cells are referred to as unary cells, to distinguish them from the unit cell discussed above.

To make the layout of a unary weighted design more manageable, a common approach is to lay out the cells in a grid and drive the cells in rows and columns, rather than each cell individually [55]. Figure 5.7 shows the way in which this was applied to this design. As can be seen in the figure, the 9-bit input signal m is split into three parts. As explained above, $m[2..0]$ are used to control the binary weighted cells. For the unary weighted part, $m[5..3]$ and $m[8..6]$ are converted to the row and column select signals, respectively, using two 3-to-7 bit thermometer encoders. These signals are then used within the unary cells to decide whether a cell should be on or off. D flip-flops (DFFs) are used before and after the thermometer encoders to prevent glitches in the encoders themselves as well as any preceding digital logic. Two DFF stages were also inserted in the binary path to ensure that it has the same delay. The digital logic required for this design will be discussed in Section 5.2.6.

5.2.4 Unit Cell Sizing

Both the unit cell transistor sizes and the LSB current I_0 are important parameters for any DAC. In this design these parameters were selected using an iterative process to arrive at the final values. The final iteration could only be done after the layout was largely complete, since the parasitic capacitors at the DAC output nodes had to be known to correctly size several parameters.

To make an informed decision, it is important to understand the relationships between the various design parameters. First of all, considering that the HR-RF-DAC has a total of six DACs, two of which are scaled by a factor $\sqrt{2}$, the relationship between I_0 and the total current I_{tot} that is drawn from the supply is given by

$$I_{tot} = (4 + 2\sqrt{2}) \cdot (2^n - 1) \cdot I_0 \tag{5.4}$$

where n is the resolution of each DAC (9 bits for this design). Furthermore, the combination of I_{tot} and the load resistors R (see Figure 5.1) sets the output common mode voltage $V_{o,cm}$. This voltage should be in the middle of the available voltage headroom to achieve the most linear performance. The relationship between these parameters is given by

$$V_{o,cm} = V_{DD} - \frac{1}{2} R I_{tot} \tag{5.5}$$

The factor $\frac{1}{2}$ in this equation stems from the fact that from the point of view of the common mode voltage, the resistors in the differential RC filter are connected in parallel. The supply voltage V_{DD} for the 40 nm TSMC technology is 1.1 V. The resistor value R, however, cannot be set arbitrarily to get the desired output common mode voltage, since it also affects the cut-off frequency f_c of the filter:

$$f_c = \frac{1}{2\pi RC} \tag{5.6}$$

Although expressions for hand calculations can also be derived for the relationships between various transistor parameters using the traditional square-law model, these expressions tend to be less accurate for nano-scale CMOS technologies. For this reason, it was chosen to set the transistor parameters based on simulation results instead.

Initially, it was assumed that the total current I_{tot} could be set to approximately 1 mA. Using Equation 5.4, this resulted in a LSB current of $I_0 = 287\,\text{nA}$. These values were assumed for the majority of the design process. Based on these values, the initial sizing was made for the current source and cascode transistors in the unit cell. The most important parameter for a current source is its output impedance, but with short-channel devices, this is difficult to achieve. While the width of both devices was kept at $W = 120\,\text{nm}$, the length was increased from minimum size (40 nm) to $L = 400\,\text{nm}$. This gave a DC output impedance of more than 20 MΩ if the drain voltage of the cascode transistor was greater than 250 mV. Using Equation 5.5, the required resistor to set the output common mode to $V_{o,cm} = 675\,\text{mV}$ (the midpoint between 250 mV and V_{DD}) was found to be $R = 850\,\Omega$. Following Equation 5.6, for a filter cut-off frequency of $f_c = 2.4\,\text{GHz}$, this corresponds to a capacitor value of $C = 78\,\text{fF}$.

As was discussed in Section 5.2.3, the mixing DAC was segmented into a 6-bit unary weighted part and a 3-bit binary weighted part. This means that in order to have DNL < 0.5 LSB, the mismatch of the unary cells (which have an output current of $8I_0$) had to be less than ± 0.5 LSB as well. The area requirement for this condition was computed through Monte-Carlo simulations of the current sources in which the W/L ratio was kept constant while the area was changed. It

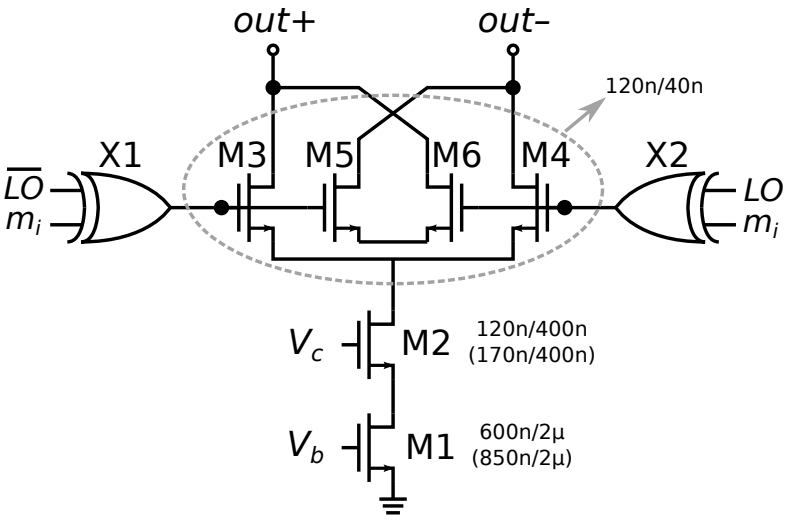

Figure 5.8: Mixing DAC unit cell showing transistor W/L values (values in parentheses are for I-DAC 2 and Q-DAC 2 which have a gain scaled by $\sqrt{2}$)

was concluded that to satisfy the 3σ DNL requirement, the area of the current source transistor M1 had to be increased to $W/L = 600\,\text{nm}/2\,\mu\text{m}$. Since the matching of the cascode transistor M2 had a much smaller effect on the current, its area could remain unchanged. This was beneficial, as making M2 bigger would also result in a higher parasitic capacitance at the output. The W/L values chosen for these two transistors are reflected in Figure 5.8.

As mentioned above, the output currents of I-DAC 2 and Q-DAC 2 had to be scaled by a factor of $\sqrt{2}$ for the harmonic rejection operation to work. This was achieved by scaling the widths of the M1 and M2, proportionally, resulting in widths of 850 nm and 170 nm, respectively. In Figure 5.8, the sizes of the scaled DACs are shown in parentheses.

With I_0 set to such a low value, the switches M3-M6 could safely be made minimum size. The on-resistance of the switches caused a negligible voltage drop, even for higher currents. In an effort to keep the parasitic capacitance to a minimum, it was decided that each cell, be it unary or binary, would have a single set of switches despite the fact that these cells would have up to eight unit current sources connected in parallel.

While these values for the various parameters showed decent results at the schematic level, some difficulties arose when the layout was nearing completion. The layout of the mixing DAC and the HR-RF-DAC will be covered in Section 5.2.7 and Section 5.5, respectively. For this discussion, however, it is important to know that both output nodes of the HR-RF-DAC (to which all six mixing DACs are connected) had a parasitic capacitance of approximately 400 fF. This capacitance value was significantly higher than what was initially expected and it essentially set the minimum value of the RC filter capacitance. Using Equations 5.4 - 5.6, it can be found that for $f_c = 2.4\,\text{GHz}$ the other parameters must be $R = 165\,\Omega$, $I_{tot} = 5.12\,\text{mA}$ and $I_0 = 1.47\,\mu\text{A}$.

Unfortunately, it was found that with the W/L values of M1 and M2 described above, the current sources could not provide the desired I_0. The highest I_0 that could be achieved while keeping the bias voltages and output impedance at reasonable levels was found to be approximately 80 µA. Furthermore, for the unary cells, where eight such current sources were connected in parallel, the assumption that the voltage drop over the switches was negligible no longer held

true. By that point, however, it was too late to change the transistor sizes since that would require modifying the entire layout of the DAC. Although it was possible to set R to the desired value, with the limited current that the current sources could provide, the common mode voltage would be significantly higher. More importantly, the output voltage swing would also be significantly reduced.

In the end, a compromise was struck by setting the load resistor R to 300 Ω and I_0 to 800 nA. This unfortunately lowered the filter cut-off frequency to $f_c = 1.3\,\text{GHz}$, which did somewhat attenuate the signal. However, due to the gentle roll-off of a 1st-order filter, the differential signal swing at the output was still approximately $300\,\text{mV}_{\text{pk-pk}}$, which was higher than it would have been if the resistor had been left at $R = 165\,\Omega$. This swing proved to be high enough that the limiting amplifier could successfully amplify it and still achieve adequate performance.

5.2.5 Bias Circuit

To bias the current sources in the unit cell, a standard bias circuit for cascoded current sources was used. The structure of this circuit is depicted in Figure 5.9. This circuit consists of two independent branches where the right branch generates the bias voltage for the main current source transistors (M1) and the left branch biases the cascode transistors (M2). By keeping the transistor sizes in the bias circuit the same as those used in the unit current sources, the bias current I_0 flowing through the right branch gets mirrored to the unit current sources. In this design, as explained above, this current was set to $I_0 = 800\,\text{nA}$.

The bias voltage for the cascode transistors is generated in the left branch, where a transistor with the same size as the cascode transistor is driven by a current which is a factor c higher than I_0. Following the square-law model, this factor should be set to $c = 4$, assuming both the main current source transistor and the cascode have the same W/L ratio. This sets the gate voltage of the cascode transistors to a voltage that is one saturation voltage $V_{ds,sat}$ higher than what would be required if the source of the transistor would be connected directly to GND. In practice, the factor c usually has to be tuned to get the best performance. For this design, since the chosen bias current I_0 was close to the limit of what a transistor with this W/L ratio could handle, c had to be set to 10 to achieve accurate current mirroring. This had the effect of increasing the drain-source voltage V_{dc} of the main current source transistor, which helped it reach this higher current.

The bias currents used in this circuit were derived from a 100 µA off-chip current reference through current division techniques using current mirrors. The choice for an off-chip reference was made for two reasons. First, this would save design time since it eliminates the need to design an accurate on-chip current reference. Second, it allows to fine-tune the current as required to always get the best performance, in spite of PVT variations of the various system components, such as the load resistors.

5.2.6 Digital Logic

As was discussed briefly in Section 5.2.3, some digital logic is required to achieve the desired DAC operation. The majority of this logic is needed for the unary weighted part of the DAC. As shown in Figure 5.7, this includes two thermometer encoders with DFF-based registers at both the input and output. These encoders convert a 3-bit binary signal into a 7-bit thermometer code. The value of the binary signal determines how many of the output bits are turned on. These thermometer codes are the row and column select signals. A row/column decoder contained within each unary cell then uses these signals to determine whether that cell needs to be turned

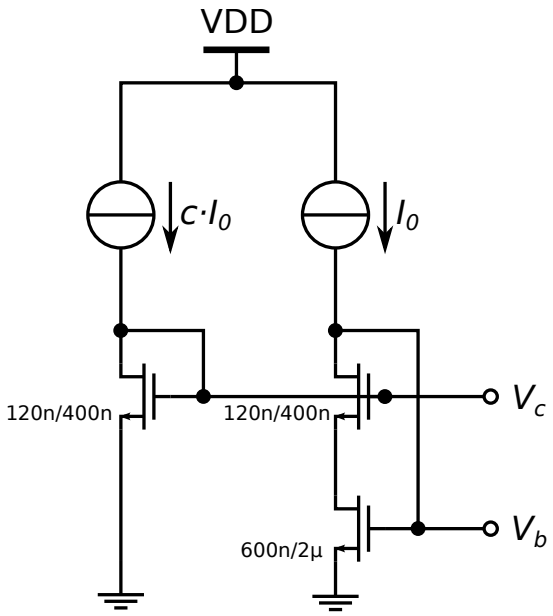

Figure 5.9: The bias circuit for the DAC

on or off. Additionally, this logic block needs to convert the single-ended digital signal to a complementary one in order to drive the two XOR gates contained within each cell.

The binary weighted part of the DAC requires much less digital logic, since each bit of the binary input data maps directly to a single binary weighted cell. Nevertheless, since the propagation delay of the unary and binary data paths must be approximately equal, the binary path also needs two stages of DFFs. Furthermore, the requirement that the data must be supplied to the cells as a complementary signal still applies.

The majority of the digital circuits presented in this section are heavily based on the digital logic used in [49] for the implementation of an RF-DAC. These circuits have proved to be fast enough for operation at 2.4 GHz in 65 nm CMOS. Since the performance of digital circuits is know to improve with technology scaling, this was considered a safe option to achieve high speed and low power consumption, while simultaneously reducing the design time.

An overview of the digital logic used to generate the row and column select signals is shown in Figure 5.10. At the core of this logic is a 3-to-7 thermometer encoder, with both the inputs and outputs of the encoder synchronized to a clock signal using DFFs[1]. The clock signal is derived from the LO but unlike the LO itself, this clock signal is the same for all six DACs. This ensures that the I and Q data gets updated simultaneously, rather than gradually over the course of an LO period. For the row and column encoders, the inputs $A[2..0]$ correspond to $m[5..3]$ and $m[8..6]$, respectively. The outputs $B[6..0]$ are fed to the cells in their respective row or column within the array of unary cells. During the layout phase of the design, it was found that, in this configuration, the interconnect parasitics presented a load capacitance of approximately 36 fF

[1]The thermometer encoder and the D flip-flop presented in this section were designed and laid out by Yiyu Shen and Lei Zhou, respectively, as part of their work on the digital PA for this transmitter system. Since the requirements of the digital logic were very similar, they were reused in this phase modulator design. The schematics of these two blocks are discussed in this section for the sake of completeness.

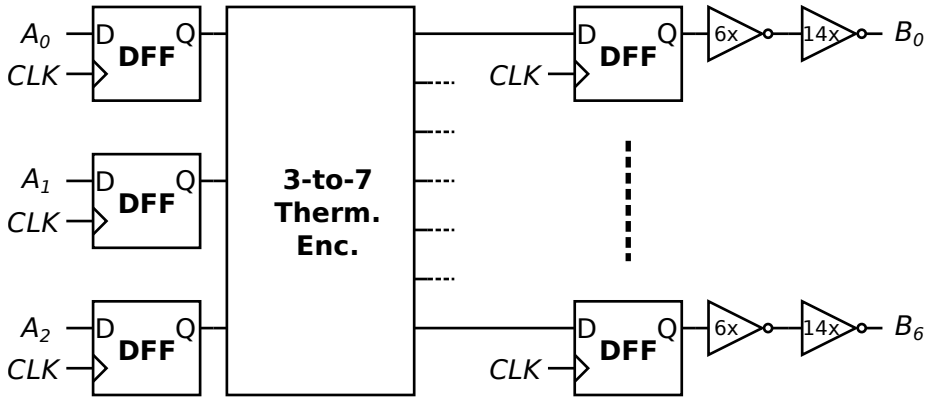

Figure 5.10: Overview of the digital logic used to generate the row and column select signals

at the output of this digital block. Since this capacitance was too high to drive directly with a DFF, a buffer consisting of two inverter stages was inserted at the output of the second DFF stage. The size of the inverters was tuned so that the rise and fall times at the output was approximately 50 ps. As shown in the figure, appropriate sizes for the first and second inverter were found to be 6× and 14× minimum size[2], respectively.

The 3-to-7 thermometer encoder in this design consists of a 2-to-3 thermometer encoder, which computes intermediate signals based on the 2 LSBs, and some additional digital gates which use these intermediate signals and the MSB to compute the 7-bit output. The schematics of the 2-to-3 encoder and the 3-to-7 encoder are shown in Figure 5.11a and Figure 5.11b, respectively. As can be seen from Figure 5.11a, the 2-to-3 thermometer encoder consists of only three gates: an OR gate, a buffer and an AND gate. These gates operate on A_0 and A_1 and are used to produce the intermediate signals BB_0, BB_1 and BB_2, respectively. In the 3-to-7 bit thermometer encoder the 3 LSBs (B[2..0])are computed by ORing these intermediate signals with A_2 while the 3 MSBs (B[6..4]) are computed by ANDing them with A_2. Since the middle bit (B_3) is equal to A_2, a simple buffer is used. With all the gates set to minimum size, the 3-to-7 thermometer encoder was still fast enough to complete the computation well within one period of the clock.

The DFF that is used to synchronize the inputs and outputs of the thermometer encoders consists of two latches in the traditional master-slave configuration as shown in Figure 5.12a. The latches are based on a multiplexer consisting of two pass gates driven by the clock signal. The schematic of the latch is shown in Figure 5.12b. When the clock is low, the four outer transistors essentially form two cascaded inverters with the D signal at the input. When the clock goes high, the connection to D is severed and the inverters are connected back to back, thereby maintaining the previous value of the signal. Note that since the latch uses pass gates, the clock signal needs to be made differential. Since only a single-ended clock signal is distributed throughout the HR-RF-DAC, the negative clock signal is generated within the DFF using an inverter.

The schematic of the row/column decoder logic that included in each unary is shown in

[2]For an NMOS device, minimum size implies $W/L = 120\,\text{nm}/40\,\text{nm}$. Since PMOS devices have a lower mobility (μ) then NMOS devices, they need to be made wider for the same speed. In this design, the PMOS devices used for digital logic were scaled by a factor of 2.25 in an effort to achieve approximately equal rise and fall times. Thus, for a PMOS device, minimum size implies $W/L = 270\,\text{nm}/40\,\text{nm}$.

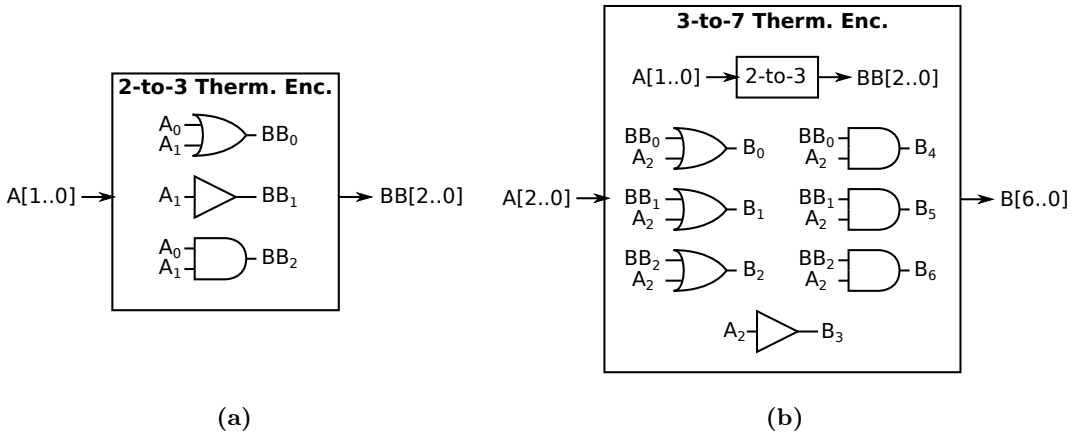

(a) **(b)**

Figure 5.11: Schematics of **(a)** the 2-to-3 thermometer encoder and **(b)** the 3-to-7 thermometer encoder

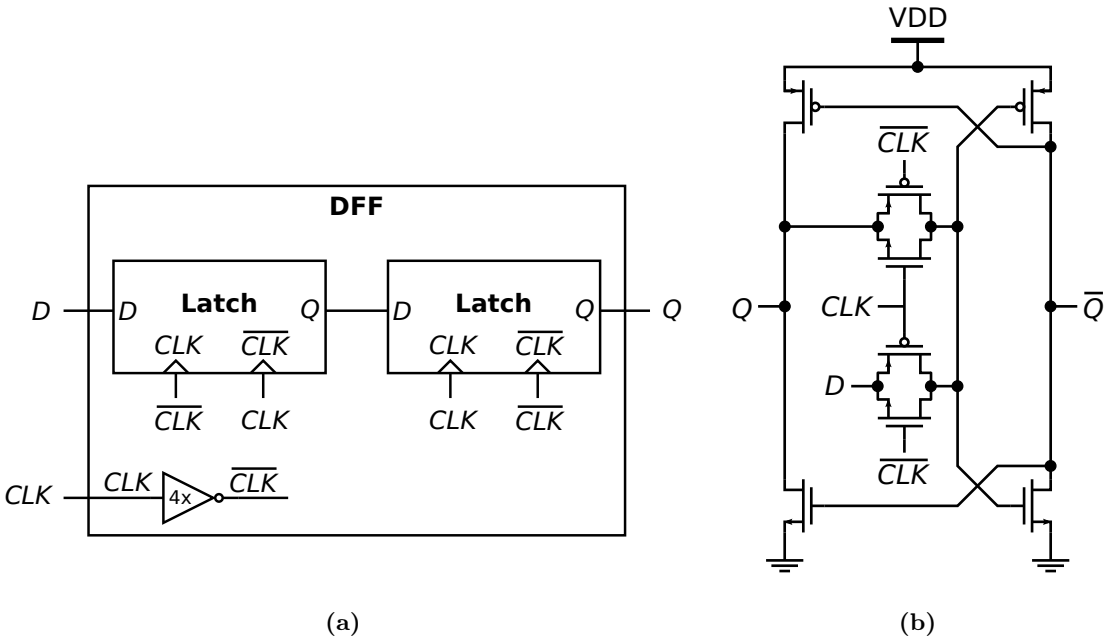

(a) **(b)**

Figure 5.12: Schematics of **(a)** the D flip-flop and **(b)** the latch that is used in it

Figure 5.13. A custom digital gate is used to implement an AND-NOR operation, which serves to determine if the unary cell in question has to be on or off based on the select signals of its row and column as well as the following column select signal. The devices used in this gate were minimum size.

As mentioned above, this logic block must also convert the unary cell control signal from single-ended to complementary. This is accomplished by means of the inverter configuration shown in the figure. Using a single inverter to generate the negative output from the positive output would result in one of these signal lagging the other. This, in turn, would result in even-order harmonic distortion and intermodulation products at the output of the DAC. By

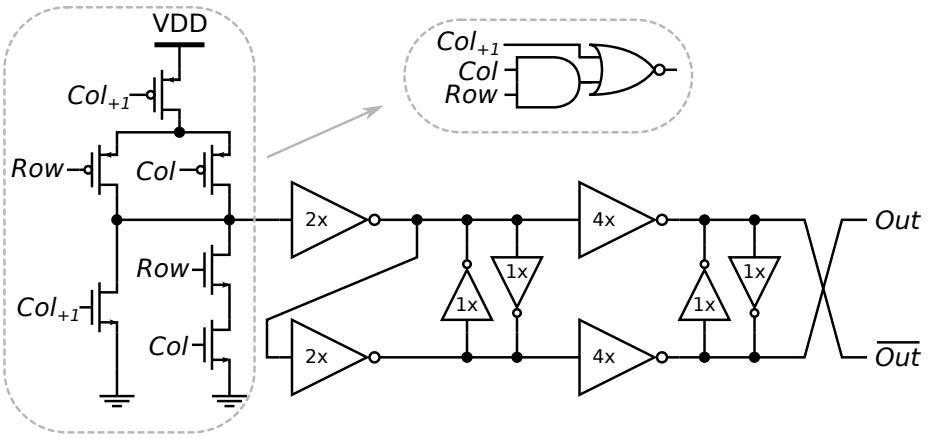

Figure 5.13: Schematic of the row/column select logic within each unary cell

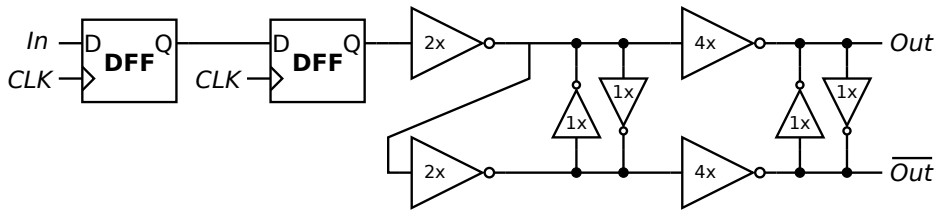

Figure 5.14: Schematic of the digital logic for the binary cells

employing phase aligners, consisting of two weak inverters connected back to back, the timing difference between these signals can be reduced. With two stages of inverters with phase aligners the difference was reduced to less than 10 ps in a post-layout simulation.

The desired operation for the row/column decoder is AND-OR, but since CMOS logic is inverting, an AND-NOR gate was implemented instead. To correct for this, the polarity of the complementary outputs was reversed, as shown in the figure.

As explained above, the binary weighted cells also required some digital logic, both to match the delay in the unary path and to convert the input signal from single-ended to complementary. The schematic of this logic is shown in Figure 5.14. As the figure shows, it consists of two cascaded DFFs followed by the single-ended to complementary converter used for the unary path. Since signal coming from the DFF was not inverted, switching the polarity of the output complementary output signal was not necessary in this case.

5.2.7 Mixing DAC Layout

Layout is a critical part of any RF design. Due to the high-frequency nature of RF circuitry, parasitic coupling between devices, metal interconnect lines and the substrate will have a much greater effect on the performance of the system than for circuits operating at lower frequencies. In general, the layout should be made as compact as possible, since this not only saves on cost but also reduces the parasitics associated with long interconnects.

Layout is also important for data converters, such as DACs, since a key factor that will

determine the performance of a data converter is the mismatch between the cells. By following good design practices and rules of thumb, the mismatch can be kept to a minimum, benefiting the differential non-linearity (DNL) and integral non-linearity (INL) performance. Perhaps the most important design practice to follow, regarding mismatch, is to ensure that key parts of the cells (in this case, the current sources) all see the same environment. This can be achieved by placing the cells in a regular grid, with dummy cells around the edges. In this way, the mismatch between devices on the manufactured chip should be close to the mismatch that can be modeled through schematic-level Monte-Carlo simulations.

Unfortunately, small area and short interconnect lines on the one hand, and a regular grid of current sources surrounded by dummies on the other are somewhat conflicting requirements. For instance, dividing the control signals of the unary weighted part of the DAC into rows and columns has little to no benefit in terms of parasitics and layout complexity if the switches and the row/column select logic blocks are located far away from the current sources. On the other hand, inserting those circuits into the grid makes it less regular, thus contributing to the mismatch. As a compromise, the DAC layout for this design consists of alternating rows of current sources and digital logic. This helps keep the layout compact and the interconnect length relatively short, and since the digital logic is identical for all unary cells, the environment seen by each cell is still roughly the same, which, in turn, improves the matching.

The DAC layout is described here in a bottom-up fashion, starting with the unit current sources and working up to the full mixing DAC design. As discussed in Section 5.2.4, two different unit current sources are used in this design. Both current sources are composed of two transistors, M1 and M2, with the same lengths but with different widths. For I-DAC 1, I-DAC 3, Q-DAC 1 and Q-DAC 3, the size of M1 is $W/L = 600\,\text{nm}/2\,\mu\text{m}$ and the size of M2 is $W/L = 120\,\text{nm}/400\,\text{nm}$. These current sources are termed "narrow" current sources, and are referred to as such throughout the rest of this thesis work. For I-DAC 2 and Q-DAC 2, the widths of M1 and M2 were scaled by a factor of $\sqrt{2}$ leading to the sizes $W/L = 850\,\text{nm}/2\,\mu\text{m}$ and $W/L = 170\,\text{nm}/400\,\text{nm}$, respectively. These are referred to as "wide" current sources.

The layout of the narrow and wide current sources can be seen in Figure 5.15a and Figure 5.15b, respectively. In both figures, M1 is the large transistor oriented vertically and M2 is the smaller transistor oriented horizontally. The only difference between the two is width of the active area of the two transistors (shown in red). A p-strap guard ring connected to GND surrounds the current sources, with regular vias connecting the substrate to metal layer 1 or m1 (light blue). This serves several purposes. First and foremost, this guard ring provides a low-impedance path from the substrate around the devices to GND. Second, it isolates the current source from its surroundings, making sure that no GND current from the surrounding circuitry can pass under the devices. If this isolation is not provided, stray current flowing through the substrate can affect the bias point of the device, which in turn can lead to an error at the output. Finally, if multiple current sources are laid out side by side such that their guard rings touch, these rings form a fine mesh spanning the entire DAC, which provides a convenient way to ensure a low-impedance path to ground for all of the connected transistors.

The current sources are designed in such a way that they can be stacked together horizontally to form continuous rows. Horizontal bias lines routed on m2 (yellow) provide the bias voltages to V_b and V_c to M1 and M2, respectively. A solid sheet of m3 (green) covers the entire current source and is connected with vias to the GND on m1. This serves to shield both the bias lines and the devices themselves from potential high-frequency signals that can be routed over the current sources on higher metal layers.

As mentioned above, the only difference between the narrow and wide current sources is the

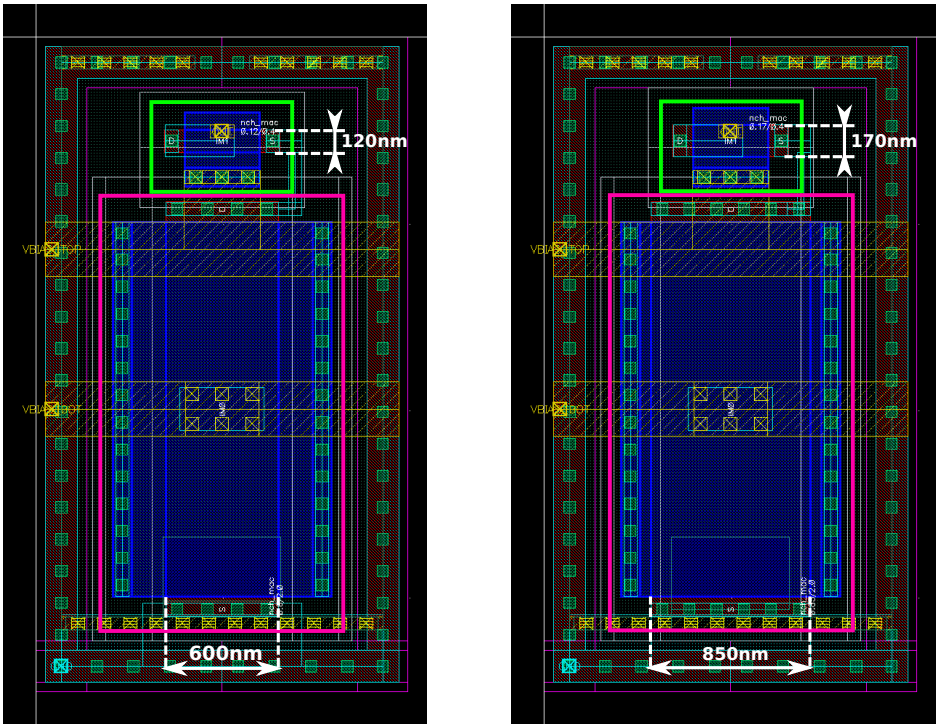

(a) **(b)**

Figure 5.15: Layout of **(a)** the narrow current source (used in I-DAC 1, I-DAC 3, Q-DAC 1 and Q-DAC 3) and **(b)** the wide current source (used in I-DAC 2 and Q-DAC 2), with M1 marked in magenta and M2 marked in green

width of the active area of the transistors. The guard ring, however, is the same size in both cases which means that both current sources occupy the same area and have the same external connections. This means that apart from which current sources are used, all six DACs in this design can have an identical layout. As such, the remainder of this section covers only the DAC layout with narrow current sources.

The layout of the unary cell is shown in Figure 5.16. Its eight current sources are juxtaposed side by side, with the switch transistors located above them and the XOR gates and the row/column select logic located below. By laying these components out in this way, the differential output lines can be kept far away from high speed digital signals, such as the LO, such that the capacitive coupling between them is minimized.

Much like the current sources, the unary cells are designed to be connected in rows. Both the output and the LO span the length of the cell so that placing multiple cell side by side forms a continuous connection. In an effort to keep the parasitic capacitance at the output to a minimum, the output lines are routed on m5, which for a minimum track width has the lowest parasitic capacitance to the substrate per unit length. Although the same is also preferable for the LO lines, m5 was occupied by some of the connections used in the XOR gates. Therefor m6 was used instead.

The digital logic was laid out in such a way that its total width was exactly half the width of the eight current sources. In this way, the remaining area could be used by the digital logic of the cell below it. Similarly, the area just to the left of the switch transistors could be used

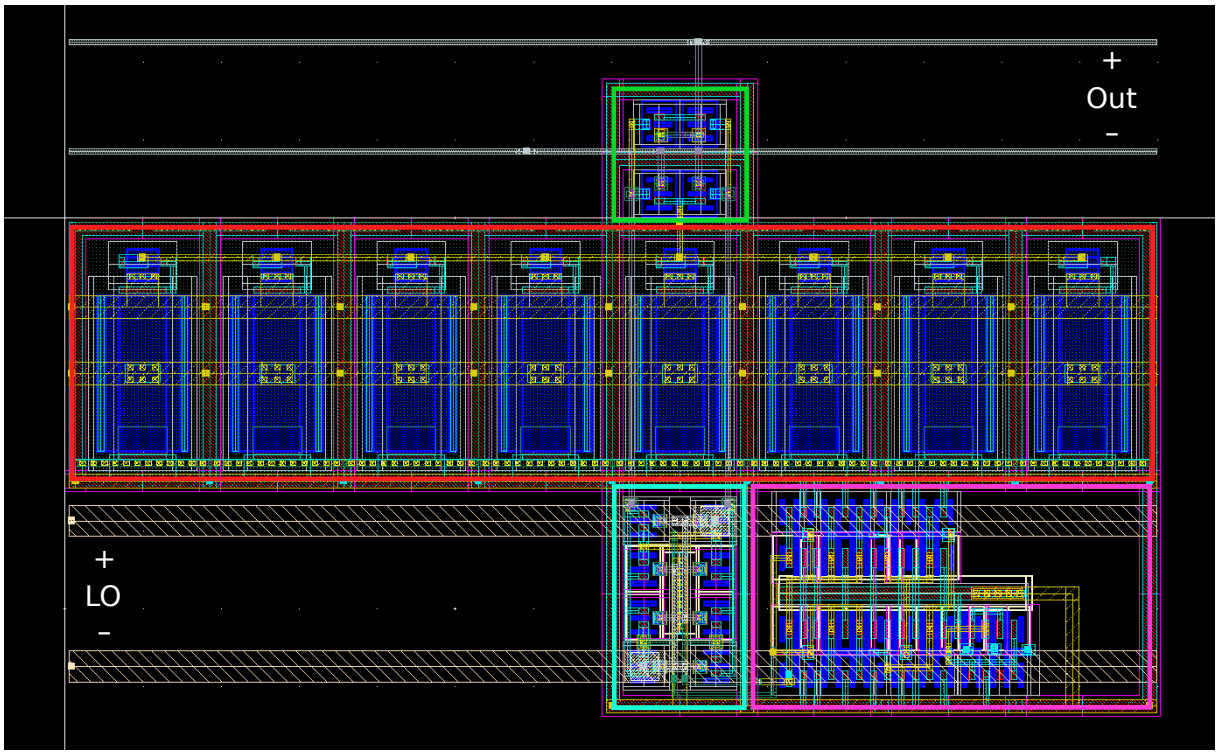

Figure 5.16: Layout of the (top) unary cell showing eight current sources (red), switch transistors (green), two XOR gates (cyan) and the row/column select logic (magenta)

for the switches of the cell above. This allows both for a very compact layout, and it reduces the parasitic capacitances of both the LO and the output lines by half, due to these lines being shared by two neighboring rows. To achieve this alternating rows must be rotated by 180° and the positive and negative connections to the output and LO must be reversed. To distinguish between the two, the unary cell depicted in Figure 5.16 is termed the "top cell", whereas the rotated version is termed the "bottom cell". Although not shown in the figure, the row select signals for the top and bottom cells are routed horizontally on m4 in between the LO lines, whereas the column select signals run vertically round the cell edges on m5.

As discussed in Section 5.2.3, the three LSBs of this DAC design are binary weighted. Together, the three binary cells contain seven unit current sources. These cells are laid out as a single block, similar to a single unary cell, which completes the cell grid of the DAC (see Figure 5.7). This block, laid out as a "bottom cell", is shown in Figure 5.17. In the figure, the current sources, switches and XOR gates of the three different cells are marked individually. An additional dummy current source is added to mirror the eight-current-source unary cell structure.

Unfortunately, due to the particular arrangement of this block, the digital logic (consisting in this case only of XOR gates) did not fit in the space above the left-most four current sources, as would be the case for a unary "bottom cell", without significantly increasing the length of the XOR output lines. For this reason, the XOR gates were placed below the output lines, as shown in the figure. This is expected to contribute somewhat to the mismatch between the current sources in that area of the DAC, but it was considered a necessary measure for RF operation.

The layout of the full mixing DAC is shown in Figure 5.18. The majority of the area is taken up by the cell array (marked in red), consisting of four pairs of rows (top and bottom), each

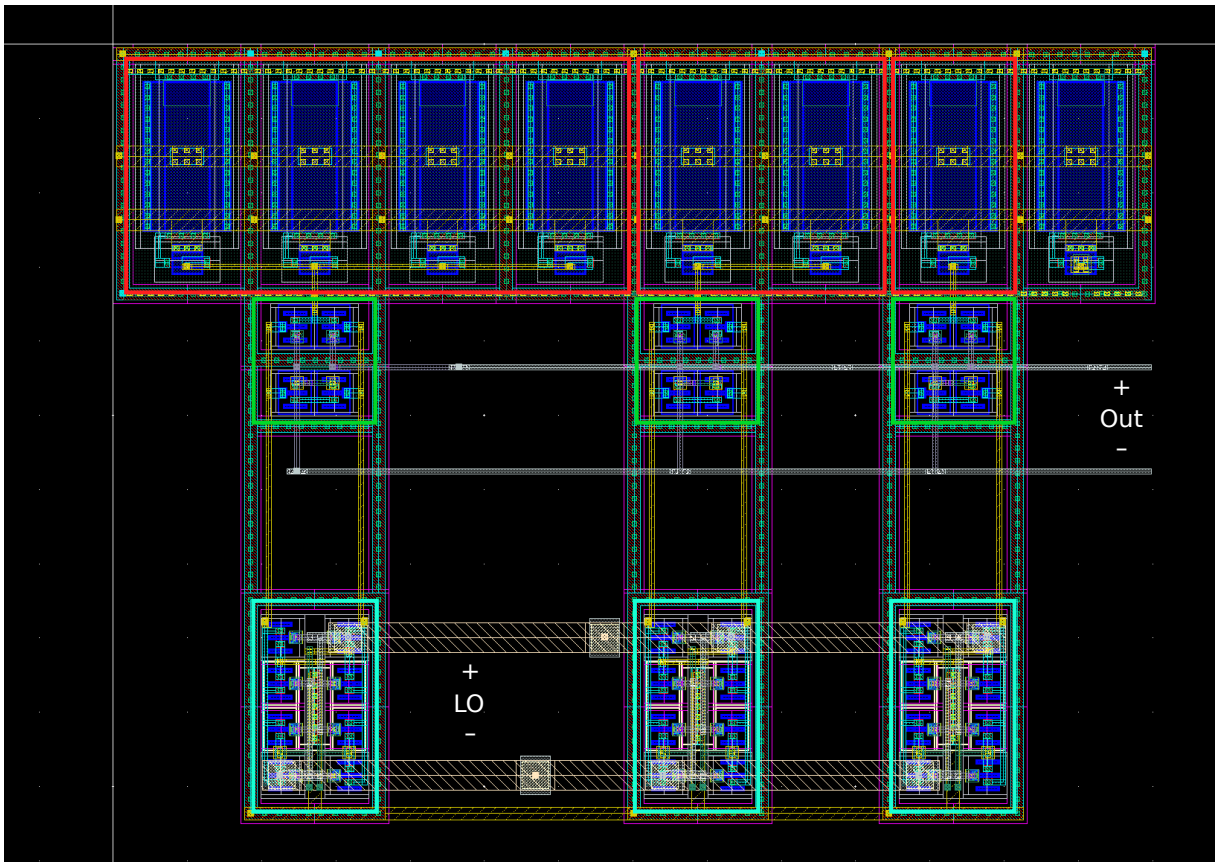

Figure 5.17: Layout of the block of binary cell showing the current sources (red), switch transistors (green) and XOR gates (cyan)

with eight columns. The outputs of all the rows are connected on the right to two vertical lines routed on m6. Dummy current sources were added to the left and right edges of the rows to improve the matching. To the left and bottom of the array are the column and row encoders, respectively (cyan), which drive the array of unary cells. The logic for the binary cells (magenta) is in the bottom-right corner.

Both the LO and the digital input data are supplied to the DAC from the left side, as indicated in the figure. Due to the long lines, these signals present a high capacitive load (approximately 30 fF for each digital input and LO row). To make it more manageable to drive these signals, buffers are inserted both for each row of the LO (green) and for each input bit and the clock signal (orange). Finally, all available area is filled up with decoupling capacitors between VDD and GND, to help suppress any supply noise coming from the rest of the system. Care was taken to leave at least 2 µm spacing between the decoupling capacitors and the critical high-frequency signals to ensure that these capacitors contribute a negligible amount to the parasitic capacitance of those lines.

5.3 Limiting Amplifier Design

The previous section described in detail the design of the mixing DAC, which, when loaded with a passive RC filter, produces an analog differential output voltage. The limiting amplifier (or

Figure 5.18: Layout of the full mixing DAC, showing the cell array (red), row and column encoders (cyan), binary logic (magenta), LO buffers (green) and data buffers (orange)

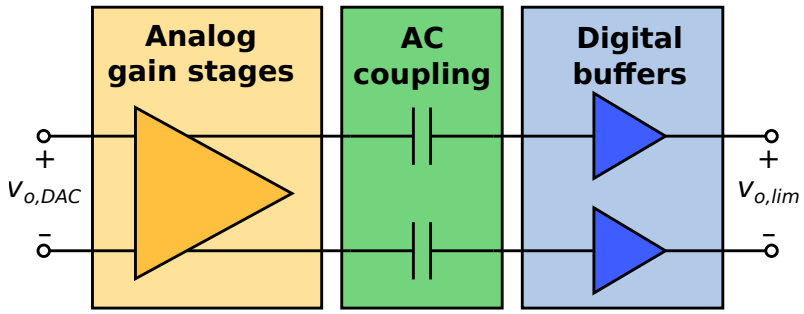

Figure 5.19: Top-level schematic of the limiting amplifier

limiter) must amplify this output voltage in order to convert the analog RF waveform into a square wave. This square wave can then be used to drive the digital PA. This section describes the design of this limiter[3] at the schematic level and in layout.

The top-level concept of the limiter is depicted in Figure 5.19. This limiter design is split into three distinct stages. At the input, an open-loop differential amplifier is employed to amplify the output voltage of the mixing DAC to the point of clipping. While this amplifier has a high gain, and its output does approximate a square wave, it is not yet at CMOS levels, since a typical amplifier cannot achieve rail-to-rail operation. For this reason, digital buffers are used at the output side. These not only provide further amplification for the signal, but also generate an output voltage that switches between VDD and GND.

It is important that the digital buffers output a square wave with 50 % duty cycle. Otherwise,

[3]While the initial design of the limiting amplifier was performed by me, both the sizing of the components and the layout described in this section are the work of Milad Mehrpoo. Without his help, I would not have been able to complete the phase modulator design within a reasonable time frame.

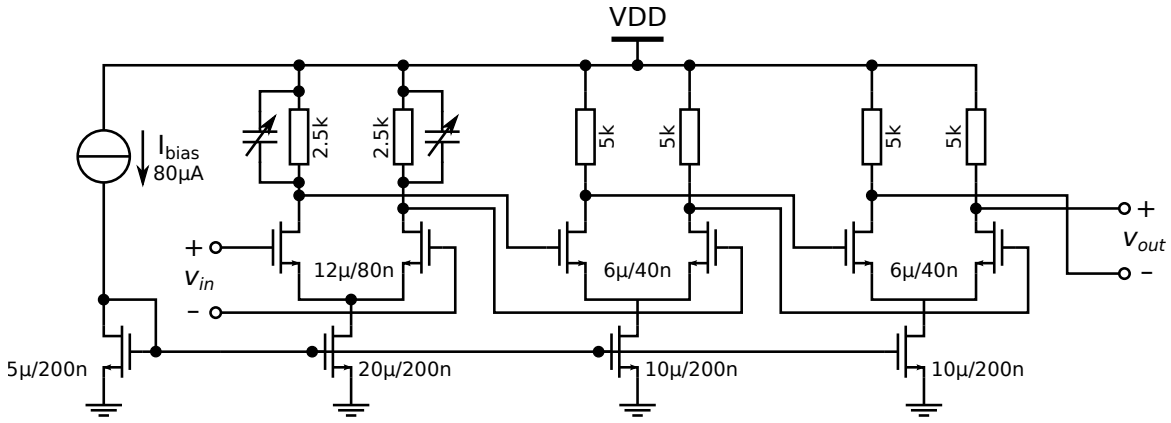

Figure 5.20: Analog part of the limiting amplifier, consisting of three gain stages

the output signal will contain significant even-order harmonic content. To achieve this, the input waveform should be centered around the threshold, or trip point, of the buffer, which is approximately half of the supply voltage. The output of the analog amplifier, however, is centered at a higher voltage, as will be discussed shortly. To solve this issue, AC coupling capacitors were inserted between the analog and digital stages. This allowed the input of the digital buffers to be biased at the desired voltage without affecting the operation of the analog amplifier.

5.3.1 Analog Gain Stages

The analog part of this limiter design consists of three cascaded common-source (CS) stages with a resistive load as shown in Figure 5.20. This open-loop structure is commonly used in optical communication systems, where limiting amplifiers are a much more common occurrence [56]. The first stage is designed to have a bandwidth of approximately 2.4 GHz in order to act as the second stage of the RC filter in this design. Programmable capacitor banks are connected in parallel with the load resistors of the first stage in order to lower the bandwidth of this stage when operating at lower frequencies. The capacitor banks have a 4-bit resolution and have a maximum capacitance of approximately 87 fF. The NMOS devices used in this design are RF MOSFETS. Unlike the baseline transistors used in the DAC, these are optimized for high-frequency analog operation at the cost of a larger overall area.

The sizing of the transistors was tuned to achieve the highest possible gain with a power budget of 1 mW. The reason to optimize for gain was that the gain directly affects the AM-to-PM distortion profile of the limiting amplifier [56]. In principle the use of a LUT to supply constant-envelope input data to the DACs is meant to relax the AM-to-PM distortion requirement, both for the HR-RF-DAC and the limiter. However, due to the fact that in this design the output of the LUT is interpolated, a constant-envelope input is no longer guaranteed. Overall, the three gain stages in this design achieve a small signal gain of approximately 30 dB at 2.4 GHz, which has shown to provide adequate system performance.

Unlike in the DAC, the current sources used to bias the differential pairs are not cascoded. This choice was primarily made due to the increased voltage overhead required for the cascodes. In this way, the output voltage of the gain stages can be made as high as possible before being fed to the digital buffers. Each stage has an output common-mode voltage of approximately

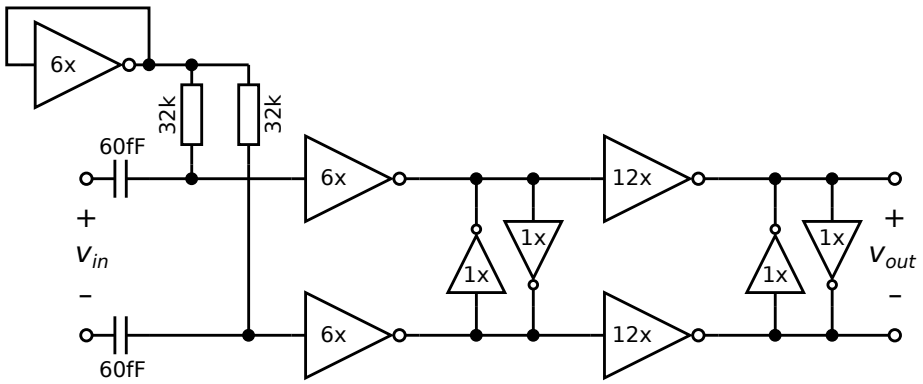

Figure 5.21: AC coupling and digital part of the limiting amplifier

0.6 V. The second and third stages reach clipping and their output swing is approximately 1 V, namely, from 100 mV up to the 1.1 V supply.

5.3.2 AC Coupling and Digital Buffers

The output of the analog gain stages is AC-coupled to the input of a set of digital buffers to provide further amplification of the signal and to convert the signal to CMOS logic levels. The schematic of the AC coupling and the buffers is depicted in Figure 5.21. The buffering operation is achieved using an inverter structure similar to the one described in Section 5.2.6. The series capacitors at the input of the inverters serve to de-couple the DC levels at the output of the gain stages and the inverter inputs. This was done to make sure that input signals that the inverters see is always centered around the inverters' theshold voltage. This voltage is generated by connecting the output of an inverter to its input, which forms a negative feedback loop. It is, then, applied as a bias to the input of the first stage of inverters in the signal path. Assuming the inverter generating the bias voltage has a negligible output impedance (a fair assumption, since the transistors in the inverter are diode-connected), the resistor-capacitor combination form a high-pass filter with a cut-off frequency of approximately 83 MHz, which is sufficiently low to not affect any of the desired high-frequency signal content.

The differential digital buffer consists of two stages of inverters with weak back-to-back inverters acting as phase aligners at the output of each stage. The output of the limiter must drive the phase path input of the DPA which is an inverter of 4 times minimum size. In addition to this, due to the physical on-chip distance between the limiter and the DPA, it was estimated that the interconnect parasitics would add an additional 10 fF of capacitance. With this load, the sizes of the inverters were tuned to provide rise and fall times of approximately 70 ps in a post-layout simulation.

5.3.3 Limiter Layout

The layout of the limiting amplifier is shown in Figure 5.22. The input and output of the amplifier are at the bottom and top respectively. The layout is made as symmetrical as possible in order to improve the matching in both the differential pairs and the capacitors. For instance, of the current sources shown in the figure (in green), only the left part is used to bias the differential

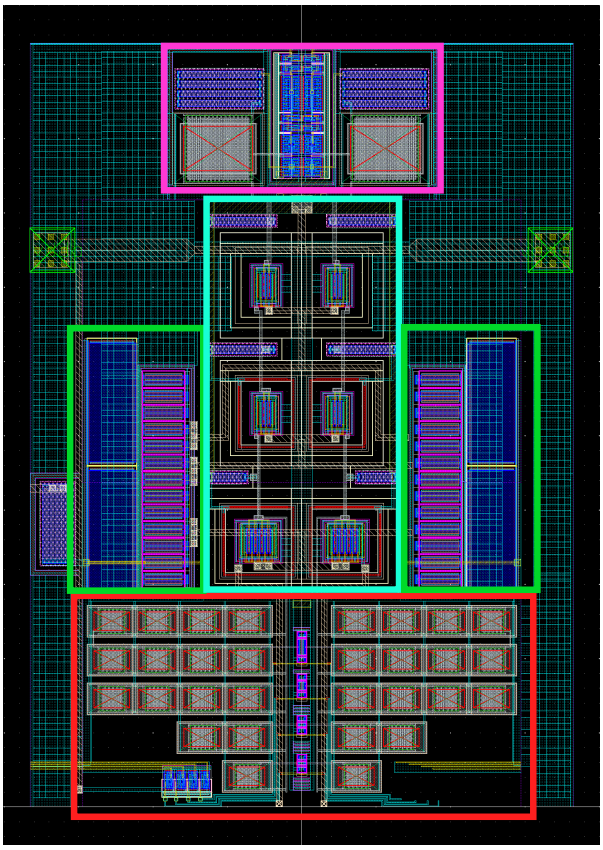

Figure 5.22: Layout of the limiting amplifier showing the programmable capacitor banks (red), the differential pairs and the load resistors of the gain stages (cyan), their current sources (green) and the AC coupling and digital buffers (magenta)

pairs in the gain stages. The right part, on the other hand, is not connected and simply serves as a dummy to make the layout symmetrical.

In order to provide a low-impedance path to GND for all transistors, all otherwise unfilled area was covered by a fine ground mesh using m1. The VDD and the RF signals were routed on m6, which is thick enough to provide a low impedance, but thin enough to minimize the cross-coupling between various signals.

5.4 Filter Design

As explained in Section 5.1, the HR-RF-DAC is loaded with a 1^{st}-order differential RC filter. While the HR-RF-DAC is designed to suppress the 3^{rd} and 5^{th} harmonics, this filter, together with the first stage of the limiting amplifier, is meant to attenuate the 7^{th} and higher harmonics. Additionally, it converts the current output of the HR-RF-DAC to a differential voltage which drives the limiting amplifier.

The basic structure of this filter has already been shown in Figure 5.1. It consists of a differential parallel combination of a resistor and a capacitor with values R and C, respectively. A more detailed schematic of this filter is shown in Figure 5.23. As discussed in Section 5.2.4, without any explicit capacitor, the parasitic load for both nodes of the filter is approximately

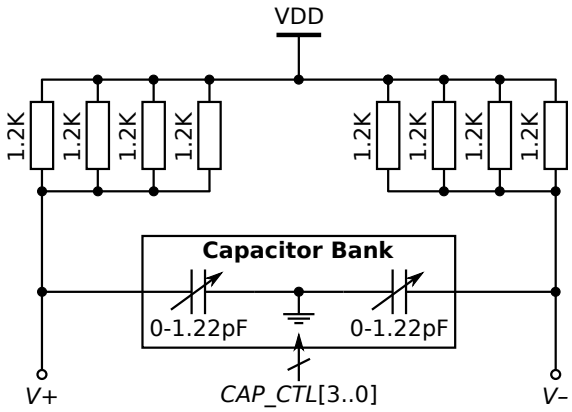

Figure 5.23: Schematic of the 1st-order RC filter

400 fF, which was significantly higher than was originally designed for. As a compromise between the cut-off frequency on the one hand and common mode voltage and signal swing on the other, the resistors were set to $R = 300\,\Omega$.

The exact value of R is not very critical, since, due the gentle roll-off of a 1st-order filter, a slight error in the cut-off frequency will have a negligible impact on performance. An error in the value of R will also change the output common-mode voltage, but a small error can be corrected for by adjusting the bias current of the current sources in the DAC. However, mismatch between the two resistors poses a more serious threat. Since the DC current flowing through both resistors is the same, a mismatch in the resistor values will introduce a DC offset at the input of the limiter. This offset would result in the limiter output having a duty cycle that is not 50 %, which, in turn, would degrade the performance of the DPA. To ensure excellent matching between the resistors, each 300 Ω resistor is implemented as a parallel combination of four resistors of 1.2 kΩ where each resistor has a size of $W/L = 2\,\mu m/12.5\,\mu m$. This results in a mismatch of 0.33 %.

One of the goals of this design is to achieve frequency-agile operation. For operation at 2.4 GHz, the parasitic capacitance is sufficient (in fact, it is too much) for filtering. Thus, no explicit capacitors are needed. For lower frequencies, however, the cut-off frequency of the filter has to be adjusted. For this reason, a programmable capacitor bank is implemented. To save on design time, the 4-bit capacitor bank that was designed for the first gain stage of the limiter was reused for this purpose. However, since the resistance of this filter is considerably lower than the output impedance of a single differential pair, a higher capacitor value is needed to achieve the same cut-off frequency. To accomplish this, several of these banks are connected in parallel. Although the lowest target frequency for this design was 900 MHz, the number of these capacitor banks was chosen to be larger than necessary to make the width of the filter roughly the same as the width of the HR-RF-DAC. The additional occupied area would otherwise have been largely unused. In total, 14 capacitor banks were connected in parallel, each with a maximum capacitance of approximately 87 fF. This brings the total maximum capacitance to 1.22 pF. Together with the 400 fF of parasitic capacitance, this results in a minimum cut-off frequency for the filter of approximately 330 MHz. This should allow for ample possibility of testing the phase modulator at lower frequencies.

Note that the center node of the capacitor banks is connected to GND. Although, so far,

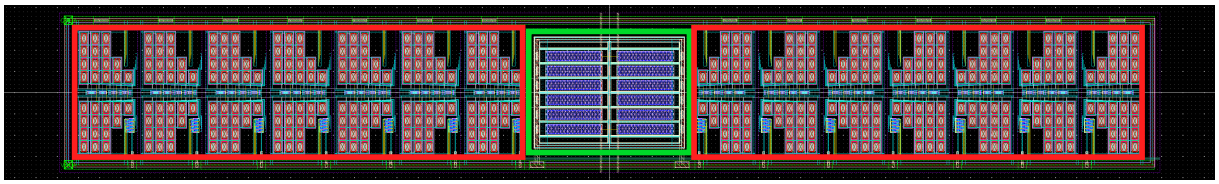

Figure 5.24: Layout of the RC showing the programmable capacitor banks (red) and the resistors (green)

the the capacitors have been shown in parallel with the resistors (i.e. the center node connected to VDD), this was done primarily to simplify the schematics. To a first order approximation, VDD and GND are equivalent when analyzing AC conditions. In practice, however, since GND is the reference node for the entire circuit, more attention is often paid to minimize the GND impedance than the impedance of VDD. Furthermore, for NMOS devices, the performance is much more dependent on the voltage relative to GND rather than the voltage relative to VDD. For this reason, whenever a voltage needs to be stabilized (i.e. low-pass filtered), it is beneficial to do so relative to GND.

The layout of the RC filter is depicted in Figure 5.24. The resistors (in green) are positioned in the middle, while the 14 capacitor banks (in red) are divided evenly across both sides to make the layout as symmetrical as possible. To improve matching between the resistors, dummy resistors are added at the top and bottom. The two nodes V_+ and V_{--} are routed vertically through the center of the filter on m6.

5.5 Top-Level Layout

All of the circuit blocks discussed in this chapter were laid out and interconnected to form the final top-level layout of the phase modulator. An overview of this layout can be seen in Figure 5.25. All of the individual blocks are labeled and the most relevant signal connections are highlighted. The layout was made to be as symmetrical as possible around the vertical axis with the I-DACs on the left and the Q-DACs on the right. The outputs of all six DACs are consolidated in the center and are routed along the vertical axis, via the filter, to the input of the limiter.

The bias block[4] is located to the left of the phase modulator. The main purpose of this block is to provide the bias currents for both the DACs and the limiter, deriving them from a 100 µA off-chip current reference. Although not discussed in detail, this is done using a series of current divisions using current mirrors. This block also contains the bias circuit for the DACs, discussed in Section 5.2.5. Care was taken to ensure that current source and cascode transistors in this circuit are surrounded by the same environment as the transistors in the DACs to ensure good matching.

To the left of the limiting amplifier is a buffer[4], meant to drive a 50 Ω off-chip load. This buffer was added to allow for testing the phase modulator separately from the rest of the TX system. During normal operation, however, this buffer is inactive, and as such, it is not discussed in this thesis.

The LO lines and the I/Q digital input signals are shown in the figure in blue and red, respectively. Due to their physical length, these lines present a large capacitive load (approximately

[4]Both the bias circuit and the 50 Ω buffer were designed and laid out by Milad Mehrpoo

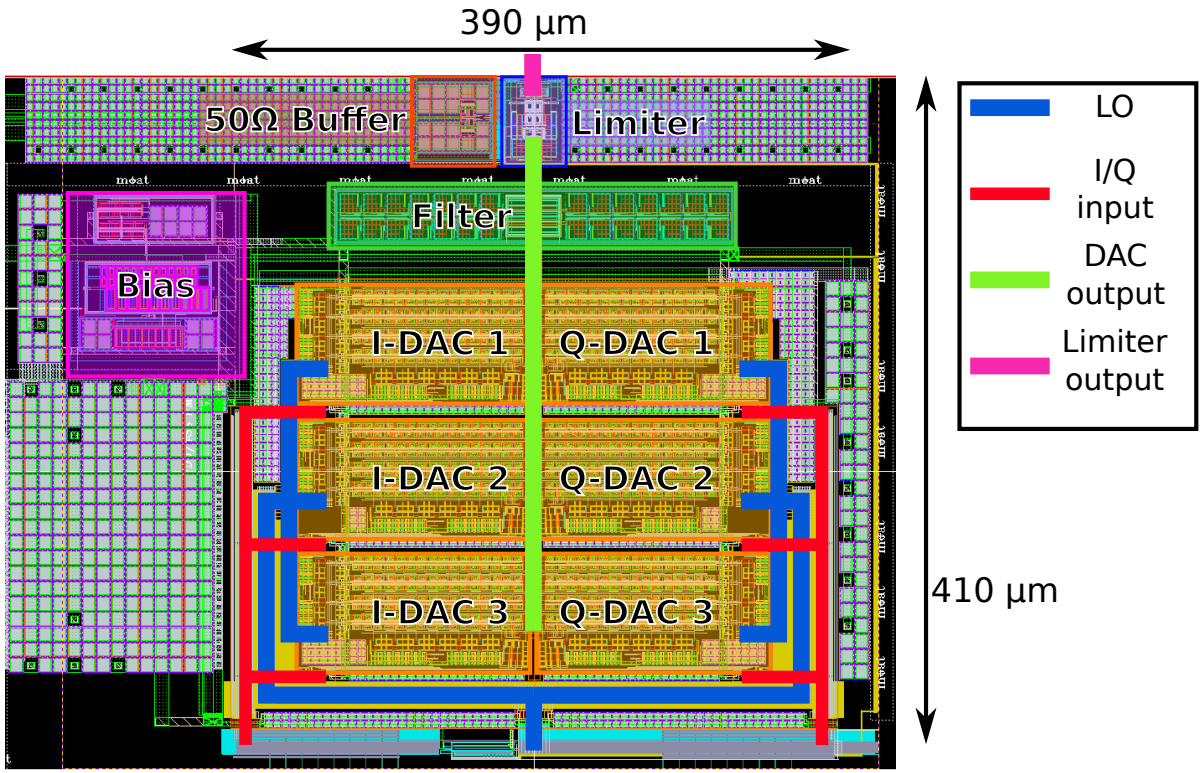

Figure 5.25: Top level layout of the phase modulator

40 fF)to the buffers that drive them. To ensure fast enough rise and fall times for these digital signals, intermediate buffers were placed at several points. For the I and Q lines, additional buffers were only necessary directly at the input of the phase modulator, since as explained in Section 5.2.7, each DAC already has an buffers for these lines. The LO lines are laid out as a tree to achieve as little delay mismatch between the six DACs as possible. Buffers were added at each point where the lines branch off: at input of the phase modulator, to the left of I-DAC 2 and to the right of Q-DAC 2.

In order to minimize the delay mismatch between individual signals in both the I/Q and LO buses, attention was paid to make sure each signal line had the same parasitic capacitance. To achieve this these buses were surrounded by dummy lines connected to GND, so that the parasitics of the outer lines would match those of the inner lines. Furthermore, the lines were routed as a cohesive bus for the entire length of the lines shown in Figure 5.25 even though each of the DACs only needed two of the eight LO phases.

Overall, the phase modulator, excluding the bias block is 390 µm wide and 410 µm high, bringing its total area to 0.16 mm^2. The bias block is a square of approximately 110 µm by 110 µm, requiring an additional area of 0.012 mm^2.

5.6 Conclusion

This chapter described the major design steps to arrive at the final implementation of this phase modulator concept in 40 nm TSMC technology. The majority of the design effort was focused on the HR-RF-DAC which consists of six mixing DACs. These DACs are based on the current

steering architecture which allows them to achieve high-speed operation. Furthermore, they incorporate the mixing operation, which in this design is performed in the digital domain, before the digital-to-analog conversion. This was chosen over the traditional double-balanced mixer architecture to reduce the parasitic capacitance at the output, to make it easier to achieve fast switching and to leave more voltage headroom for signal swing. One of the main disadvantages of this decision is the high level of LO feed-through due to charge injection at switches in the DAC cells. However, this disadvantage was countered by adding cross-connected dummy switches to cancel out the injected charge.

Besides the HR-RF-DAC, the other main components of this phase modulator design are the RC low-pass filter and the limiting amplifier. The limiting amplifier consists of an analog and a digital part. The former provides the majority of the gain whereas the latter serves to convert the signal to CMOS levels in order to drive the DPA. The first analog gain stage of the limiting amplifier additionally serves as the second stage of the low-pass filter, to better suppress the higher harmonics at the output of the HR-RF-DAC. To allow for frequency-agile operation, both filter stages are equipped with a programmable capacitor bank. This capacitor bank is disabled for operation at 2.4 GHz but can be enabled to lower the filter cut-off frequency for operation at lower frequencies.

As with any RF system, the layout plays a critical part in the overall system performance. A large part of the design effort focused on making a compact layout with as little coupling as possible between the various high-frequency signals. Nevertheless, when the layout was nearing completion, it was found that the parasitic capacitance at the output of the HR-RF-DAC was severely underestimated in the initial stages of the design. To compensate for this, a lower value was needed for the filter resistor, but since this resistor value also determined the output common mode voltage and the signal swing, this, in turn, prompted an increase in bias current. Due to the chosen sizes of the current source transistors, the required increase current could not be accommodated without modifying the whole layout. As a compromise, the current was increase as much as the transistor sizing allowed, which set the filter cut-off frequency to approximately 1.3 GHz. However, despite this frequency being so low, the phase modulator was found to still have acceptable performance, as will be shown in the following chapter.

Chapter 6

Simulation Results

In Chapter 5, the design process was discussed in detail to arrive at the finalized phase modulator implementation. This chapter presents the simulation results to verify the system performance. All simulations were performed at the post-layout level, unless stated otherwise.

First, the static phase error is considered in Section 6.1, both without mismatch and with. Section 6.2 deals with the dynamic performance of the phase modulator, which is ascertained through a single-tone simulation and a simulation of the step response. Then, in Section 6.3, the phase modulator is tested as part of a full TX system both at 2.4 GHz and at 900 MHz LO frequency. Finally, some conclusions are drawn in Section 6.4.

6.1 Static Performance

The static performance of the phase modulator is evaluated by computing the phase error versus input phase of the phase modulator. To achieve this, a constant input codeword is supplied to the phase modulator in a transient simulation and the output is allowed to settle, thus eliminating any dynamic effects, such as inter-symbol interference (ISI). In order to decrease the time required for the simulations, the phase error is computed using a single transient simulation where the input is a staircase function, where each step has a duration of 10 ns. This eliminates the need to run a separate simulation for each input code.

6.1.1 Static Phase Error

A post-layout static simulation of the phase modulator was performed, in which the input phase was swept from 0° to 360° in 4°-steps. The resulting phase error is depicted in Figure 6.1. As can be seen from the figure, the maximum phase error is approximately 1.7°, which is significantly higher than was predicted in Section 4.3. Furthermore, the shape of the phase error is approximately sinusoidal, with a period of 180° degrees. This indicates that this error is neither caused by insufficient filtering, nor by insufficient harmonic rejection, as in these cases the period of the error would be 45° and 90°, respectively. Instead, this error is most likely caused by a systematic mismatch between the I and Q parts of the HR-RF-DAC. Since at the schematic level, the HR-RF-DAC is perfectly symmetrical with respect to I and Q, it is likely that this error is caused by layout parasitics.

Although this error is significantly higher than designed for, since it is of a static nature, it is well suited to be corrected for using the static DPD functionality provided by the LUTs

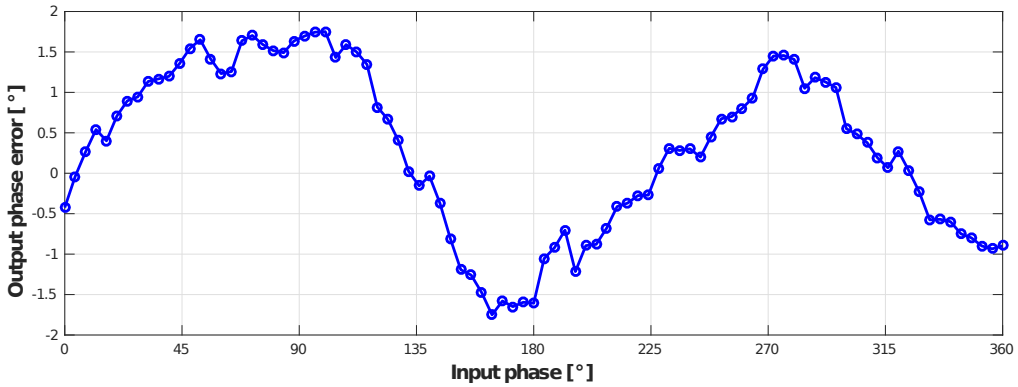

Figure 6.1: Static phase error at the output of the limiter

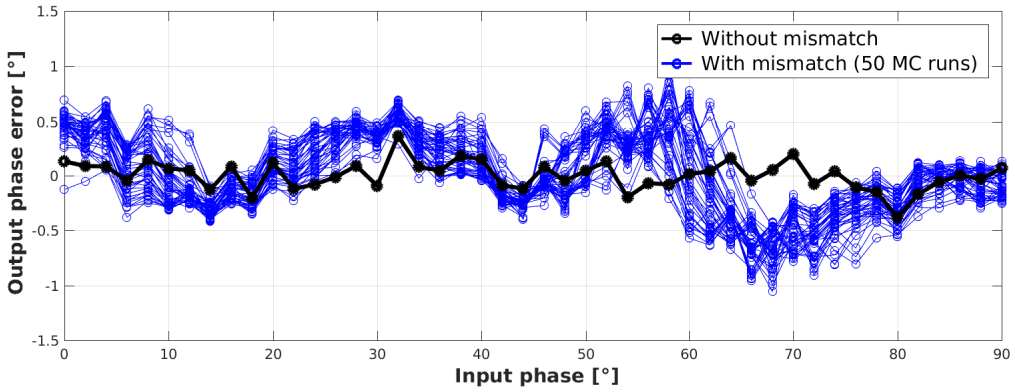

Figure 6.2: Static phase error at the output of the limiter with and without mismatch (schematic level)

at the phase modulator input. However, this option has not been investigated further, and the remainder of the simulation results presented in this chapter show the performance without DPD.

6.1.2 Monte-Carlo Analysis

To estimate the effect of mismatch on the static phase error, a Monte-Carlo analysis was performed. In order to complete this analysis within a reasonable time-frame, the simulations had to be done at the schematic level. The static phase error as a result of 50 Monte-Carlo runs is shown in Figure 6.2, in blue. Additionally, the black curve shows the performance without mismatch.

As the figure shows, without mismatch, the maximum phase error at the schematic level is approximately 0.4°, which is in agreement with the error presented in Section 4.3. This confirms that the large error seen in the post-layout simulation is due to layout effects, rather than the schematic-level design. When mismatch is taken into account, it can be seen that the maximum phase error can be as high as ±1°. This error is expected to be more severe if layout effects are to be considered. Nevertheless, since the error is still of a static nature, through careful calibration, it could still be corrected for to a large extent through the use of DPD.

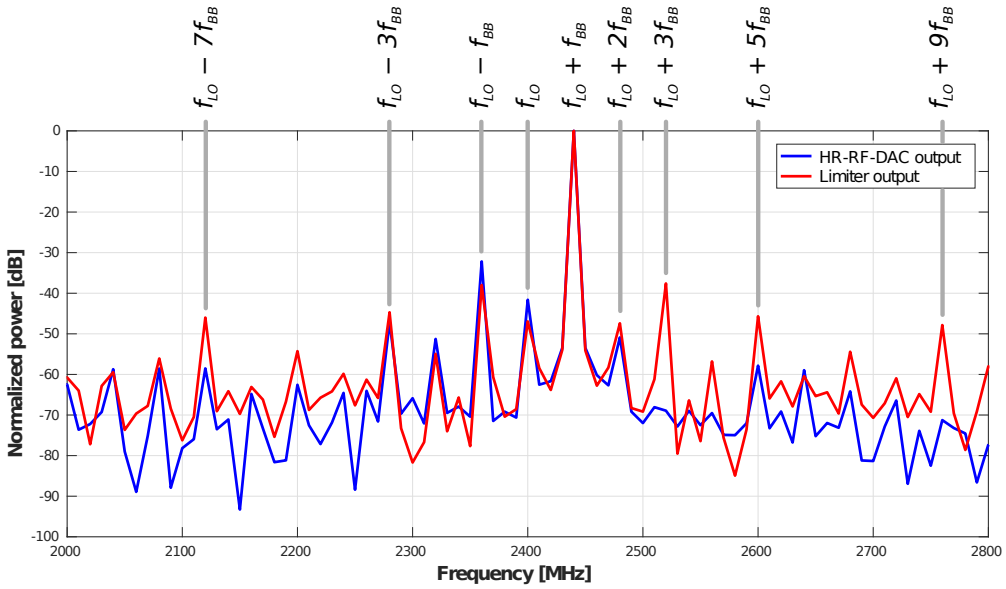

Figure 6.3: Output spectrum of the HR-RF-DAC and the limiter with a single-tone stimulus

6.2 Dynamic Performance

The dynamic performance of the phase modulator describes how well the stand-alone modulator handles a varying input signal. Evaluating the dynamic performance allows one to quickly estimate several relevant figures of merit before testing the modulator in a full TX system. In order to determine the frequency-domain performance of the modulator, a single-tone simulation was performed. The time-domain performance of the modulator was evaluated by simulating the step response.

6.2.1 Single-tone Performance

A single-tone simulation is a commonly-used method to determine the spectral performance of the phase modulator. In such a simulation, sine and cosine waveforms are applied to the I and Q inputs, respectively, which translates to a ramp of constant slope in terms of the input phase. This results in a constant frequency shift away from the LO frequency. In this case, the LO frequency is $f_{LO} = 2.4\,\text{GHz}$ and the tone distance or baseband frequency is chosen to be $f_{BB} = 40\,\text{MHz}$, as this is half of the maximum channel bandwidth that the modulator is designed for. The resulting spectrum, both at the output of the HR-RF-DAC and at the limiter output is depicted in Figure 6.3. In the figure, several relevant frequencies are labeled for easy reference.

The most dominant tone in the spectrum is, of course, the upconverted baseband tone, which is present at $f_{LO} + f_{BB}$. Two relevant performance metrics are the LO feed-through and the image rejection ratio. The tones associated with these metrics fall on f_{LO} and $f_{LO} - f_{BB}$, respectively. The figure shows, that, at the output of the limiter, the LO feed-through is $-48\,\text{dBc}$, which demonstrates the effectiveness of the charge-injection cancellation circuit proposed in Section 5.2.2. The image at the limiter output has a power of $-38\,\text{dBc}$, and is caused by phase mismatch between I and Q due to layout effects. Note that both these tones fall in-band, and the image, in particular, will limit the achievable EVM.

Another tone that is present in the spectrum of both the HR-RF-DAC and the limiter is the tone located at $f_{LO} - 3f_{BB}$. This tone is caused by third-order intermodulation of the upconverted basedband tone ($f_{LO} + f_{BB}$) and its third harmonic, which appears at $3f_{LO} - f_{BB}$. To simplify the discussion, these tones will be referred to as f_1 and f_2, respectively. Although the phase modulator employs the harmonic rejection architecture to suppress the third harmonic, this suppression is not ideal due to layout effects. As is typical for third-order intermodulation, mixing products are produced at $2f_1 - f_2$ and $-f_1 + 2f_2$. The latter product falls on a very high frequency far outside the band of interest, namely $5f_{LO} - f_{BB}$. On the other hand, the former product presents itself at $-f_{LO} + 3f_{BB}$. Folding this product to a positive frequency results in the tone at $f_{LO} - 3f_{BB}$ as is shown in the phase modulator's output spectrum. This type of distortion is known as third-order counter-intermodulation (C-IM3). [57], [58]

The above discussion has covered the most relevant tones in the HR-RF-DAC output, as the other tones are considered to be too close to the simulation noise floor. However, the output of the limiting amplifier shows additional tones in the spectrum. These tones are caused by two separate effects. Due to the limiting amplifier's inherent non-linearity, intermodulation of the upconverted baseband tone ($f_{LO} + f_{BB}$) and the C-IM3 tone ($f_{LO} - 3f_{BB}$) causes additional tones at $f_{LO} - 7f_{BB}$ and $f_{LO} + 5f_{BB}$ [58]. Furthermore, since a limiter eliminates any amplitude modulation in the input signal such that only phase modulation remains, this results in an output spectrum which is symmetrical around the most dominant tone, namely $f_{LO} + f_{BB}$ [18]. As a result, the LO feed-through is mirrored to $f_{LO} + 2f_{BB}$, the image is mirrored to $f_{LO} + 3f_{BB}$ and finally, the C-IM3 tone at $f_{LO} - 7f_{BB}$ is mirrored to $f_{LO} + 9f_{BB}$.

Apart from the upconverted baseband signal, all the tones discussed above will degrade the performance of the phase modulator, either in-band or out-of-band. However, the analysis above indicates that these tones are all generated due to either LO feed-through, the image, or the third-order non-linearity of the HR-RF-DAC. As such, improving any of these performance characteristics will improve the phase modulator's overall performance in future designs.

6.2.2 Step Response

Whereas the single-tone test gives a good indication of the phase modulator's frequency-domain performance, a simulation of the step response is instrumental to characterizing the time-domain behavior. To simulate the step response of the phase modulator, a single transient simulation was run in which steps of various sizes in the input phase were applied at 10 ns intervals. The output waveform of the phase modulator was then down-converted to baseband and filtered with a 20^{th}-order Bessel filter with a constant group delay up to approximately 2.4 GHz. This filter effectively suppressed the higher harmonics in the waveform (except the 2^{nd} harmonic, which falls on 2.4 GHz after down-conversion), while maintaining the signal's time-domain characteristics.

To test the response for small input phase steps, the input step size was swept from $-10°$ to $10°$ in $2°$ increments. The resulting step responses can be seen in Figure 6.4. As shown in the figure, for each step, the output phase settles within 1 ns. However, it can be seen that the accuracy of the phase modulator is limited, both in terms of the static error, as discussed in Section 6.1.1, and in terms of phase noise. Furthermore, both the overshoot as well as the propagation delay vary somewhat between different steps, which is due to the phase modulator's non-linear and time-variant nature.

The step response was also simulated for larger step sizes. Figure 6.5 shows the response for steps between $20°$ and $340°$ in $40°$ increments. As can be seen from the figure, for steps larger than $20°$, the output phase experiences wrapping and becomes negative. To illustrate the cause

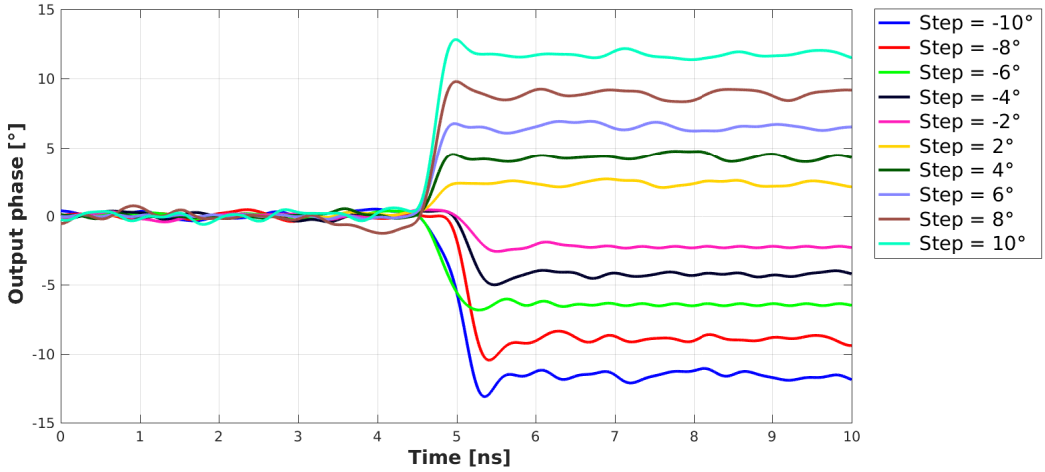

Figure 6.4: Step response for small phase steps

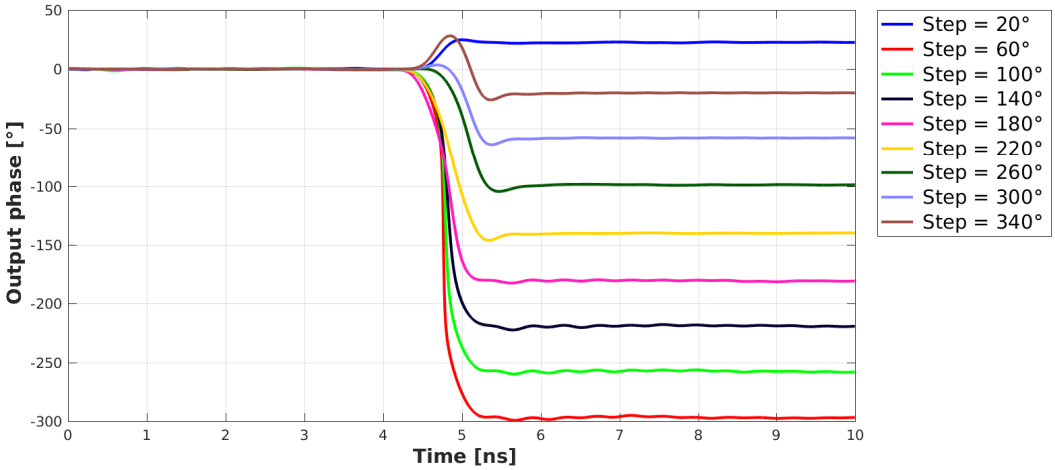

Figure 6.5: Step response for large phase steps

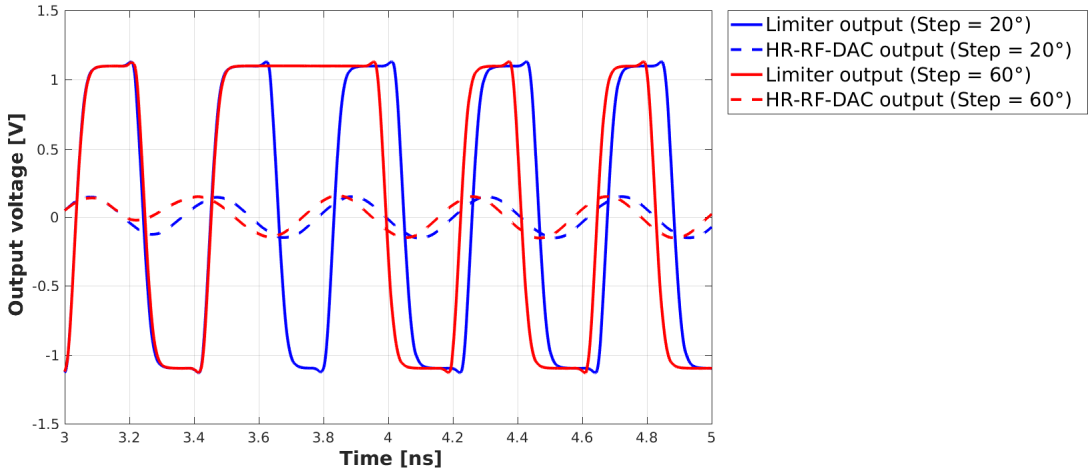

Figure 6.6: Time-domain waveforms for 20° and 60° input phase steps

of this wrapping behavior, the time-domain waveforms, both at the output of the HR-RF-DAC and at the output of the limiter, are depicted in Figure 6.6 for the 20° step and the 60° step. As the figure shows, when a phase step is applied (just before the 3.2 ns mark), the HR-RF-DAC output waveform is temporarily disturbed to accommodate the phase shift. For larger steps this disturbance is stronger, and it can be seen that for a step of 60°, the HR-RF-DAC output does not cross zero during one period. Due to this fact, during that period, the output of the limiter does not toggle, resulting in one skipped pulse, which corresponds to a negative phase step. This phase wrapping behavior will degrade the phase modulator performance. However, because large phase steps are most likely to occur when the signal crosses the origin, i.e. when the signal envelope is small, the degradation is expected to be minimal.

Figures 6.4 and 6.5 show that, regardless of the step size, the 10 %-90 % rise/fall time of the modulator's output phase is $t_r \leq 0.8$ ns. By assuming that the phase-modulator behaves like a 2^{nd}-order linear system, a rough estimate of the modulation bandwidth can be obtained using [59]:

$$t_r \approx \frac{1.8}{\omega_n} \tag{6.1}$$

where ω_n is the system's bandwidth in radians. Solving for ω_n results in a (single-sideband) modulation bandwidth of approximately 360 MHz. However, in reality, the bandwidth of the phase modulator will be limited to the Nyquist frequency of the input data, which for a 600 MS/s input signal is 300 MHz. This corresponds to a DSB bandwidth of 600 MHz, which is sufficient, according to the requirement set in Section 3.2.

6.3 QAM-Modulated Signal

To test the performance of the phase modulator within the context of the full TX system, several transient simulations were run with the phase of a QAM-modulated signal as the input. The envelope path was assumed to be ideal, apart from using the same interpolation filter as in the phase path. After the simulation, the envelope and phase paths were recombined to produce the RF output.

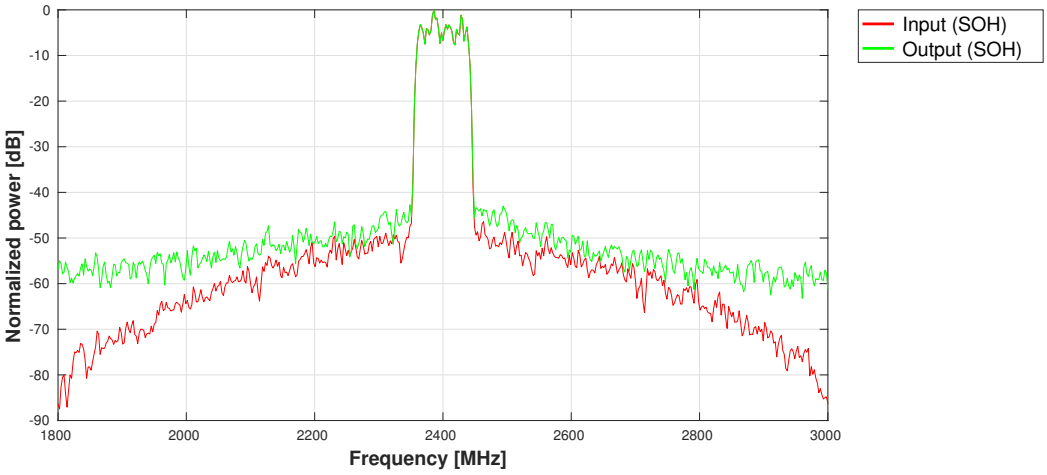

(a)

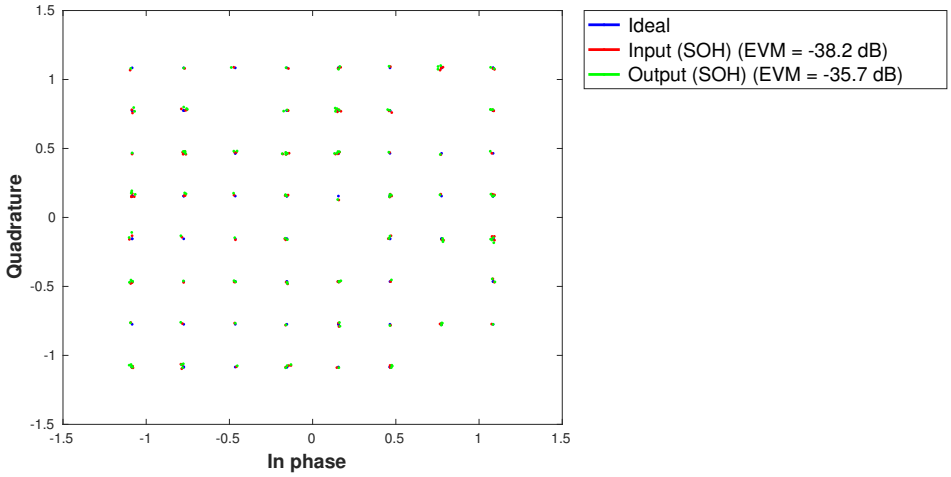

(b)

Figure 6.7: (a) Power spectra and **(b)** constellation diagrams of the RF output and input for a 80 MHz 64-QAM signal at 2.4 GHz with SOH interpolation

So far, 256-QAM has been used in system-level simulations. However, with random input data, significantly more than 256 symbols need to be simulated to fill the entire constellation diagram. For a signal with 80 MHz bandwidth, 256 symbols translate to 3.2 µs of simulation time, which excludes any overhead due to the pulse-shaping filter. Since with the available simulation infrastructure, each microsecond took approximately one day to simulate, it was chosen to perform these simulations with a 64-QAM signal, instead. In this way, the time required for simulation could be reduced, while still covering most of the constellation diagram.

6.3.1 80 MHz 64-QAM with SOH Interpolation at 2.4 GHz

The first of the simulations with a modulated signal involved a 64-QAM signal with 80 MHz bandwidth with a center frequency of 2.4 GHz. The input data was up-sampled to 2.4 GS/s using

Table 6.1: Power consumption of the phase modulator for a 80 MHz 64-QAM signal at 2.4 GHz with SOH interpolation

Component	Power consumption [mW]
HR-RF-DAC + filter (analog supply)	3.4
HR-RF-DAC (digital supply)	28.1
Limiting amplifier	1.8
Total	33.3

SOH interpolation, which, as explained in Section 3.4, is the default configuration of the system. The signal spectra and constellation diagrams both at the input and output of the TX system are depicted in Figure 6.7a and Figure 6.7b, respectively. As the figure shows, in terms of the close-in spectrum, the output signal follows the input quite closely, with the non-linearity of the output is only approximately 5 dB higher than that of the input. The ACPR of the output (assuming 40 MHz channel spacing) for the lower and higher adjacent channels is $\text{ACPR}_L = -48\,\text{dBc}$ and $\text{ACPR}_H = -47\,\text{dBc}$, respectively, which meets the requirement of $-45\,\text{dBc}$. Unfortunately, further out of band, the noise floor does not reach the target of $-60\,\text{dBc}$. Instead, the level of the noise floor is at approximately $-56\,\text{dBc}$.

In terms of the in-band performance, it can be seen that the constellation diagram of the output signal resembles that of the ideal signal quite closely. The EVM of the output signal is $-35.7\,\text{dB}$, which is less than 3 dB higher than that of the input and meets the $-35\,\text{dB}$ requirement.

It should be noted that these simulation results are achieved without the use of DPD or other calibration techniques. If some form of calibration is applied, such as predistorting the input phase to compensate for the static error, or slightly skewing the various LO clocks to improve the image rejection, it should be possible to achieve better performance.

The power consumption of the various components of the phase modulator is shown in Table 6.1. The table shows that for this modulated signal, the phase modulator uses a total of 33.3 mW, the majority of which (28.1 mW) is consumed by the digital circuitry. This circuitry includes mixing DAC logic that drives the differential pair switches in the unit cells, as well as any buffers in the LO and data paths.

6.3.2 80 MHz 64-QAM with ZOH Interpolation at 2.4 GHz

As was discussed in Section 3.4, while applying SOH interpolation to the input signal does suppress spectral replicas at 600 MHz away from the carrier, it also degrades the signal's EVM and ACPR. As such, the option to bypass the SOH interpolation was included in the design, such that the signal would exhibit a ZOH characteristic. In this way, it can be determined whether or not the performance of the phase modulator is limited by the input signal. The bandwidth and center frequency were kept the same as in the case of SOH interpolation.

The resulting signal spectra and constellation diagrams can be seen in Figure 6.8a and Figure 6.8b, respectively. The spectral replicas at 600 MHz away from the carrier due to ZOH are clearly visible, and the noise floor reaches only approximately $-50\,\text{dBc}$ at its lowest point. In terms of close-in nonlinearity, $\text{ACPR}_L = -49\,\text{dBc}$ and $\text{ACPR}_H = -44\,\text{dBc}$. The EVM is $-35.4\,\text{dB}$, so the in-band performance is marginally worse than with SOH interpolation.

Overall, the performance can be seen to be somewhat worse than before, despite the fact that the ZOH-interpolated input signal is much cleaner in terms of both EVM and ACPR.

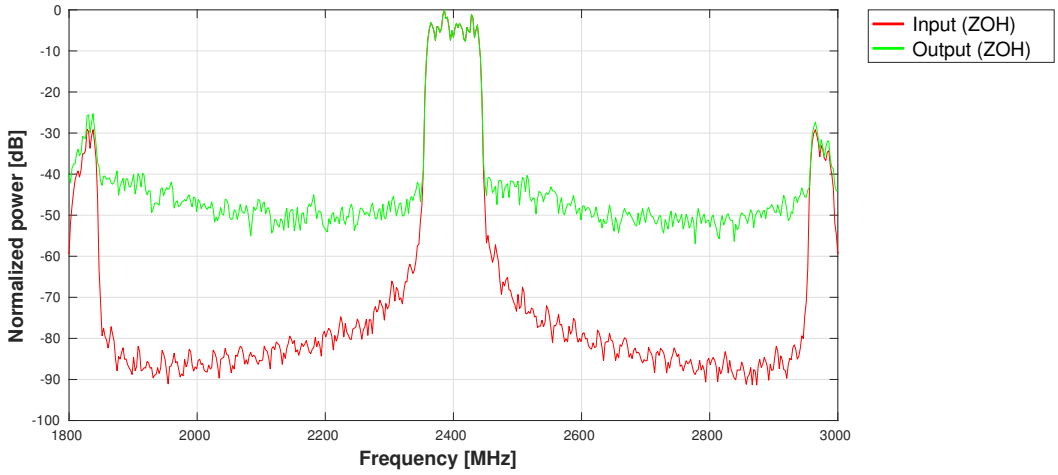

(a)

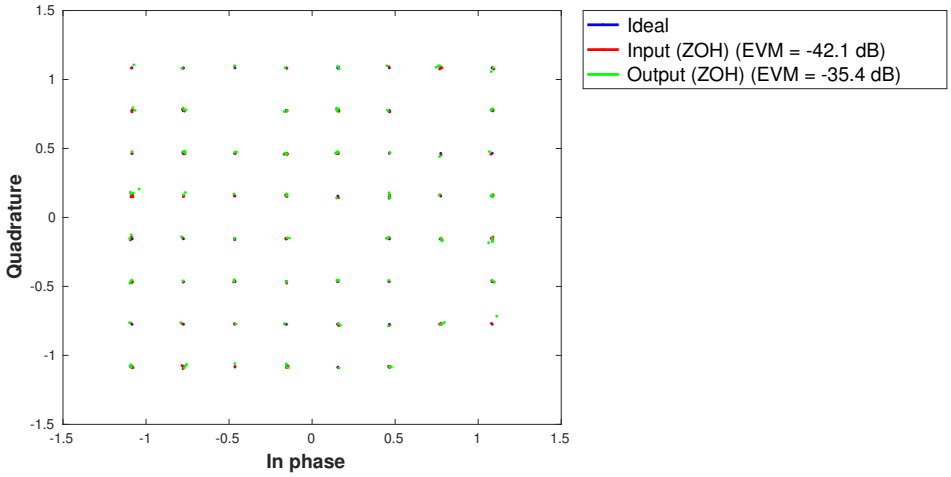

(b)

Figure 6.8: (a) Power spectra and **(b)** constellation diagrams of the RF output and input for a 80 MHz 64-QAM signal at 2.4 GHz with ZOH interpolation

This indicates that the phase modulator performance is not limited by the input signal, but rather by its own non-linearity. The far-out performance degradation with respect to the SOH case is caused by spectral regrowth around the spectral replicas. The close-in and in-band performance degradation is most likely the result of the larger phase steps that are inherent to ZOH interpolation, such that phase wrapping (as discussed in Section 6.2.2) is more common in the signal transitions.

The phase modulator's power consumption, when used with this input signal is summarized in Table 6.2. It can be seen that the consumption is almost unchanged with respect to the SOH case, apart from a marginally lower consumption of the digital circuitry in the HR-RF-DAC. This is due to the less frequent switching of the I/Q data lines as a result of the ZOH interpolation.

Table 6.2: Power consumption of the phase modulator for a 80 MHz 64-QAM signal at 2.4 GHz with ZOH interpolation

Component	Power consumption [mW]
HR-RF-DAC + filter (analog supply)	3.4
HR-RF-DAC (digital supply)	26.6
Limiting amplifier	1.8
Total	31.8

Table 6.3: Power consumption of the phase modulator for a 30 MHz 64-QAM signal at 900 MHz with SOH interpolation

Component	Power consumption [mW]
HR-RF-DAC + filter (analog supply)	3.4
HR-RF-DAC (digital supply)	11.0
Limiting amplifier	1.1
Total	15.6

6.3.3 30 MHz 64-QAM with SOH Interpolation at 900 MHz

As the goal of this thesis work is to develop a phase modulator which support frequency-agile operation, it is important to show the system performance at lower frequencies. For this, another simulation was run, with the center frequency set to 900 MHz. Again, SOH interpolation was used to up-sample the signal to the carrier frequency, while the bandwidth was scaled proportionally to 30 MHz. While in the case of operation at 2.4 GHz, the parasitic capacitance at the output of the HR-RF-DAC limited the cut-off frequency of the filter to a much lower value than desired, this is no longer true for operation at 900 MHz. As such, for this simulation, it was possible to set the filter cut-off frequency equal to the carrier frequency.

The resulting signal spectra and constellation diagrams are shown in Figure 6.9a and Figure 6.9b, respectively. As can be seen from the spectrum, the output very closely the input signal, with $\text{ACPR}_L = -47\,\text{dBc}$ and $\text{ACPR}_H = -49\,\text{dBc}$. Further out, the noise floor can be seen to be at approximately $-60\,\text{dBc}$, as per the requirement, which suggests that in the 2.4 GHz case, the higher noise floor is caused by dynamic effects, which are more prevalent at higher frequencies. In terms of the in-band performance, the EVM is 1.6 dB worse than when operating at 2.4 GHz. This is likely caused by the somewhat weaker attenuation of the harmonics relative to the fundamental tone, due to the higher cut-off frequency of the filter. However, the exact cause of this degradation has not been investigated in detail.

The power consumption of the phase modulator operating under these conditions is shown in Table 6.3. As can be seen from this table, while the consumption of the analog portion of the HR-RF-DAC remains unchanged, the consumption of the other two blocks has decreased significantly. In the case of the digital part of the HR-RF-DAC, the consumption is 61 % lower than when operating at 2.4 GHz, which is almost proportional to the 62.5 % decrease in frequency. The power consumption of the limiter is only 39 % lower than at 2.4 GHz, since the limiter also contains analog gain stages, the consumption of which is unaffected by the frequency decrease. Overall, the digital portion of the HR-RF-DAC has the highest impact on the overall power consumption, decreasing the total consumption to 15.6 mW.

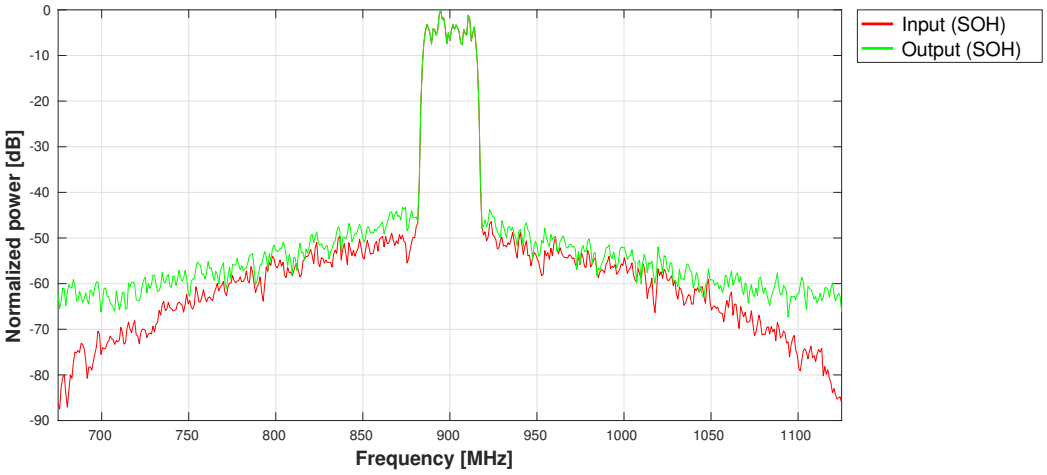

(a)

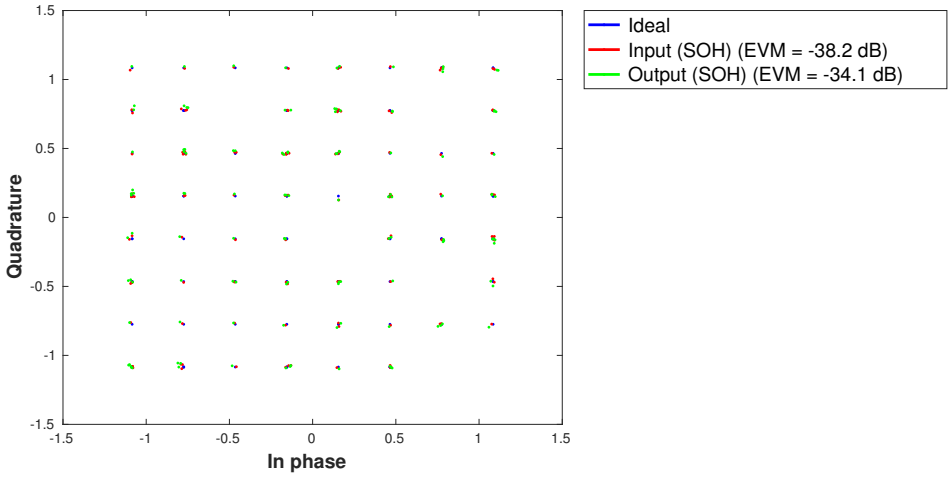

(b)

Figure 6.9: (a) Power spectra and **(b)** constellation diagrams of the RF output and input for a 30 MHz 64-QAM signal at 900 MHz with SOH interpolation

6.4 Conclusion

In this chapter, the post-layout simulation results of the finalized phase-modulator design were presented. A static simulation of the output phase error with respect to the input phase revealed the error to be as high as 1.7°, which is significantly higher than was predicted in Section 4.3. This error is caused by layout effects and is not inherent to the concept or the schematic-level design. Furthermore, additional error is expected to present itself in the fabricated circuit, as illustrated by the (schematic-level) Monte-Carlo simulation. This static error, however, is well suited to be corrected for through predistortion by utilizing the LUTs preceding the phase modulator. As such, better performance can be expected after calibration, although this has not been investigated in simulation.

The phase modulator's dynamic performance was investigated with both a single-tone sim-

ulation as well as a simulation of the step response. An analysis of the single-tone simulation results revealed a number of tones in the phase modulator's output spectrum, the most relevant of which can be traced back to phase mismatch between I and Q, LO feed-through and the HR-RF-DAC's third-order non-linearity. The highest unwanted tone is the image which has a power of $-38\,\mathrm{dBc}$ and since it falls in-band, it limits the achievable EVM.

The step response simulation showed the modulator's time-domain characteristics. It was shown that the output phase settles to its final value, albeit with limited accuracy, within approximately 1 ns, and that the rise/fall time of the phase signal is less than 0.8 ns. The latter indicates that the phase modulator's analog modulation bandwidth is in the order of 360 MHz which is high enough to not be the limiting factor in the overall modulation bandwidth. For steps larger than approximately 20°, the phase wraps to the negative side, which in the time domain is represented by a single skipped pulse in the limiter's output voltage. However, because such large phase steps are much more likely to occur when the I/Q signal crosses the origin (i.e. when the signal power is low), this wrapping behavior has little effect on the overall system performance.

Finally, the phase modulator was also tested as part of a full TX system, which was otherwise considered ideal. It was shown that with an 80 MHz 64-QAM signal interpolated with SOH, and with an LO of 2.4 GHz, the modulator meets the target requirements for ACPR and EVM while consuming 33.3 mW of power from a 1.1 V supply. However, wideband distortion causes the effective noise floor to be higher than desired. Bypassing the SOH interpolation (such that ZOH results) does not improve the performance in any way, despite the seamingly cleaner input signal. In fact, the spectral replicas and the additional regrowth that occurs around them further raises the effective noise floor.

The phase modulator was also tested at 900 MHz center frequency (and a proportionally scaled signal bandwidth), to demonstrate the frequency-agile nature of this design. At this frequency, the system showed better ACPR and the target for the noise floor of $-60\,\mathrm{dBc}$ was reached. However, the EVM performance was marginally worse, presumably due to the filter cut-off frequency being set differently, relative to the LO frequency. Since the majority of the power is consumed by the digital logic in this design, operating at a lower frequency resulted in a significantly lower total power consumption, namely 15.6 mW.

Chapter 7

Conclusion

In this thesis work, the need for a wideband phase modulator for polar TX systems was motivated, and a promising phase modulator concept was brought forward as a possible solution. This concept was implemented in TSMC's 40 nm CMOS technology and although there was no time to perform measurements on the fabricated chip in the context of this thesis work, extensive simulations were performed to characterize the expected system performance.

In Section 7.1, the simulation results presented in the previous chapter are put into perspective by comparing them to state-of-the-art phase modulators in recent literature. Furthermore, some benefits and shortcomings of the current design are pointed out, in order to aid in trade-offs in future designs. Section 7.2 summarizes the thesis outcome and reflects on whether or not the goals that were set at the beginning of this thesis have been achieved. Finally, in Section 7.3, some recommendations are made for future work, based on the insights gained during the course of this project.

7.1 Discussion

The post-layout simulation results presented in Chapter 6 provide a reasonable prediction of the performance that can be expected of this phase modulator implementation. To put phase modulator performance into perspective, it is important to compare it to other phase modulators reported in recent literature. It should be noted, however, that although there is no shortage of reported phase modulator designs, the majority of these are evaluated stand-alone as they are targeted at constant-envelope modulation schemes such as phase-shift keying (PSK), Gaussian frequency-shift keying (GFSK), or similar. Generally speaking, standards that use these schemes have much less demanding bandwidth requirements, allowing other performance metrics such as power consumption to be optimized for, instead. As such, a comparison to these designs would be of limited use, because the goals of these designs are too dissimilar to the goals set for this thesis work.

Nevertheless, a number of transmitter systems for complex-modulated signals have been reported in recent years, which implement wideband phase modulators for use in either polar or outphasing configurations. For instance, [42] presents an outphasing modulator that supports both 20 MHz 802.11g and 40 MHz 802.11n. The phase modulators used in this design are implemented using the coarse-fine delay line architecture discussed in Section 2.2.3. In [44], a polar transmitter for 20 MHz 802.11g signals is presented in which the phase modulator is implemented using an analog-intensive variant of the constant-envelope direct up-conversion architecture. Fi-

Table 7.1: Comparison of phase modulator performance to literature data

Parameter	**[42]**[a]	**[44]**[a]	**[45]**[a]	**This work**[b]	
Process	32 nm CMOS	65 nm CMOS	45 nm CMOS SOI	40 nm CMOS	
Carrier frequency	2.4 GHz	2.2 GHz	10 GHz	2.4 GHz	900 MHz
Modulation	64-QAM OFDM	64-QAM OFDM	256-QAM OFDM	64-QAM	
Signal bandwidth	40 MHz	20 MHz	133 MHz	80 MHz	30 MHz
EVM	−28 dB	−28 dB	−29 dB	−35 dB	−34 dB
ACPR	−50 dBc[c]	−50 dBc[c]	−36 dBc[c]	−47 dBc	−47 dBc
Noise floor	−50 dBc[c]	−80 dBc[c]	−40 dBc[c]	−56 dBc	−60 dBc
Power consumption	20 mW[d]	9 mw	860 mW[d]	33 mW	16 mW
Area	0.05 mm^2 [e]	0.07 mm^2 [e]	2 mm^2	0.17 mm^2	

[a] Measured performance after calibration
[b] Simulated performance with no calibration
[c] Estimated based on reported plots of the output spectrum
[d] Estimated based on half of the outphasing modulator power consumption
[e] Estimated based on reported chip micrograph

nally, [45] presents an outphasing modulator that targets LTE carrier aggregation in the 10 GHz band and supports bandwidths of up to 133 MHz. This modulator also implements the phase modulators using an analog-intensive constant-envelope direct up-conversion structure. A comparison of various performance metrics of these modulators to those of the modulator developed during this thesis work is summarized in Table 7.1.

Overall, the phase modulator in [42] is the smallest of the modulators, while meeting all the requirements of 802.11g/n, which demonstrates how compact a fully digital implementation can be in an advanced nanometer-scale technology. The modulator developed in [44] has by far the smallest power consumption and despite being heavily based on analog techniques and containing two inductors its total area is only marginally bigger than that of [42]. However, of the modulators presented here, it supports the lowest bandwidth, and despite being designed for 802.11g signals, the measurement results were only reported at a center frequency of 2.2 GHz rather than the standard 2.4 GHz. Furthermore, at −80 dBc, it has by far the lowest noise floor. However, the publication did not make it clear whether or not the measurement was performed with a bandpass filter at the TX output. If this is the case, the unfiltered noise floor could be significantly higher. The modulator reported in [45], supports the largest signal bandwidth, and with its 256-QAM OFDM modulation, it was the first outphasing modulator to surpass 1 Gb/s data rate. However, both its ACPR and noise floor performance is worse than the other modulators, while its power consumption and area are two orders of magnitude higher than the modulator of [44]. The latter, however, is primarily due to the fact that this modulator was designed to deliver 23 dBm peak output power without any external PA.

Of the phase modulators presented in Table 7.1, the one developed during this thesis work supports the largest signal bandwidth, relative to its operating frequency. It also has superior EVM performance, while having a very reasonable ACPR and noise floor. Both the area and the power consumption (at 2.4 GHz) are somewhat larger than [42] and [44], which demonstrates the main drawback of using the harmonic rejection architecture, with its requirement for eight different LO phases. However, because it is a frequency-independent solution, it facilitates true multi-band/multi-standard operation.

Unfortunately, although the modulators compared here share similar goals, the validity of this comparison is still limited. This is because the performance of the modulator developed during this thesis work can only be predicted based on simulation results. Additional non-idealities present in the fabricated circuit could degrade the performance further, whereas calibration techniques such as DPD could result in a dramatic performance improvement. Moreover, this comparison assumes that the phase modulator performance is the dominant factor for the performance of the overall TX system, which is often not the case. Finally, since OFDM modulation is a fundamentally different type of signal than standard QAM, further differences are to be expected. All in all, a true comparison can only be made when the fabricated chip is measured, after proper calibration, ideally with an OFDM-modulated signal. Nevertheless, the data available so far appears promising, and shows that this phase modulator concept warrants further development.

7.2 Thesis Outcome

Throughout the course of this thesis work a phase modulator for digital polar transmitters was designed and implemented, guided by the goals stated in Section 1.5. The novelty of this design lies in the application of the harmonic rejection architecture, originally developed to improve analog mixer performance, and several digitally-intensive techniques to achieve a wide bandwidth and frequency-agile operation. The modulator has been shown to support 80 MHz QAM-modulated signals at a center frequency of 2.4 GHz with adequate performance, despite the significant bandwidth expansion that occurs when converting from Cartesian to polar coordinates. Post-layout simulations of the phase modulator, combined with an ideal envelope path, have shown that the design meets the EVM requirement of -35 dB and the ACPR requirement of -45 dBc. Unfortunately, the noise floor performance is approximately 4 dB worse than desired, but since it is dominated by wideband distortion rather than quantization noise, it is expected that this can be improved using calibration techniques.

One of the major advantages of this phase modulator concept is its ability to support a large range of operating frequencies. To demonstrate this, its performance has also been simulated at 900 MHz, with the bandwidth of the input signal scaled proportionally to 30 MHz. At this frequency, the ACPR remained almost unchanged, and the noise floor improved, to meet the -60 dBc requirement. Unfortunately, the EVM performance was showed a degradation of approximately 1.6 dB, but this can most likely be improved by selecting a slightly lower filter cut-off frequency.

One of the design goals for this phase modulator is to be as digital as possible, to take advantage of the ongoing technology scaling of CMOS technology. This has been achieved to a certain extent, as the interface to this modulator to the outside world is fully digital, and the design relies heavily on digital logic to achieve the mixing operation, to drive the DACs in an effective manner and to support calibration through DPD. However, the core of the modulator, consisting of the DACs, the filter and the first three stages of the limiter, is still analog in nature.

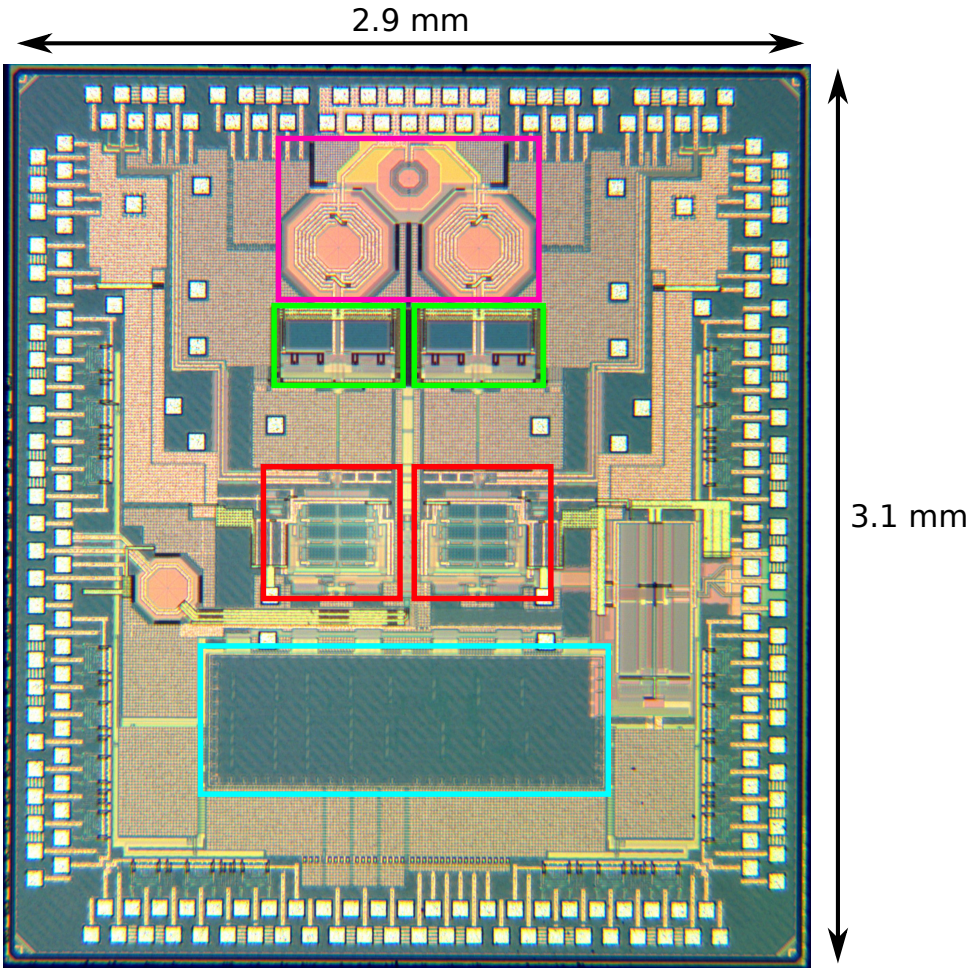

Figure 7.1: Chip micrograph of the full digital polar Doherty transmitter showing two instances of the phase modulator (red), two DPAs (green), the matching network (magenta) and the DSP logic and SRAM LUTs (cyan)

Nevertheless, since the sizing of the DAC unit cells was primarily determined based on mismatch and the tolerances on the physical dimensions of transistors generally improve in finer technology nodes, this design is still expected to benefit from technology scaling.

In order to be cost-effective, it is vital that the modulator occupies a small area on-chip and its power consumption is low. Unfortunately, the use of the harmonic rejection architecture is disadvantageous for both of these performance metrics, due to the additional DACs and LO frequencies that are required. However, despite these issues, and especially considering the added benefits of this design choice in terms of frequency-independent operation and the simplicity of the required filter, the design remains relatively competitive with the other phase modulators considered in the previous section. Furthermore, the use of inductors has been avoided in this design, in part to prevent disturbances from a high-power DPA from inductively coupling to the analog signals in the phase-modulator core. This design choice allows for better integration of the full TX system on a single chip, further reducing the overall costs. Furthermore, avoiding the use of inductors is also beneficial for moving towards finer technology nodes, since inductors do not benefit from scaling.

The design has been fabricated in TSMC's 40 nm CMOS technology, and the finished chip micrograph can be seen in Figure 7.1. As described in Section 1.3 (and in more detail in Section 3.1), the phase modulator is used in a polar Doherty configuration, which contains a main branch and a peak branch. The two instances of the phase modulator can be seen marked in red. Also highlighted are the two DPAs including the thermometer encoder-based envelope modulator (green), the on-chip matching network (magenta) and the DSP logic and SRAM LUTs that precede the envelope and phase modulators. At the time of writing, the fabricated chip has been received from the foundry and the measurement process has begun, in order to bring this project to its conclusion.

7.3 Recommendations for Future Work

Although this thesis work has yielded some interesting results, and the proposed phase modulator concept has been shown to be viable in simulation, there are still numerous areas that warrant further investigation. First of all, the performance of fabricated circuit should be measured, both stand-alone and in the context of the full TX system. This would allow to test how much the performance can be improved with proper calibration. It also allows for testing the system with OFDM signals, which are more relevant in many modern communication standards. Only then can a truly fair comparison be made to the state-of-the-art phase modulators reported in literature.

At the start of this thesis work, the highest target frequency was set to 2.4 GHz, whereas the lowest target frequency was chosen to be 900 MHz. In the broader scope of the SEEDCOM project, however, the goal is to support all frequencies in the low-GHz band, up to approximately 6 GHz. Although it has been shown that very low frequencies can be supported with relative ease, provided that there is enough capacitance in the RC filter, for very high frequencies, this is most definitely not the case. The limiting factor at high frequencies is the parasitic capacitance that is present at the output node of the HR-RF-DAC. In the case of this design, it limits the cut-off frequency of the first filter stage to approximately 1.3 GHz. In order to operate at higher frequencies, methods of either reducing this parasitic capacitance or suppressing its effect should be investigated. This could, for instance, be accomplished using wideband buffers to isolate the output nodes of the individual DACs, or by using a transimpedance amplifier in the output path. Regardless of the chosen method, making this design suitable for higher frequencies without degrading the performance while minimizing the increase in power consumption presents an interesting challenge.

Another research direction worth considering is the choice of interpolation method for the input data. It has been shown that, if no interpolation is performed on the input data, the spectral replica's in the output spectrum prove detrimental to the performance of the TX system. On the other hand, the chosen interpolation method, namely, SOH interpolating in the I/Q domain, as described in Section 3.4, has its own limitations. First of all, both the in-band and out-of-band characteristics of the input signal are degraded. If the intrinsic performance of the phase modulator is improved in future designs, it is quite possible that the chosen interpolation scheme will become the limiting factor in meeting the performance requirements. Furthermore, because the interpolation occurs in the I/Q domain, rather than the phase domain, the I and Q inputs provided to the phase modulator are no longer constant-envelope, which can introduce further errors. As such, an investigation of more advanced interpolation schemes could prove beneficial for future digital polar TX systems.

Finally, due to the bandwidth expansion that occurs when converting from Cartesian to

polar coordinates, polar transmitters will always be at a disadvantage when compared to their Cartesian counterparts with respect to the signal bandwidths that can be achieved. Several methods of reducing the required modulation bandwidth have been investigated in literature. For instance, [60] presents a technique which alters the signal trajectory such that it avoids crossing the origin of the constellation, which would otherwise result in a discontinuity in the phase and a sharp corner in the envelope. In this way, the required modulation bandwidth can be traded off for a slight degradation of other performance characteristics such as EVM and ACPR.

An alternative idea, which surfaced during the course of this thesis work, is to allow the envelope signal to take on a negative value. In this way, if the signal crosses the constellation origin, both the phase discontinuity and the sharp corner of the envelope signal can be eliminated. Of course, if the signal trajectory does not come near the origin, toggling the sign of the envelope signal would exacerbate the bandwidth expansion problem. Thus, with every signal transition, a decision would need to be made, whether toggling the sign of the envelope is beneficial or not. So far, this idea has not been developed beyond an intuitive analysis. However, if it can be shown to work as expected, it could facilitate the development of very wideband polar transmitters in the future.

Bibliography

[1] "Cisco visual networking index: Global mobile data traffic forecast update, 2015-2020", Cisco Systems, Inc., Tech. Rep., 2016.

[2] S. C. Cripps, *RF power amplifiers for wireless communications*, 2nd ed. Artech House, 2006.

[3] B. J. Minnis, P. A. Moore, P. N. Whatmough, P. G. Blanken, and M. P. van der Heijden, "System-efficiency analysis of power amplifier supply-tracking regimes in mobile transmitters", *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 56, no. 1, pp. 268–279, Jan. 2009, ISSN: 1549-8328. DOI: 10.1109/TCSI.2008.927013.

[4] F. Richter, A. J. Fehske, and G. P. Fettweis, "Energy efficiency aspects of base station deployment strategies for cellular networks", in *Vehicular Technology Conference Fall (VTC 2009-Fall), 2009 IEEE 70th*, IEEE, 2009, pp. 1–5.

[5] G. Fettweis and E. Zimmermann, "ICT energy consumption-trends and challenges", in *Proceedings of the 11th International Symposium on Wireless Personal Multimedia Communications*, vol. 2, 2008, p. 6.

[6] D. Rudolph, "Out-of-band emissions of digital transmissions using Kahn EER technique", *IEEE Transactions on Microwave Theory and Techniques*, vol. 50, no. 8, pp. 1979–1983, Aug. 2002, ISSN: 0018-9480. DOI: 10.1109/TMTT.2002.801349.

[7] R. B. Staszewski and M. S. Alavi, "Digital I/Q RF transmitter using time-division duplexing", in *Radio-Frequency Integration Technology (RFIT), 2011 IEEE International Symposium on*, Nov. 2011, pp. 165–168. DOI: 10.1109/RFIT.2011.6141744.

[8] L. R. Kahn, "Single-sideband transmission by envelope elimination and restoration", *Proceedings of the IRE*, vol. 40, no. 7, pp. 803–806, Jul. 1952, ISSN: 0096-8390. DOI: 10.1109/JRPROC.1952.273844.

[9] E. McCune and W. Sander, "EDGE transmitter alternative using nonlinear polar modulation", in *Circuits and Systems, 2003. ISCAS '03. Proceedings of the 2003 International Symposium on*, vol. 3, May 2003, III-594-III-597 vol.3. DOI: 10.1109/ISCAS.2003.1205089.

[10] M. S. Alavi, R. B. Staszewski, L. C. N. de Vreede, A. Visweswaran, and J. R. Long, "All-digital RF I/Q modulator", *IEEE Transactions on Microwave Theory and Techniques*, vol. 60, no. 11, pp. 3513–3526, Nov. 2012, ISSN: 0018-9480. DOI: 10.1109/TMTT.2012.2211612.

[11] W. H. Doherty, "A new high-efficiency power amplifier for modulated waves", *The Bell System Technical Journal*, vol. 15, no. 3, pp. 469–475, Jul. 1936, ISSN: 0005-8580. DOI: 10.1002/j.1538-7305.1936.tb03563.x.

[12] H. Chireix, "High power outphasing modulation", *Proceedings of the Institute of Radio Engineers*, vol. 23, no. 11, pp. 1370–1392, Nov. 1935, ISSN: 0731-5996. DOI: 10.1109/JRPROC.1935.227299.

[13] L. C. N. de Vreede, M. Acar, D. A. Calvillo-Cortes, M. P. van der Heijden, R. Wesson, M. de Langen, and J. Qureshi, "Outphasing transmitters, enabling digital-like amplifier operation with high efficiency and spectral purity", *IEEE Communications Magazine*, vol. 53, no. 4, pp. 216–225, Apr. 2015, ISSN: 0163-6804. DOI: 10.1109/MCOM.2015.7081097.

[14] G. E. Moore, "Cramming more components onto integrated circuits", *Electronics*, pp. 114–117, Apr. 1965, ISSN: 0018-9219.

[15] D. A. Calvillo-Cortes and L. C. N. de Vreede, "Analysis of pure- and mixed-mode class-B outphasing amplifiers", in *Circuits and Systems (LASCAS), 2014 IEEE 5th Latin American Symposium on*, Feb. 2014, pp. 1–4. DOI: 10.1109/LASCAS.2014.6820271.

[16] M. P. van der Heijden and M. Acar, "A radio-frequency reconfigurable CMOS-GaN class-E Chireix power amplifier", in *2014 IEEE MTT-S International Microwave Symposium (IMS2014)*, Jun. 2014, pp. 1–4. DOI: 10.1109/MWSYM.2014.6848251.

[17] N. Nidhi, P.-E. Su, and S. Pamarti, "Open-loop wide-bandwidth phase modulation techniques", *Journal of Electrical and Computer Engineering*, vol. 2011, p. 6, 2011.

[18] B. Razavi, *RF microelectronics*, 2nd ed. Prentice Hall New Jersey, 2011.

[19] A. Behzad, *Wireless LAN radios: System definition to transistor design*. John Wiley & Sons, 2007, vol. 11.

[20] T. A. D. Riley and M. A. Copeland, "A simplified continuous phase modulator technique", *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 41, no. 5, pp. 321–328, May 1994, ISSN: 1057-7130. DOI: 10.1109/82.287002.

[21] N. M. Filiol, T. A. D. Riley, C. Plett, and M. A. Copeland, "An agile ISM band frequency synthesizer with built-in GMSK data modulation", *IEEE Journal of Solid-State Circuits*, vol. 33, no. 7, pp. 998–1008, Jul. 1998, ISSN: 0018-9200. DOI: 10.1109/4.701242.

[22] S. Pamarti, L. Jansson, and I. Galton, "A wideband 2.4-GHz $\Delta\Sigma$ fractional-N PLL with 1-Mb/s in-loop modulation", *IEEE Journal of Solid-State Circuits*, vol. 39, no. 1, pp. 49–62, Jan. 2004, ISSN: 0018-9200. DOI: 10.1109/JSSC.2003.820858.

[23] P.-E. Su and S. Pamarti, "A 2-MHz bandwidth $\Delta\Sigma$ fractional-N synthesizer based on a fractional frequency divider with digital spur suppression", in *2010 IEEE Radio Frequency Integrated Circuits Symposium*, May 2010, pp. 413–416. DOI: 10.1109/RFIC.2010.5477398.

[24] M. H. Perrott, T. L. Tewksbury, and C. G. Sodini, "A 27-mW CMOS fractional-N synthesizer using digital compensation for 2.5-Mb/s GFSK modulation", *IEEE Journal of Solid-State Circuits*, vol. 32, no. 12, pp. 2048–2060, Dec. 1997, ISSN: 0018-9200. DOI: 10.1109/4.643663.

[25] M. J. Underhill and R. I. H. Scott, "Wideband frequency modulation of frequency synthesisers", *Electronics Letters*, vol. 15, no. 13, pp. 393–394, Jun. 1979, ISSN: 0013-5194. DOI: 10.1049/el:19790283.

[26] C. Durdodt, M. Friedrich, C. Grewing, M. Hammes, A. Hanke, S. Heinen, J. Oehm, D. Pham-Stabner, D. Seippel, D. Theil, S. V. Waasen, and E. Wagner, "A low-IF RX two-point ΣΔ-modulation TX CMOS single-chip Bluetooth solution", *IEEE Transactions on Microwave Theory and Techniques*, vol. 49, no. 9, pp. 1531–1537, Sep. 2001, ISSN: 0018-9480. DOI: 10.1109/22.942563.

[27] R. Yu, T. T. Yeo, K. H. Tan, S. Mou, Y. Cui, H. Wang, H. S. Yap, E. Ting, and M. Itoh, "A 5.5mA 2.4-GHz two-point modulation Zigbee transmitter with modulation gain calibration", in *2009 IEEE Custom Integrated Circuits Conference*, Sep. 2009, pp. 375–378. DOI: 10.1109/CICC.2009.5280802.

[28] P. Vaananen, N. Mikkola, and P. Helio, "VCO design with on-chip calibration system", *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 53, no. 10, pp. 2157–2166, Oct. 2006, ISSN: 1549-8328. DOI: 10.1109/TCSI.2006.883844.

[29] M. Youssef, A. Zolfaghari, H. Darabi, and A. Abidi, "A low-power wideband polar transmitter for 3G applications", in *2011 IEEE International Solid-State Circuits Conference*, Feb. 2011, pp. 378–380. DOI: 10.1109/ISSCC.2011.5746362.

[30] R. B. Staszewski, K. Muhammad, D. Leipold, C.-M. Hung, Y.-C. Ho, J. L. Wallberg, C. Fernando, K. Maggio, R. Staszewski, T. Jung, J. Koh, S. John, I. Y. Deng, V. Sarda, O. Moreira-Tamayo, V. Mayega, R. Katz, O. Friedman, O. E. Eliezer, E. de-Obaldia, and P. T. Balsara, "All-digital TX frequency synthesizer and discrete-time receiver for Bluetooth radio in 130-nm CMOS", *IEEE Journal of Solid-State Circuits*, vol. 39, no. 12, pp. 2278–2291, Dec. 2004, ISSN: 0018-9200. DOI: 10.1109/JSSC.2004.836345.

[31] R. B. Staszewski, J. L. Wallberg, S. Rezeq, C.-M. Hung, O. E. Eliezer, S. K. 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 transmitter for mobile phones", *IEEE Journal of Solid-State Circuits*, vol. 40, no. 12, pp. 2469–2482, Dec. 2005, ISSN: 0018-9200. DOI: 10.1109/JSSC.2005.857417.

[32] R. B. Staszewski, J. Wallberg, C.-M. Hung, G. Feygin, M. Entezari, and D. Leipold, "LMS-based calibration of an RF digitally controlled oscillator for mobile phones", *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 53, no. 3, pp. 225–229, Mar. 2006, ISSN: 1549-7747. DOI: 10.1109/TCSII.2005.858750.

[33] G. Marzin, S. Levantino, C. Samori, and A. L. Lacaita, "A 20 Mb/s phase modulator based on a 3.6 GHz digital PLL with -36 dB EVM at 5 mW power", *IEEE Journal of Solid-State Circuits*, vol. 47, no. 12, pp. 2974–2988, Dec. 2012, ISSN: 0018-9200. DOI: 10.1109/JSSC.2012.2217854.

[34] C. M. Hsu, M. Z. Straayer, and M. H. Perrott, "A low-noise wide-BW 3.6-GHz digital ΔΣ fractional-N frequency synthesizer with a noise-shaping time-to-digital converter and quantization noise cancellation", *IEEE Journal of Solid-State Circuits*, vol. 43, no. 12, pp. 2776–2786, Dec. 2008, ISSN: 0018-9200. DOI: 10.1109/JSSC.2008.2005704.

[35] Y. H. Liu and T. H. Lin, "A wideband PLL-based G/FSK transmitter in 0.18 μm CMOS", *IEEE Journal of Solid-State Circuits*, vol. 44, no. 9, pp. 2452–2462, Sep. 2009, ISSN: 0018-9200. DOI: 10.1109/JSSC.2009.2022994.

[36] P. Andreani, A. Bonfanti, L. Romano, and C. Samori, "Analysis and design of a 1.8-GHz CMOS LC quadrature VCO", *IEEE Journal of Solid-State Circuits*, vol. 37, no. 12, pp. 1737–1747, Dec. 2002, ISSN: 0018-9200. DOI: 10.1109/JSSC.2002.804352.

[37] H. Mair and L. Xiu, "An architecture of high-performance frequency and phase synthesis", *IEEE Journal of Solid-State Circuits*, vol. 35, no. 6, pp. 835–846, Jun. 2000, ISSN: 0018-9200. DOI: 10.1109/4.845187.

[38] A. A. Abidi, "Phase noise and jitter in CMOS ring oscillators", *IEEE Journal of Solid-State Circuits*, vol. 41, no. 8, pp. 1803–1816, Aug. 2006, ISSN: 0018-9200. DOI: 10.1109/JSSC.2006.876206.

[39] P. E. Su and S. Pamarti, "A 2.4 GHz wideband open-loop GFSK transmitter with phase quantization noise cancellation", *IEEE Journal of Solid-State Circuits*, vol. 46, no. 3, pp. 615–626, Mar. 2011, ISSN: 0018-9200. DOI: 10.1109/JSSC.2010.2099450.

[40] S. Zheng and H. C. Luong, "A WCDMA/WLAN digital polar transmitter with low-noise ADPLL, wide-band PM/AM modulator and linearized PA in 65nm CMOS", in *European Solid State Circuits Conference , ESSCIRC 2014 - 40th*, Sep. 2014, pp. 375–378. DOI: 10.1109/ESSCIRC.2014.6942100.

[41] B. T. Thiel and R. Negra, "Digitally controlled 3GHz 360° phase modulator in an 1.2V 65nm CMOS technology", in *2011 German Microwave Conference*, Mar. 2011, pp. 1–4.

[42] A. Ravi, P. Madoglio, H. Xu, K. Chandrashekar, M. Verhelst, S. Pellerano, L. Cuellar, M. Aguirre-Hernandez, M. Sajadieh, J. E. Zarate-Roldan, O. Bochobza-Degani, H. Lakdawala, and Y. Palaskas, "A 2.4-GHz 20-40-MHz channel WLAN digital outphasing transmitter utilizing a delay-based wideband phase modulator in 32-nm CMOS", *IEEE Journal of Solid-State Circuits*, vol. 47, no. 12, pp. 3184–3196, Dec. 2012, ISSN: 0018-9200. DOI: 10.1109/JSSC.2012.2216671.

[43] A. Kavousian, D. K. Su, and B. A. Wooley, "A digitally modulated polar CMOS PA with 20MHz signal BW", in *IEEE International Solid-State Circuits Conference. Digest of Technical Papers*, Feb. 2007, pp. 78–588. DOI: 10.1109/ISSCC.2007.373596.

[44] L. Ye, J. Chen, L. Kong, E. Alon, and A. M. Niknejad, "Design considerations for a direct digitally modulated WLAN transmitter with integrated phase path and dynamic impedance modulation", *IEEE Journal of Solid-State Circuits*, vol. 48, no. 12, pp. 3160–3177, Dec. 2013, ISSN: 0018-9200. DOI: 10.1109/JSSC.2013.2281142.

[45] M. S. Mehrjoo, S. Zihir, G. M. Rebeiz, and J. F. Buckwalter, "A 1.1-Gbit/s 10-GHz outphasing modulator with 23-dBm output power and 60-dB dynamic range in 45-nm CMOS SOI", *IEEE Transactions on Microwave Theory and Techniques*, vol. 63, no. 7, pp. 2289–2300, Jul. 2015, ISSN: 0018-9480. DOI: 10.1109/TMTT.2015.2435005.

[46] J. E. Volder, "The CORDIC trigonometric computing technique", *IRE Transactions on Electronic Computers*, vol. EC-8, no. 3, pp. 330–334, Sep. 1959, ISSN: 0367-9950. DOI: 10.1109/TEC.1959.5222693.

[47] S. Luschas, R. Schreier, and H.-S. Lee, "Radio frequency digital-to-analog converter", *IEEE Journal of Solid-State Circuits*, vol. 39, no. 9, pp. 1462–1467, Sep. 2004, ISSN: 0018-9200. DOI: 10.1109/JSSC.2004.829377.

[48] M. J. Pelgrom, *Analog-to-digital conversion*. Springer, 2010.

[49] M. S. Alavi, "All-digital I/Q RF DAC", PhD thesis, Delft University of Technology, 2014.

[50] Z. Hu, L. C. N. de Vreede, M. S. Alavi, D. A. Calvillo-Cortes, R. B. Staszewski, and S. He, "A 5.9 GHz RFDAC-based outphasing power amplifier in 40-nm CMOS with 49.2% efficiency and 22.2 dBm power", in *2016 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, May 2016, pp. 206–209. DOI: 10.1109/RFIC.2016.7508287.

[51] H. Jin, D. Kim, S. Jin, H. Lee, K. Moon, H. Kim, and B. Kim, "Efficient digital quadrature transmitter based on IQ cell sharing", in *2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers*, Feb. 2015, pp. 1–3. DOI: 10.1109/ISSCC.2015.7062979.

[52] J. A. Weldon, J. C. Rudell, L. Lin, R. S. Narayanaswami, M. Otsuka, S. Dedieu, L. Tee, K.-C. Tsai, C.-W. Lee, and P. R. Gray, "A 1.75 GHz highly-integrated narrow-band CMOS transmitter with harmonic-rejection mixers", in *2001 IEEE International Solid-State Circuits Conference. Digest of Technical Papers.*, Feb. 2001, pp. 160–161. DOI: 10.1109/ISSCC.2001.912586.

[53] C.-H. Lin and K. Bult, "A 10-b, 500-MSample/s CMOS DAC in 0.6 mm^2", *IEEE Journal of Solid-State Circuits*, vol. 33, no. 12, pp. 1948–1958, Dec. 1998, ISSN: 0018-9200. DOI: 10.1109/4.735535.

[54] M. J. M. Pelgrom and A. C. J. Duinmaijer, "Matching properties of MOS transistors", in *Solid-State Circuits Conference, 1988. ESSCIRC '88. Fourteenth European*, Sep. 1988, pp. 327–330.

[55] T. Miki, Y. Nakamura, M. Nakaya, S. Asai, Y. Akasaka, and Y. Horiba, "An 80-MHz 8-bit CMOS D/A converter", *IEEE Journal of Solid-State Circuits*, vol. 21, no. 6, pp. 983–988, Dec. 1986, ISSN: 0018-9200. DOI: 10.1109/JSSC.1986.1052639.

[56] B. Razavi, *Design of integrated circuits for optical communications*. John Wiley & Sons, 2012.

[57] M. Ingels, Y. Furuta, X. Zhang, S. Cha, and J. Craninckx, "A multiband 40nm CMOS LTE SAW-less modulator with -60dBc C-IM3", in *2013 IEEE International Solid-State Circuits Conference Digest of Technical Papers*, Feb. 2013, pp. 338–339. DOI: 10.1109/ISSCC.2013.6487760.

[58] Y. Furuta and M. M. A. Ingels, *Circuit for baseband harmonic rejection*, US Patent App. 14/108,745, Dec. 2013.

[59] G. F. Franklin, J. D. Powell, A. Emami-Naeini, and J. D. Powell, *Feedback control of dynamic systems*, 6th ed. Prentice Hall New Jersey, 2010.

[60] J. Zhuang, K. Waheed, and R. B. Staszewski, "A technique to reduce phase/frequency modulation bandwidth in a polar RF transmitter", *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 57, no. 8, pp. 2196–2207, 2010.

List of Figures

List of Tables

List of Acronyms

TX transmitter 1

VCO voltage-controlled oscillator 11

ZOH zero-order hold 18