

Millimeter-Wave Circuit Design for Radar Transceivers

Paul Swirhun

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2013-192

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-192.html>

December 1, 2013



Copyright © 2013, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Millimeter-Wave Circuit Design for Radar Transceivers

by Paul Swirhun

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science, Plan II.**

Approval for the Report and Comprehensive Examination:

Committee:

Professor Ali Niknejad
Research Advisor

(Date)

* * * * *

Professor Elad Alon
Second Reader

(Date)

Copyright 2013, Paul Swirhun

All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

I would like to thank the BWRC faculty and sponsors for their support, and Nokia Research Center Berkeley for funding the design and fabrication of the millimeter-wave integrated circuit described herein. I would also like to thank Andrew Townley for his extensive collaboration in design, measurement, and learning throughout this radar project.

Contents

CHAPTER 1 INTRODUCTION.....	6
MILLIMETER-WAVE INTEGRATED CIRCUITS.....	6
FMCW RADAR	9
PHASED ARRAYS	14
ARRAY FACTOR SIMULATION	19
CHAPTER 2 RECEIVER DESIGN	26
THE VALUE OF SIMULATORS	26
LNA DESIGN AND METHODOLOGY	28
MEASUREMENT RESULTS AND ANALYSIS.....	30
MIXER DESIGN.....	36
CHAPTER 3 ASYMMETRIC TRANSFORMER BALUN DESIGN	38
BALUNS AND THEIR APPLICATIONS	38
INTEGRATED TRANSFORMER BALUNS.....	40
ASYMMETRIC TRANSFORMER BALUN.....	44
SIMULATION AND OPTIMIZATION.....	48
CHAPTER 4 RADAR TRANSCEIVER FRONT-END	53
RADAR UNIT CELL DESIGN	53
TRANSMITTER LEAKAGE ANALYSIS.....	59
FABRICATED CHIP	64
CHAPTER 5 HIGH FREQUENCY MODELING	68
GETTING SIGNALS OFF THE CHIP	68
PORT EXCITATIONS IN ELECTROMAGNETIC SIMULATORS.....	69
MEASURED PASSIVE STRUCTURE	73
CHAPTER 6 CONCLUSION	77
PROPOSED CHANGES	77
SUMMARY.....	81
BIBLIOGRAPHY	83

List of Figures

FIGURE 1-1: RANGE AND VELOCITY MEASUREMENT USING FMCW RADAR [7]	11
FIGURE 1-2: CONCEPTUAL BLOCK DIAGRAM FOR IMPLEMENTED FMCW RADAR	14
FIGURE 1-3: FLIP-CHIP-ON-BOARD ASSEMBLY WITH PLANAR ANTENNA.....	15
FIGURE 1-4: LINEAR ARRAY OF LINEARLY-PHASED DIPOLES ORIENTED ALONG THE Z-AXIS.....	16
FIGURE 1-5: DUAL-POLARIZATION PATCH ANTENNA WITH SIMULATED RADIATION PATTERN.....	19
FIGURE 1-6: EXAMPLE ARRAY PATTERN FOR 8 ELEMENTS.....	20
FIGURE 1-7: 3-BIT, CARTESIAN I/Q GRID OF ARRAY PHASES AND CORRESPONDING AZIMUTHAL PATTERN	21
FIGURE 1-8: NEAREST-CHOICE SELECTION FROM CARTESIAN GRID AND CORRESPONDING AZIMUTHAL PATTERN.	21
FIGURE 1-9: 3-BIT 8-ELEMENT CARTESIAN BEAM-STEERING ARRAY SUMMARY	22
FIGURE 1-10: 4-BIT 32-ELEMENT CARTESIAN BEAM-STEERING ARRAY SUMMARY	23
FIGURE 1-11: 2-BIT 8-ANTENNA CARTESIAN BEAM-STEERING ARRAY SUMMARY	24
FIGURE 2-1: EM MODELING TO THE DEVICE LEVEL INTERCONNECT (LEFT) AND TO THE BLOCK-LEVEL (RIGHT)	29
FIGURE 2-2: LNA SCHEMATIC OF TRANSCIEVER UNIT CELL.....	30
FIGURE 2-3: LNA DESIGN MEASUREMENT VERSUS NOMINAL SIMULATION.....	33
FIGURE 2-4: LNA MEASUREMENT VERSUS SIMULATION WITH 35% MIM CAPACITOR REDUCTION.....	34
FIGURE 2-5: LNA NOISE FIGURE MEASUREMENT AND SIMULATIONS.....	35
FIGURE 2-6: MIXER SCHEMATIC.....	37
FIGURE 3-1: SCHEMATIC OF IDEAL TRANSFORMER BALUN WITH PORT CONNECTIONS.....	41
FIGURE 3-2: DIAGRAM OF SINGLE-TURN ASYMMETRIC CONCENTRIC SINGLE TURN TRANSFORMER BALUN	45
FIGURE 3-3: ASYMMETRIC TRANSFORMER BALUN WITH ARBITRARY OFFSET ANGLE	47
FIGURE 3-4: FOUR ROTATIONAL AND CONCENTRIC VARIATIONS OF TRANSFORMER BALUNS	48
FIGURE 3-5: SCHEMATIC OF OPTIMIZATION TEST-BENCH FOR TRANSFORMER BALUNS	49
FIGURE 3-6: TABLE OF OPTIMIZATION RESULTS FOR FOUR TRANSFORMER BALUN STRUCTURES	51
FIGURE 4-1: SYSTEM LEVEL DIAGRAM OF SINGLE-ANTENNA RADAR TRANSCIEVER.....	54
FIGURE 4-2: IMPLEMENTED RADAR FRONT-END SHOWING CHIP BOUNDARY	58
FIGURE 4-3: IDEALIZED SIGNAL PATH SHOWING PERFECT SYMMETRY	60
FIGURE 4-4: ACTUAL SIGNAL PATH SHOWING FAILURE TO MAINTAIN SYMMETRY.....	60
FIGURE 4-5: UPPER ANTENNA INSERTION LOSS TO LNA	61
FIGURE 4-6: UPPER PA INSERTION LOSS TO ANTENNA	62
FIGURE 4-7: TRANSMITTER LEAKAGE ANALYSIS UNDER LO PORT IMPEDANCE MISMATCH.....	63
FIGURE 4-8: FABRICATED 3X2.5MM CHIP SHOWING TEST-COMPONENTS TRANSCIEVER UNIT-CELL	64
FIGURE 4-9: WIRE-BONDED CHIP-ON-BOARD USED FOR PROBING.....	65
FIGURE 4-10: IF WAVEFORMS VERSUS CONDUCTING TARGET DISTANCE	67
FIGURE 4-11: EXAMPLE IF WAVEFORMS WITH TX LEAKAGE ASSUMING IDEAL CIRCUITS	67
FIGURE 5-1: EXCITATION OF ON-CHIP GSG WITH LATERAL LUMPED PORTS.....	70
FIGURE 5-2: EXCITATION OF ON-CHIP GSG WITH VERTICAL LUMPED PORT	71
FIGURE 5-3: SIMULATED GSG STRUCTURES WITH ORIGINAL AND CORRECTED MODELING METHODS	72
FIGURE 5-4: COMPARISON OF GSG INPUT IMPEDANCE WITH ORIGINAL AND CORRECTED MODELING METHODS ...	73
FIGURE 5-5: COMPARISON OF SIMULATED AND MEASURED THRU-LINE GSG STRUCTURES.....	74
FIGURE 5-6: REDESIGNED GSG WITH CAPACITIVE STUB TUNING	76
FIGURE 5-7: SIMULATION OF REDESIGNED GSG STRUCTURE.....	76
FIGURE 6-1: PROPOSED REVISED UNIT CELL ARCHITECTURE.....	79

Chapter 1 Introduction

Millimeter-Wave Integrated Circuits

Due to advancements in transistor technology, silicon integrated circuits are pushing the boundaries to bring high-frequency electronics to a larger market at lower cost. In particular, the high cutoff frequency of today's advanced CMOS and SiGe bipolar devices are enabling communication, radar, and imaging circuits operating at millimeter-wave frequencies (30-300GHz) to be successfully integrated in commercial silicon processes [1–4]. It seems likely that the commercial success of wireless communications below 6GHz will cause a large portion of future millimeter-wave circuits to be communication transceivers. However, radar techniques invented and developed for military applications are finding their way into consumer products as well. They will be increasingly used in automobile collision avoidance systems and high-speed sensor systems that may one day replace human drivers. To that end, a large number of millimeter-wave radars (and sub-circuits to be used in radars) have been published that operate in the 77GHz automotive radar band [3], [5–7]. Other applications of millimeter-wave circuits already include security screening and may soon extend to medical imaging at a large scale. There are bands allocated for other (non-automotive) forms of millimeter-wave sensing at 94GHz and 120GHz; work at these frequencies can approach and exceed $f_T/2$ of even the most advanced silicon-germanium bipolar transistors that are commercially available today [8],

[9]. A variety of design challenges at this frequency still exist, not least of which is the lack of a unified design environment for the circuit designer to use. The tool flow involves a combination of electromagnetic simulation, RF design and matching network design, and traditional circuit schematic and layout tools.

In addition to the segmentation among several design tools, a major challenge for millimeter-wave circuit design is that device and circuit modeling have more uncertainty at millimeter-wave frequencies than they do at low frequencies. This is in part due to the fact that millimeter-wave circuits are still niche-products compared to digital CMOS circuits, and the process characterization and model development at frequencies approaching 110GHz are less widely supported (and less easily measured) than the process and model development that supports high-volume CMOS processes. The sensitivity of designs to device parameters and parasitics is also a challenge in very high frequency design. Consider for example the sensitivity of a resonant LC tank to its capacitance value $S_C^{\omega_0}$ (or conversely to its inductance):

$$\omega_0 = \frac{1}{\sqrt{LC}}, \quad \frac{\partial \omega_0}{\partial C} = \frac{-1}{2(LC)^{\frac{3}{2}}} = \frac{-\omega_0^3}{2}$$

$$S_C^{\omega_0} = \frac{C}{\omega_0} \frac{\partial \omega_0}{\partial C} = -\frac{C\omega_0^2}{2} = -\frac{1}{2L}$$

This shows that the sensitivity of the tank's resonant frequency to capacitance increases as inductance decreases. In many cases, operating a circuit at a higher frequency will be achieved by reducing both inductance *and* capacitance. This is because very large devices (whose parasitics constitute the capacitor in many cases), as well as very large inductors, can have too much parasitic capacitance to be used at millimeter-wave frequencies. For example inductors may self-resonate and active devices with a lot of wiring capacitance may have inadequate f_T to be useful.

To see why capacitance typically also decreases when designing for higher frequencies, consider the bandwidth of a parallel RLC circuit, which can be derived as $BW = 1/RC$. This means that if capacitance is not also reduced when moving to higher frequencies, then the fractional bandwidth will decrease, or equivalently the quality-factor of the resonant circuit will increase. This may negate some of the benefits of moving to millimeter-wave frequencies, which are attractive among other reasons because of large fractional bandwidth availability contiguously. Reducing the inductance and capacitance of resonant circuits is potentially a problem since the sensitivity of the center-frequency to component and parasitic values is increased at the same time: a high-quality-factor resonance means the penalty for misaligning the center frequency is greater, and the gain reduces more quickly as you deviate from the center frequency. Because of this, millimeter-wave designs can suffer from higher design- and parasitic-parameter sensitivities than do lower frequency designs. This is in addition to other forms of uncertainty in the design process, including modeling uncertainty that arises because reliable measurement is more difficult at millimeter-wave frequencies.

The difficulties in designing and modeling at millimeter-wave frequencies place requirements on the design methodology and level of detail that are required of the designer. As frequency increases, inductance becomes increasingly important. Similarly as device and circuit dimensions decrease, accurately capturing the electrical interactions between metal traces will increasingly require 3-dimensional field solvers as opposed to empirical or “rule-based” parasitic extractors that have pervaded low frequency designs for decades. The shrinking dimensions of passive devices such as inductors, transmission line matching circuits, and transformers imply that even the parasitic electrical behavior of short wires, underpasses, vias, and other small layout features can have dramatic effects on the overall circuit performance—even at the micron-scale and device-level. Uncertainty in how to account for the sensitivity of larger-scale circuits to minute layout choices is one of the barriers to achieving good performance results on

the first attempt. To be conservative, millimeter-wave designs can be modeled at high levels of detail, but this can increase the design time significantly. Like other endeavors, experience helps; and this research project was as much about learning how to design at millimeter-wave frequencies (what to ignore, what to model in detail) as it was about investigating the tradeoffs and complications presented by a polarimetric transceiver architecture.

FMCW Radar

The basic elements of a radar transceiver are a transmitted signal with some modulation, reception of a reflected signal that has bounced off one or more targets, and signal processing to compare the received signal with a known (or directly observed) transmitted signal and derive a useful result, such as the object’s range, lateral position, or velocity in one or more directions. Radar modulation schemes have been devised and tested for decades and include time-domain electromagnetic pulses, frequency-domain pulses or “chirps”, and—more recently—code and phase domain chirps. In all cases, a simple equation referred to as “the radar equation” combines free-space path loss for a wave traveling round-trip to the target and back with antenna gain and radar cross section to estimate the (free-space) power ratio between transmitted and received signals [10]. This equation is shown below, and relates the received and transmitted powers P_R and P_T using the receive and transmit antenna gains G_R and G_T , the target radar cross section σ , the range R to the target, and constants.

$$\frac{P_R}{P_T} = \frac{G_T G_R \lambda^2 \sigma}{(4\pi)^3 R^4}$$

In most cases this loss is high due to the fourth-order dependence on range. While there is apparently a penalty to using higher frequency signaling, it allows miniaturization of phased arrays which can raise the effective antenna gains to compensate. In the architecture described herein, the transmit and receive antennas would physically be the same array, so $G_T = G_R$.

Finally, if azimuthal resolution is important, as it is for some applications like automotive radar, then the larger arrays afforded at higher frequencies allow denser spatial imaging, simply because a higher-gain phased antenna array can fit more non-overlapping pixels or “beam-widths” within a given field of view.

The principle of frequency-modulated continuous wave radar is that a continuous sinusoidal tone is swept in frequency around some nominal center frequency. The sweep is typically linear, and to allow the transceiver front-end to be designed only using narrow-band circuits, the bandwidth of frequency modulation is small compared to the center frequency, typically around 1% fractional bandwidth. The transmitted signal is reflected off distant targets and returns to the receiver after a round-trip time delay, indicated by the grey delayed waveform in Figure 1-1 [7]. At the receiver, a comparison is made between the frequency currently being transmitted and the frequency received; this is done by mixing the two signals and examining the low frequency content of the mixer IF output. The frequency difference Δf that appears at baseband is related to the distance R to the target by the speed of light c , the sweep bandwidth B and the modulation period T_m as shown in the equation below. Since the reflected signal must be received and compared to the transmitted signal before the frequency sweep changes direction (or stops its ramp), the round-trip delay must be less than half the modulation period. Delays longer than half the modulation period are aliased if the frequency modulation is actually triangular, or they are translated to some other frequency if the frequency modulation has finite dwell times at its maximum and minimum excursions, as is typically the case. The modulation rate is typically slow (approximately 1-10kHz) to increase the unambiguous range of detection, essentially lowering the frequency/distance-mapping constant in exchange for using less RF bandwidth, while maintaining a given unambiguous target range. Stationary targets produce a frequency difference proportional to distance, so peak discrimination on an FFT of sampled IF data allows the system to deduce target range.

$$\Delta f = \frac{2B}{T_m} \frac{2R}{c}$$

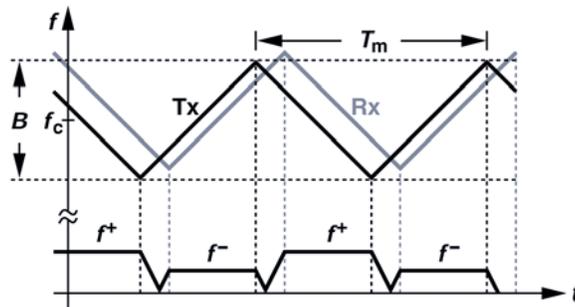


Figure 1-1: Range and velocity measurement using FMCW radar [7]

The IF waveform in FMCW radar also “glitches” or has some undesirable frequency content around the time that the ramp ceases and changes direction. To ignore this, each one-direction ramp is triggered, then IF data is acquired during only a fraction of the transmitted ramp bandwidth. In this way, the behavior around frequency-ramp-direction transitions is ignored and any erroneous target detection around that time is avoided. If the target is moving radially with respect to the radar transceiver then the frequency ramp slope incident on the target is increased or decreased, causing the frequency difference Δf at the transceiver to be slightly different during ramp-up and ramp-down parts of the triangle wave shown in Figure 1-1. This allows radial velocity information to be deduced in addition to target range.

Transmitter leakage is a big problem in radar systems of this type because the transmitted signal contains many harmonics of the modulation frequency all translated up in frequency around the carrier. If the transmitted signal leaks into the receiver (rather than just being used to drive the down-converter) then it can mix with a time-delayed version of itself and produce a large number of frequency components at the modulation frequency and its harmonics. This is illustrated below where the phase of the transmitted signal is derived in terms of a Fourier expansion of the triangle-shaped frequency modulation around the center frequency ω_0 . Here

the transmitted signal is $x(t)$ and produces a transmitter leakage component $y(t)$ that mixes with $x(t)$. Amplitude pre-factors and units are neglected since only the spectral characteristics at baseband (which are due to transmitter leakage) are of interest.

$$\phi(t) = \int_0^t \omega(t) dt = \int_0^t \omega_0 + \sum_{n=-N}^N \Omega_n e^{\frac{j2\pi n}{T_m} t} dt = \omega_0 t + \sum_{n=-N}^N \frac{\Omega_n T_M}{j2\pi n} e^{\frac{j2\pi n}{T_m} t}$$

$$x(t) \propto \cos(\phi(t))$$

$$y(t) \propto \cos(\phi(t - \Delta t))$$

$$IF(t) \propto x(t)y(t)$$

$$\begin{aligned} &\propto \cos\left(2\omega_0 t + \sum_{n=-N}^N \frac{\Omega_n T_M}{j\pi n} e^{\frac{j2\pi n}{T_m} t} e^{-\frac{j\pi n}{T_m} \Delta t} \cos\left(\frac{j\pi n}{T_m} \Delta t\right) - \omega_0 \Delta t\right) \\ &\quad + \cos\left(\sum_{n=-N}^N \frac{\Omega_n T_M}{\pi n} e^{\frac{j2\pi n}{T_m} t} e^{-\frac{j\pi n}{T_m} \Delta t} \sin\left(\frac{j\pi n}{T_m} \Delta t\right) + \omega_0 \Delta t\right) \end{aligned}$$

To simplify this expression, assume that the modulation frequency is much smaller than the carrier frequency and that the time delay Δt is small, as it will be if it is caused by an on-chip, on-board, or on-array leakage signal rather than some longer path out to a nearby target and back to the receiver. Also assume that the mixer or IF buffers filter the high frequency components around $2\omega_0$. Finally, neglect all phase terms since a zero phase reference is arbitrary. The receiver signal processing algorithm will typically consider only the FFT magnitude since the phase change along the signal path to the target and back is random. Then the simplified result is the following:

$$IF(t) \cong \cos\left(\sum_{n=-N}^N j\Omega_n \Delta t e^{\frac{j2\pi n}{T_m} (2t - \Delta t)}\right)$$

This analysis made many simplifications, including that the LNA is perfectly linear, the mixer is ideal and only down-converts at the LO fundamental, and amplitude and constant-phase

terms are neglected. Still, we see that a low frequency term appears that is a cosine of the phase accumulation that occurs in the triangle-wave modulation of the frequency. The Fourier coefficients of the phase are almost the same as those of the triangle wave, except they are scaled by phase constant $j\Omega_n$ and the phase is shifted by a constant amount that depends on the delay of the leakage path Δt . Also notice that this IF leakage signal is at twice the rate of the modulation frequency, meaning that if a full-triangle modulation period is T_m , then the leakage signal produces interferers at $f_{leak} = \frac{1}{2T_m}$ as well as harmonics of f_{leak} . Since this modulation frequency is typically on the same order as the desired IF signal that conveys target range, the two interfere and removing the leakage signal at baseband (for example by digital filtering) can be extremely difficult or impossible.

A common form of FMCW radar, and a conceptual summary of the radar architecture described in this report, is illustrated in Figure 1-2. A lower-frequency frequency synthesizer generates the ramped frequency signal described previously, which is then multiplied up in frequency using a frequency multiplier. This frequency multiplier drives a power amplifier which connects to an antenna through an isolating structure of some kind that acts as a circulator to direct signals returning to the antenna from distant targets towards the low noise amplifier. Then the amplified and received signal is down-converted using the transmit signal to produce range and velocity information of targets. A transmit phase shifter is included to steer the transmit beam in a phased array, which would include an array of many unit elements described by Figure 1-2. Although many existing automotive radar systems do not have steerable transmit beams, it may become necessary in the future when multiple radar systems will be operating in the same environment (such as on a road) and need to discriminate their own signal from the interfering transmit signals of other nearby radars [11]. Receive-side beam-forming can also be accomplished if the mixer is made of two in-phase and quadrature-phase mixers with phase shifting at RF, IF, or LO. The diagram shows that the LO signal driving the mixer originates

before the transmit phase shifter and before the power amplifier. In practice, it can be extracted anywhere along the transmit path, so long as the LO signal also experiences the frequency ramp; this is a fundamental feature of FMCW radar as seen in Figure 1-1. In particular, the architecture implemented herein extracts the LO signal from the output of the power amplifier using a waveguide coupler. For an extension of the conceptual block diagram of Figure 1-2 to the actual transceiver architecture, refer to Chapter 4.

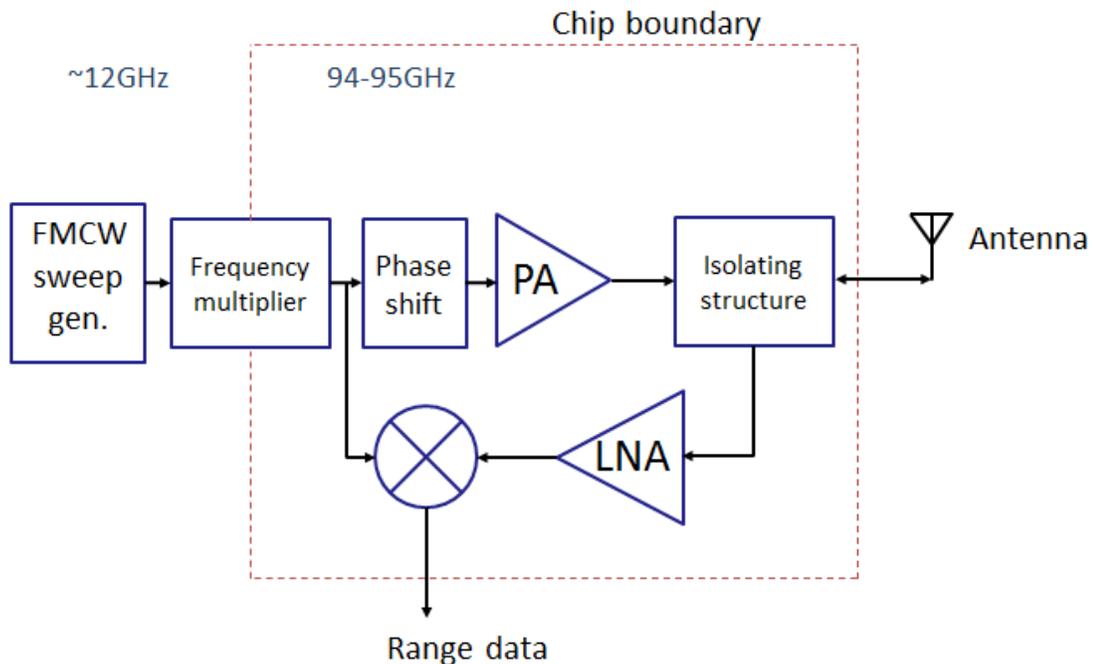


Figure 1-2: Conceptual block diagram for implemented FMCW radar

Phased Arrays

Phased arrays allow a radar or communication system to spatially localize transmitted power in one or more directions, as well as to discriminate received power that originated from one or more directions against that received from other directions. In most cases, phased arrays are designed to preferentially transmit and receive in one direction at a time, and in a way that the direction selected can be controlled electrically without modification of the positioning of the

constituent antenna elements. Perhaps the simplest phased array realization for millimeter-wave integrated circuits would be to dedicate one radar transceiver per antenna and tile sets of transceiver/antenna pairs in a grid pattern, each with a form of phase control and fed by a common reference signal. This could be achieved by flipping silicon dies onto a laminate substrate with planar antennas made of a high quality metal layer on that substrate, as shown in Figure 1-3. Alternatively, the dies could be flipped onto the back of the substrate and high-frequency signals can be coupled to the antennas by vias or apertures in the board's metal layers.

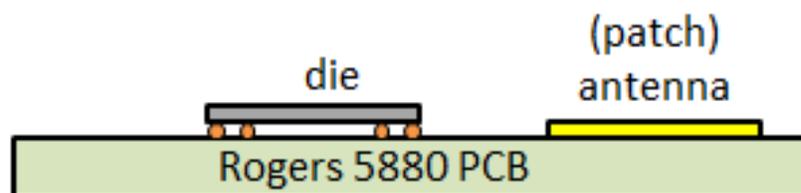


Figure 1-3: Flip-chip-on-board assembly with planar antenna

The advantages of this architecture are that the array is uniform, and each element is composed of identical dies, all with similar characteristics. The coupling to the antenna is also simple: an on-board impedance match is made to the chip and an on-board transmission line connects to the antenna. Alternative structures use multiplexers to address several antenna elements with fewer transceivers, but so far arrays of this type implemented in silicon suffer from the high loss associated with silicon series switches operating around 100GHz [2]. Using an active circuit individually assigned to each antenna allows reception, demodulation, amplification, and processing to be done all on one chip, and allows a low-frequency output to carry the radar information to a digital signal processor rather than requiring many high frequency signals to be connected to a multiplexer operating at millimeter-wave frequencies [12]. The disadvantage is clearly that more dies are needed to implement an array, and that the

power consumption of the front-end grows linearly with the number of elements if all are active simultaneously. This growth can be sub-linear for a given target range since transmit power can also be reduced per element. However receiver power consumption still scales linearly with the number of receivers because the receiver is always designed for the best achievable sensitivity.

The analysis of phased arrays is thoroughly described in introductory electromagnetics textbooks; the mathematics originates from the superposition in space of electric and magnetic fields produced by two or more antenna elements. Similar to the geometric calculations for double-slit diffraction patterns, it can be derived using simple trigonometry that linearly spaced antenna elements as shown in Figure 1-4, phased with uniform phase offset $\Delta\alpha$ will constructively interfere and produce a plane wave at an angle β (illustrated) from the y-axis of Figure 1-4 such that the following relationship is true:

$$\beta = \frac{\pi}{2} - \cos^{-1}\left(\frac{c\Delta\alpha}{\omega d}\right) = \frac{\pi}{2} - \cos^{-1}\left(\frac{\Delta\alpha}{\pi}\right) \text{ if } d = \frac{\lambda}{2}$$

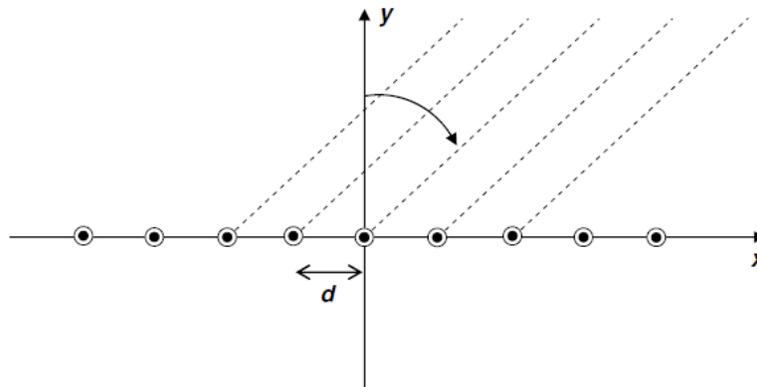


Figure 1-4: Linear array of linearly-phased dipoles oriented along the z-axis

Superposition of electric fields from many elements in an array assumes that the electrical excitation of each antenna element is unaffected—or affected in a known way—by the electromagnetic fields of other elements of the array. Under such an assumption, the far-field

electric field of the array can be computed as the product of a 3-dimensional electric field vector function $\vec{E}(r, \theta, \phi)$ describing a single radiating antenna, with a sum of phasor terms (complex exponentials) which encapsulate the positions of all elements of the array as well as their excited magnitude and phase. The first term is typically called the *element factor* since it describes the electric field of a single antenna element of the array. The latter term—denoted $F(r, \theta, \phi)$ —is a 3-dimensional scalar function describing the amount that electric fields from all elements interfere constructively or destructively at each position in space; it is called the *array factor*. The far-away electric field for an array of n-elements can be computed to be:

$$\vec{E}_{ff}(\vec{r}) = \vec{E}(r, \theta, \phi) \sum_{k=0}^{N-1} \frac{\tilde{I}_k}{\tilde{I}_0} e^{jk\hat{r}\vec{h}_k}$$

In the above equation each antenna element at position \vec{h}_k has current phasor \tilde{I}_k ; and emits a wave with wavenumber $k = 2\pi/\lambda$ in any given direction $\hat{r} = (\sin \theta \cos \phi \hat{x} + \sin \theta \sin \phi \hat{y} + \cos \theta \hat{z})$, which is the unit-length vector pointing to the evaluated point (r, θ, ϕ) . In the case of the simplest type of radiating element: a Hertzian length d dipole (meaning $d \ll \lambda$) oriented along the \hat{z} axis, the element factor is given below for element $k = 0$. This element factor could be replaced with the electric vector field of a patch antenna, for example, or any other single antenna element being used in the phased array design. The radiation pattern of a dual-polarization patch antenna simulated on a commercial substrate is shown in Figure 1-5, overlain onto the physical layout. For the Hertzian dipole, however, the pattern varies only in θ , and has only $\vec{\theta}$ directed electric field components.

$$\vec{E}_{Hertzian}(r, \theta, \phi) = \hat{\theta} \frac{j\eta_0 k I_0 d}{4\pi r} \sin \theta e^{-jkr}$$

From the equations above and an input excitation power, the gain of an antenna array can be computed using the Poynting vector (which is available to electromagnetic field solvers since

it is just the cross product of electric and magnetic vectors at each point in space). The pattern can also be computed, which shows the shape in space but not the absolute magnitude of the gain. The pattern is defined as the gain normalized by the peak gain, and therefore divides away all pre-factors inherent to the gain calculation including those dependent on radius (the distance from the array). This eliminates the exponential dependence of e^{-jkr} terms, as well as all radial dependence.

$$G(\theta, \phi) = \frac{\langle \vec{S}(r, t) \rangle \cdot \vec{r}}{\frac{P_{in}}{4\pi r^2}} \quad p(\theta, \phi) = \frac{G(\theta, \phi)}{\max_{\theta, \phi} G(\theta, \phi)}$$

In the simplest and most common phased arrays, antenna elements are all identical, spaced regularly along a one- or two-dimensional grid, and use “end-fire” phase coordination, meaning that the phase difference between any two adjacent elements is the same along that entire axis of adjacency. In practice this is typically not the case, and each element has some finite phase resolution, as well as variation in phase setting as the desired phase setting is changed. To be more specific, any real phase shifter has a nonlinear relationship between desired phase and actual phase, and typically may also have non-constant amplitude across all phase settings and non-constant phase across amplitude settings. Such variations within a phased array are akin to AM-PM distortion in communication circuits.

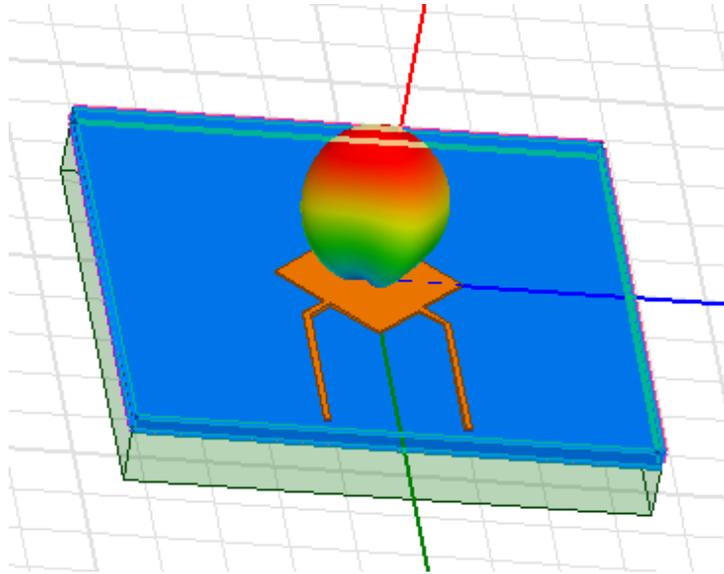


Figure 1-5: Dual-polarization patch antenna with simulated radiation pattern

Array Factor Simulation

Using the preceding formulation of the array factor and assuming a known antenna element factor (the following continues to use Hertzian dipoles) the array radiation pattern can be computed directly. In particular, an investigation into the phase and amplitude resolution required in the constituent elements of an array can help inform the design of each transceiver. Such an analysis is beneficial since it shows—at various levels of detail—the array performance that can be expected given non-ideal transceivers whose transmit amplitude or receive gain can vary across the array, as well as whose phase resolution is finite and subject to variation. The same method used to predict array factors due to quantized phase and amplitude can later be used to calculate actual array factors after, for example, transmitter phase resolution and variation are deduced from a circuit simulation.

Based on the preceding introduction to phased arrays, the array factor in the azimuthal axis is computed assuming that the phase and amplitude of an array of transmitters (or equivalently, of receivers) is quantized to a uniform Cartesian grid, as is the case if a the phase shifter is

composed by summing in-phase and quadrature-phase signals with adjustable, linearly distributed weights. The pattern for an array of Hertzian dipoles has back- and forward-pointing lobes as shown in Figure 1-6, but in most planar architectures one lobe is eliminated by a metalized reflector: for example a lower metal layer of the substrate of Figure 1-3.

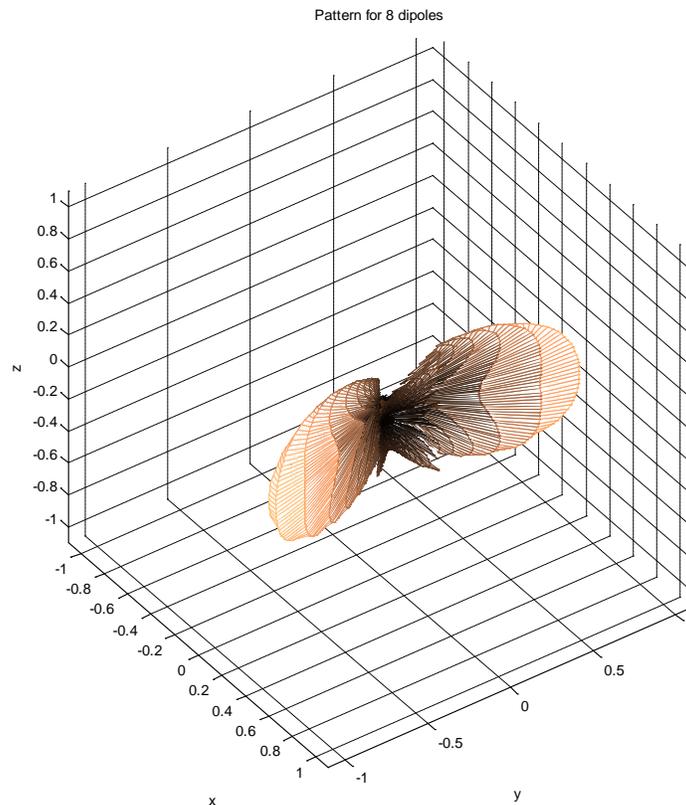


Figure 1-6: Example array pattern for 8 elements

A 3-bit sign/magnitude Cartesian grid of amplitudes in each in-phase and quadrature-phase components is shown in Figure 1-7, illustrating array steering to 60-degrees azimuthally. The azimuthal array pattern is shown at right and represents a horizontal slice of Figure 1-6. The selected beam steering direction happens to require elements in the array to alternately select a fully in-phase or fully quadrature-phase point on the Cartesian grid, as shown with blue-circles around the red grid points used for the current steering direction. As a result, no quantization error occurs.

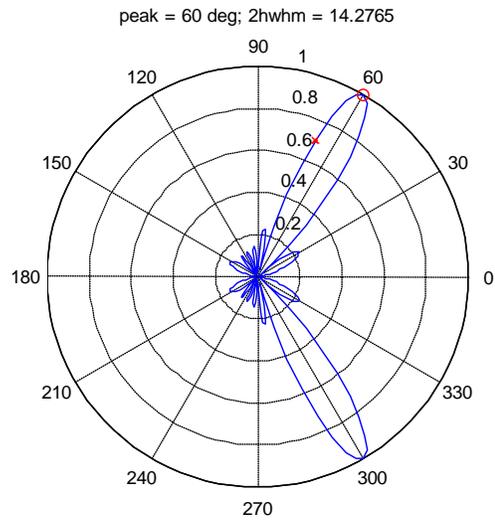
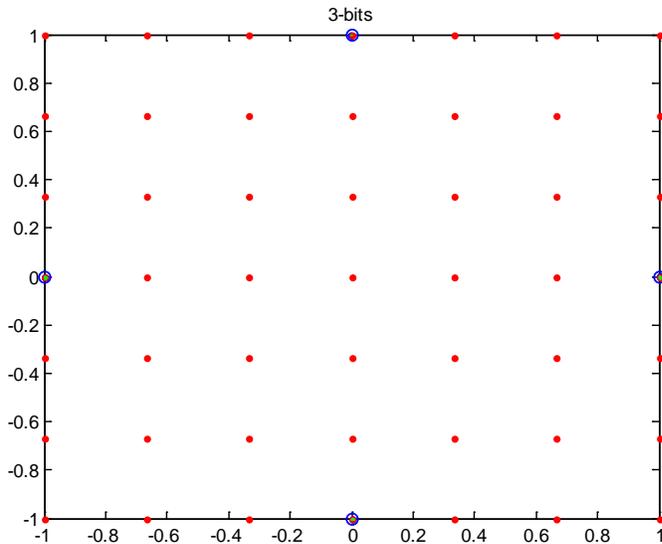


Figure 1-7: 3-bit, Cartesian I/Q grid of array phases and corresponding azimuthal pattern

In general, there will be quantization error and the phase set-points required for each element of the array (blue dots) will deviate from the available Cartesian grid set-points (red dots), as shown in Figure 1-8 and the resulting pattern at right.

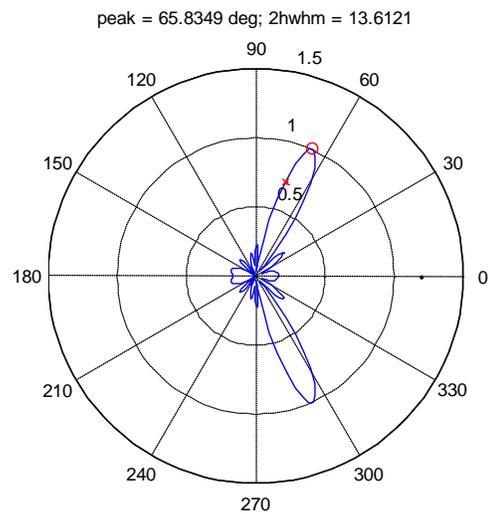
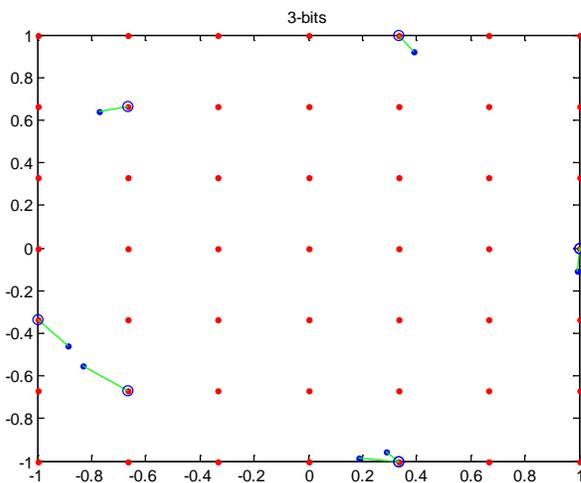


Figure 1-8: Nearest-choice selection from Cartesian grid and corresponding azimuthal pattern

Although pre-distortion can be used to compute the desired phase set-points to counteract the sinusoidal relationship between $\Delta\alpha$ and β , the array will still exhibit both magnitude and phase fluctuation as shown in Figure 1-9. Applying pre-distortion linearizes the beam direction, but 5-10% beam magnitude error is observed across a -40 to 40 degree beam direction when using 3-bit sign/magnitude quantization and 8 antenna elements. The beam magnitude error is computed as a relative array gain error compared to the gain when all elements are in-phase—which is shown by the fact that all red curves pass through (0,0) on their respective axes. The beam width also varies with steering angle and is asymmetric about the beam peak, particularly as the beam steers away from the axis normal to the array. The beam direction error is small, however, compared to the beam width. This is a feature of the number of elements in the array, as we will see through comparison later.

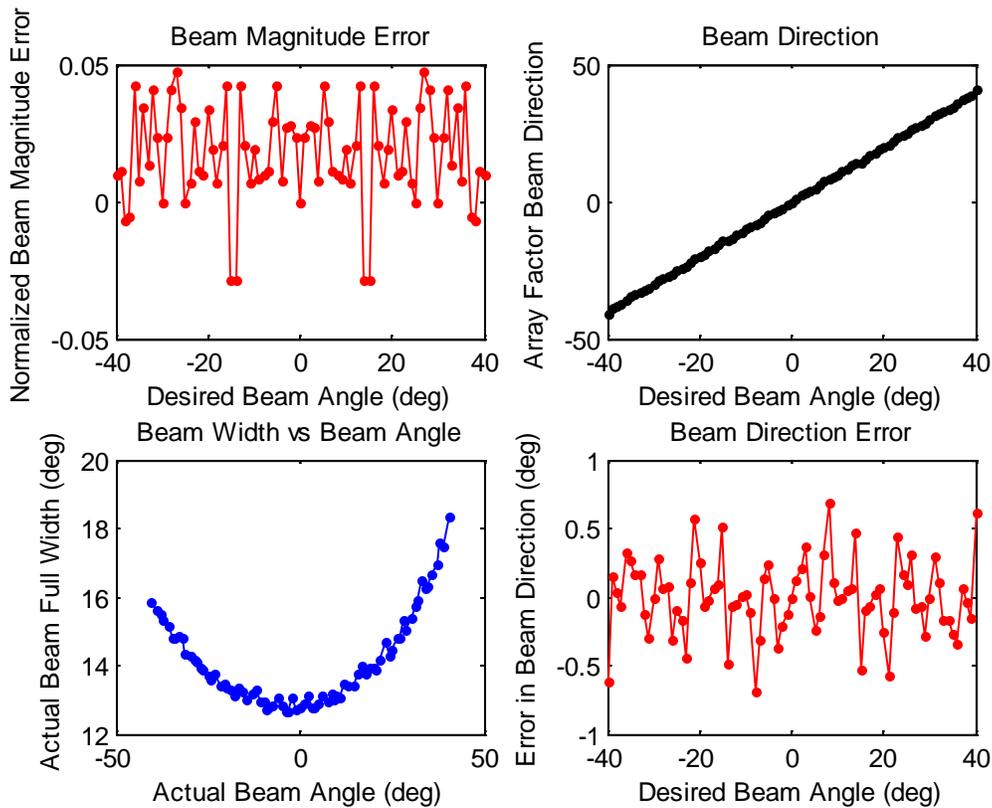


Figure 1-9: 3-bit 8-element Cartesian beam-steering array summary

Adding additional elements narrows the array beam width and also reduces the beam direction error versus desired beam direction, as shown in Figure 1-10. However, in continuous-wave phased arrays there is little use in having beam-steering resolution of ≤ 0.1 degree when the beam-width itself is an order of magnitude (or more) larger. This is because objects narrower than the beam-width will appear widened to be as large as the beam-width (akin to a spatial convolution of the object cross section with the beam pattern as the beam is scanned). The result will be an image of the antenna array pattern and not of the smaller-dimensioned object. Also note that adding more bits to the Cartesian grid improves the beam magnitude error significantly, as seen in the upper left plot of Figure 1-10.

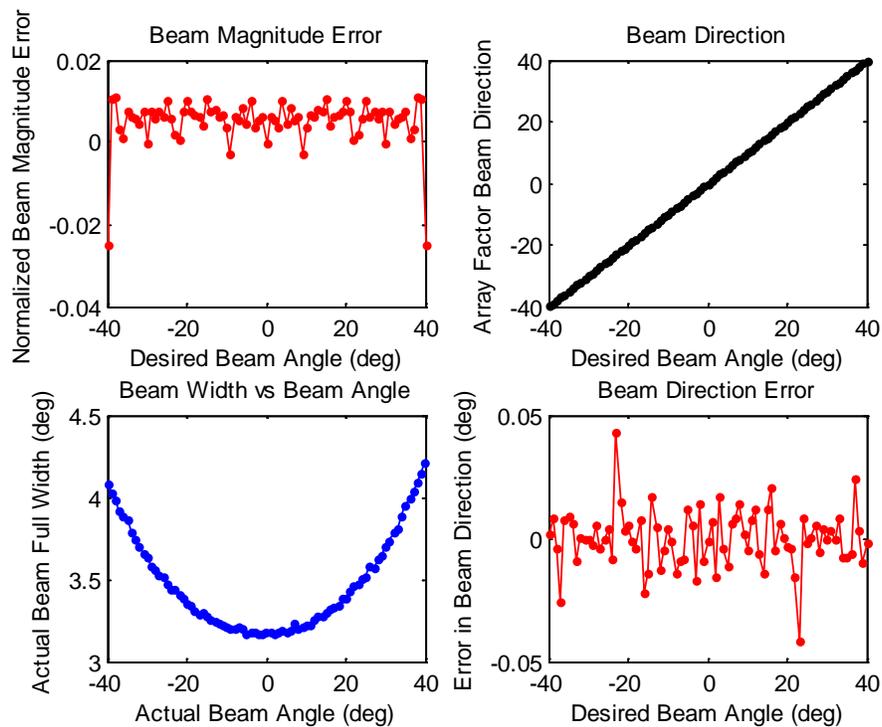


Figure 1-10: 4-bit 32-element Cartesian beam-steering array summary

From these plots it can be seen that the beam direction does not track the desired beam direction exactly; however, as long as the actual beam direction is known within some fraction of the beam-width, an azimuthal image can be constructed. As described in the previous

paragraph, it is likely that the field of view of the phased array will be segmented into angular pixels approximately equal to the beam-width, or with some overlap probably not exceeding 50% of the beam-width. For this reason, the set of beam angles that will actually be used in a phased array is only a subset of the x-coordinates in these quantization-summary plots. This is to say that in the upper left plot of Figure 1-11, we can select a subset of points with equal (normalized) beam amplitude, since extra angular resolution will not improve the azimuthal imaging capabilities of an 8-element array whose beam-width is approximately 14 degrees. This could alleviate the system from requiring different gain (or IF spectrum thresholds) at each point of the angular scan. In practice, gain variance on the order of 10-15% (0.5dB) around the normalized beam magnitude is unlikely to pose major problems for the radar system, particularly because the front-end circuits will typically have their own in-band gain variation of that order.

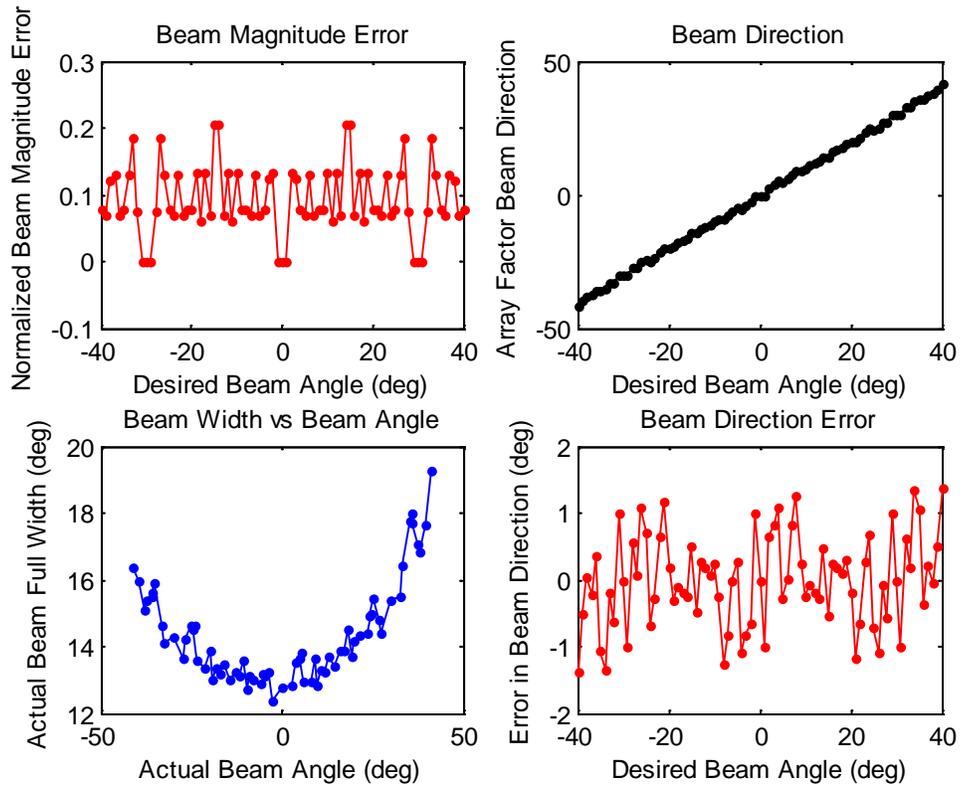


Figure 1-11: 2-bit 8-antenna Cartesian beam-steering array summary

The conclusion from this analysis—and in particular from Figure 1-11—is that if the array is moderately sized (8 elements or more), then even very coarse phase quantization between array elements is acceptable. The beam-direction error will average to zero as the array grows larger, even with coarse phase adjustment in each element. The array magnitude variation will likely not be an issue for the two reasons outlined previously. The first was that the gain variation is not a huge contributor to the overall system (transmit and receive) gain variation when considering all front-end circuits. The second was that if the array is not enormous, only a subset of desired beam-directions will actually be used, so that subset can be selected to give equal array gain; otherwise adjacent intended beam directions will actually be overlapped due to the finite beam-width of a limited-extent array.

To put the phase shifting coarseness of Figure 1-11 into perspective, each in-phase and quadrature-phase channel for phase-combining (before a PA, for example) has only 2 sign/magnitude bits, which is equivalent to saying that each channel can be on or off, and its sign can be switched. Such coarse phase resolution is much easier to achieve at a circuit level than fine-grained phase adjustment, which can be difficult to design and even more difficult to measure at millimeter-wave frequencies.

Chapter 2 Receiver Design

The Value of Simulators

The design of the low-noise amplifier (LNA) and mixers attempts to balance four requirements: low noise figure to improve the detection range of the radar receiver, a high input compression point to tolerate some transmitter leakage in this direct-coupled architecture, low power consumption to allow larger arrays within a certain power budget, and small layout area. The last of these requirements is a concern not because of the cost of silicon area but due to the non-planar architecture of the radar transceiver (described in Chapter 4), which requires that the LNA be surrounded by the RF signal path and only accessible to the chip periphery using underpasses beneath the microstrip signal path. Making the receiver smaller decreases the routing loss throughout the transceiver unit-cell and therefore improves both the system noise figure and transmitted power reaching the antenna.

LNA designs operating at W-band frequencies have been widely studied [9], [13–16] and publications frequently rely on simple linearized models to demonstrate the utility of some small architectural change, such as transformer feedback, capacitive feedback, or—most commonly—inductive emitter feedback. Algorithmic design procedures have also been developed [17], [18], although their usefulness only extends to getting an approximate first-cut design. Advanced silicon-germanium device models at millimeter-wave frequencies involve tens of parasitic

elements that can significantly impact, for example, the gain and reverse isolation of common-emitter amplifier stages, so simplified π -models or h -models can be extremely inaccurate compared to simulation results. Extracted parasitics can also significantly detune designs done in schematic or using linearized hand-calculated parameters. The result is that a nontrivial amount of iteration in the simulator- and in layout-environments is necessary, which can become a tedious task when repeated electromagnetic simulations or layout extractions are needed to tune an amplifier design. In practice, millimeter-wave designs almost always benefit from tuning in a simulator, and the designer throws away easily achievable performance improvements by settling for just schematic-simulated or hand-calculated design parameters.

To give a typical example, it is common in publications to see the well-known benefits of inductive emitter degeneration apparently calculated, demonstrating how a real input impedance appears as $\omega_T L_E$, a value that can be more easily matched to a 50Ω environment. However, the quantitative usefulness of this analysis is highly dubious around 100GHz since it neglects intrinsic emitter resistance and inductance, base resistance, and extracted parasitic capacitance networks including the base-collector capacitance C_{μ} . In practice, after deriving the theoretical concept using a simplified model, the designer would be better served by simply adjusting the emitter inductance in a simulator to achieve the desired input resistance, for example by plotting the optimum source-impedance or available-gain circles on a Smith chart. Particularly at frequencies approaching 100GHz, the models and their dependence on parasitics are complex enough that hand calculations can provide only an approximate starting point for the design.

Finally, it is important to remember that linearized models, including s-parameter simulations, are unaware of the design goals beyond gain, bandwidth, and noise figure. To handle significant transmitter leakage at the LNA input, large-signal compression simulations also need to be run; and to balance power consumption, compression point, gain, bandwidth and noise figure, the LNA designer will need to iterate between small- and large-signal

simulations, typically arriving at a non-optimal design when viewed through only the lens of one domain, such as s-parameter simulations. Algorithmic design procedures, while beneficial and often derived using realistic device simulations, may neglect to take these other concerns into account, and instead focus wholly on the lowest possible noise figure, for example.

LNA Design and Methodology

Two flows for electromagnetic simulation and parasitic extraction were used during the design and analysis of the chip. They differ in the extent to which rule-based parasitic extraction is used for active devices. Both use electromagnetic simulation for large-scale transmission-line structures after a first-draft design is completed using scalable circuit models. The two flows are illustrated in Figure 2-1, which depicts (at left) a detailed connection to the device including vertical metal stacks descending from the upper-metal layer to the lower device-connection layers, drawn in pink. Upper metal tapers are also included in the electromagnetic simulator. At right, the device-level connections and the local connections to nearby resistors and AC coupling capacitors are extracted and only the larger-scale transmission line structures are simulated in the electromagnetic simulator. The vertical connections at right are just lumped ports to excite the structure, whereas at left they are physical metal objects and the lumped ports provide short vertical excitations between a ground plane and the pink lower metal layers.

The method at left has the potential limitation that the ground plane does not actually exist directly below the lower-metal device-connections, but it is required to excite the ports off of a common node in the simulator. The method at right has the limitation that the extractor might not accurately capture inductive and capacitive effects at the larger distance-scales enclosing the entire “core” of an amplifier stage. In the case of the 130nm process used, it turns out that with even older rule-based parasitic extraction (not even using a more accurate quasi-field-solver extractor) the results of these two methods give very similar results. So unless tuning will

be done only at the core-device level, the method at right is preferred since electromagnetic simulators still require much more compute time than circuit-simulators and parasitic-extractors combined.

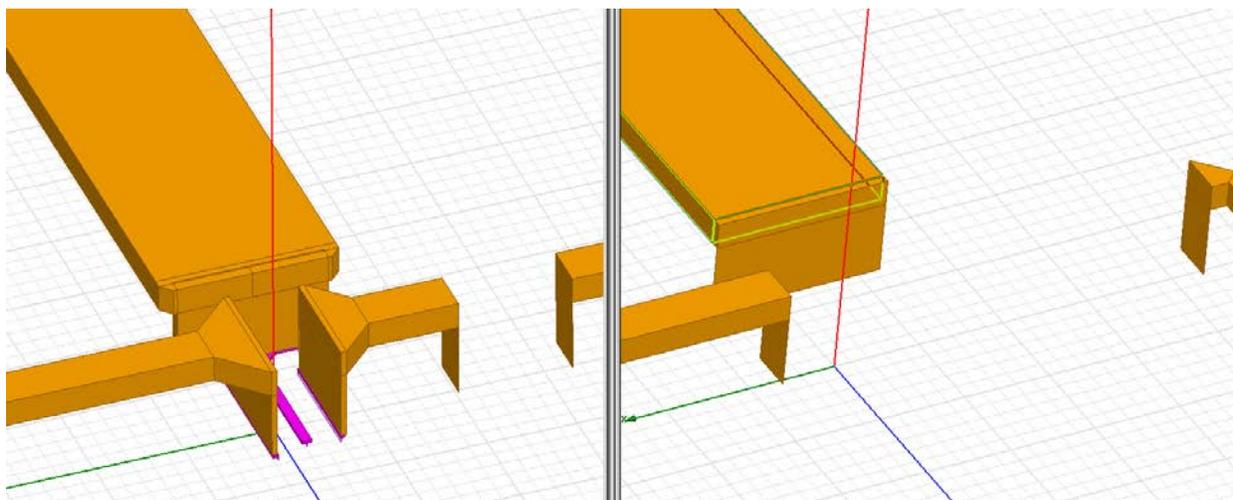


Figure 2-1: EM modeling to the device level interconnect (left) and to the block-level (right)

The LNA design implemented in the transceiver unit cell is illustrated in Figure 2-2. It uses four resistively-biased, AC-coupled common-emitter stages with varying amounts of emitter degeneration. Transmission-line based L-match circuits are used to attempt to bi-conjugately impedance match each stage, including the final stage which drives a pair of mixers in-phase with a cross-shaped microstrip splitter. The final stage was power-matched to the large-signal periodically time-varying linearized RF input impedance of the mixers (s-parameter simulation on top of a periodic steady state operating point), which is significantly lower impedance than the collector nodes of the transistors. To minimize noise injection due to the emitter degeneration in the LNA, the inductive transmission lines are made $20\mu\text{m}$ wide, while the matching networks between stages—which connect to higher impedance collector nodes—are made out of narrower $5\mu\text{m}$ wide transmission lines. The amount of emitter degeneration varies: it is high in the first stage to better align the optimum noise- and power-match input impedances

and reduce the overall LNA noise figure. The second and third stages use less emitter degeneration to achieve higher gain and they contribute the majority of the amplifier's overall gain. The fourth stage uses a medium amount of inductive emitter degeneration to reduce the gain and improve the linearity of the LNA, since the last stage experiences a much larger signal swing at its input than do the preceding stages.

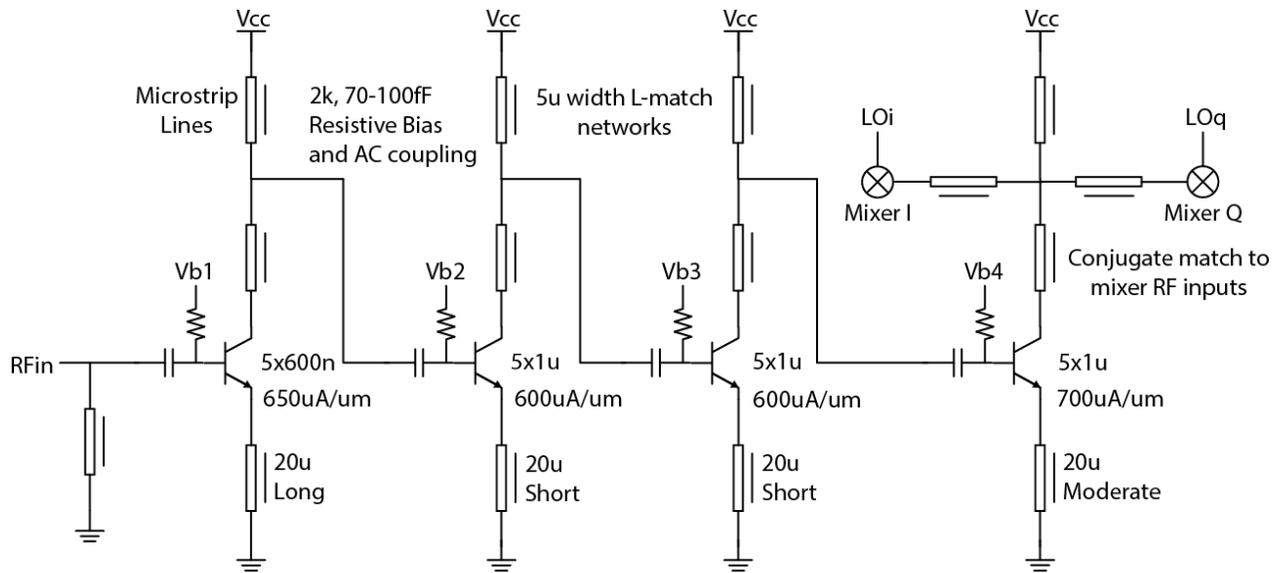


Figure 2-2: LNA Schematic of transceiver unit cell

Measurement Results and Analysis

After calibration of an Anritsu 110GHz frequency-extended vector network analyzer, the stand-alone LNA was measured to compare against simulations. Observations on the differences can help inform design adjustments and allow the designer to modify the simulation methodology and design choices used in subsequent designs. The stand-alone LNA differs from the LNA used in the radar transceiver in that its output matching stage is single ended and AC coupled with a different value capacitor. Unlike the LNA in the transceiver unit-cell, the LNA used for testing is roughly matched to 50Ω , which is the intended termination impedance of the ground-signal-ground pad structure and the microprobes used to test the chip. The LNA

implemented in the unit-cell transceiver was conjugately matched to the mixers' RF input impedances, which redesigned versions indicate is unnecessary and can be traded off to achieve flatter gain in the receiver. Chapter 5 discusses the ground-signal-ground (GSG) structure and the simulation mistakes made in designing it for operation at the 94.5GHz intended center frequency.

The measurements include both the input and output GSG structures, which each introduce around 0.5dB of loss. No de-embedding is performed, and models of the GSG structure corrected using the measurements of Chapter 5 are added schematically to the input and output of the LNA simulation to compare the raw measurement against expected results. Additionally, since the LNA in the transceiver is not accessed by pads but is instead connected to a Wilkinson combiner in an on-chip 50Ω environment, the stand-alone LNA actually has two input matching networks. This is because a GSG structure matched to 50Ω is simply prepended onto the input of the LNA used in the transceiver, which uses its own matching network to match to 50Ω . In other architectures the LNA interfaces directly with pads on the chip, so these two matching networks could be combined and the loss could be reduced. As stated previously, the fourth stage differs between the stand-alone LNA and the one used in the radar transceiver; the latter having two in-phase outputs which cannot be easily measured using only GSG microprobes.

The DC characteristics of the LNA agree well with expectations: the signal path consumes 20mW from a 1.8V supply and the biasing consumes an additional 16% or 3mW under the operating conditions used for measurements. The layout consumes roughly $500\mu m \times 450\mu m$ not including GSG pads which are only present on the stand-alone version used for testing. The simulated LNA gain is 13.8dB (or about 14.8dB with the input and output GSG structures de-embedded). The input 1dB compression point is -18dBm for the single-output LNA implemented for testing. The design achieves a very wide (>22GHz) 3dB bandwidth from around 88GHz to

over 110GHz, the limit of the VNA's measurement capabilities; this is largely due to input and output impedance matches staggered in frequency. The dual-output LNA combined with the mixers has a simulated input-referred compression point that is similar: around -20dBm, depending on the LO drive signal applied. The actual LO drive signal experienced by the mixers in the transceiver will depend on the multiplier and PA output powers, which are not directly measurable. According to simulations, the mixer design does not significantly degrade the input compression point of the receiver; this is due to its higher supply voltage and current density, its lack of an additional transconductor following the last stage of the LNA, and the power splitting that occurs at the last stage of the LNA going into the two mixers.

The results of the s-parameter measurements from 70-110GHz are illustrated with solid lines in Figure 2-3 under typical operating conditions, showing that the input and output characteristics agree fairly well with expectations, but the transmission S_{21} shows that the measured gain experienced a significant high-pass characteristic, dramatically reducing the gain below 90GHz compared to the simulation. Several potential causes for this discrepancy were investigated, including incorrect current bias, differences in modeling methodology depicted in Figure 2-1, device variation among the three corners provided by the foundry, bias resistor variation, and MIM capacitor variation. It was found that MIM capacitor variation best explained the discrepancy in gain, and that a 35% global MIM capacitor reduction in simulation produces good agreement between measurement and simulation, as shown in Figure 2-4.

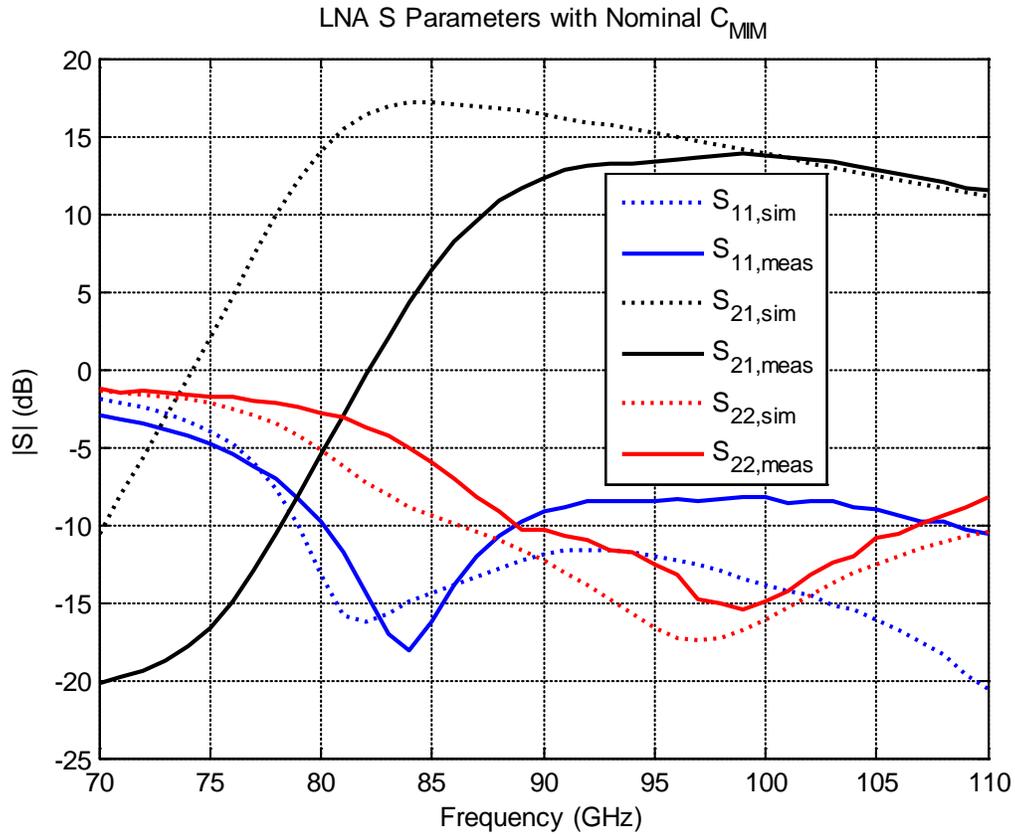


Figure 2-3: LNA design measurement versus nominal simulation

Although global MIM variation aligns measurement and simulated gain very well, it misaligns the input match frequency slightly compared to the original schematic simulated for Figure 2-3. Since variation between adjacent MIM capacitors is likely to be much smaller than global variation, this suggests that in addition to MIM variation, the parasitic extraction of the first stage is slightly wrong and underestimates the parasitic input (base) capacitance of the minimum-length first stage of the LNA. To improve this matching, a more accurate 3-dimensional parasitic extraction tool could be used—a more basic extractor was used for this design—and minimum emitter-length devices can also be avoided. Unfortunately as frequencies increase, the optimum emitter length and power consumption for noise figure decreases, as shown by [17], so increasing the input-stage emitter length could have a detrimental effect on the LNA noise figure. Redesigns currently in progress are focusing on more robust *a priori* matching between

simulation and reality, in part by using only larger AC-coupling capacitors. The design whose measurements are shown here used inter-stage MIM capacitors in the range of 70-100fF to participate in the inter-stage impedance matching networks. Under significant process variation, these impedance matches can significantly affect the amplifier transfer function. Using only 300fF capacitors (for example) for inter-stage coupling will result in minimal differences in performance even under large (35%) variations in capacitor values.

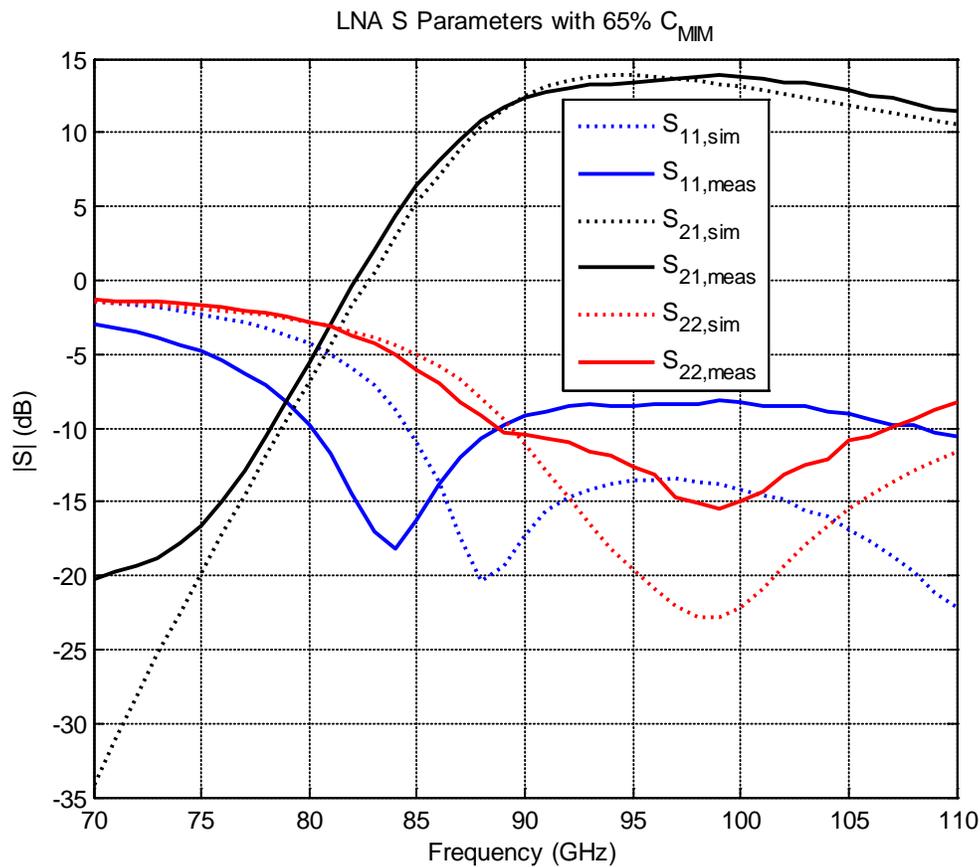


Figure 2-4: LNA measurement versus simulation with 35% MIM capacitor reduction

Two simulations of noise figure using the original design and the updated simulation with reduced global MIM values are shown in Figure 2-5, along with noise figure measurements performed using a rectangular-waveguide (WR10) down-converter, commercial amplifier, and W-band noise source. Noise figure measurements at this frequency are complicated by the non-

ideal components used to make the measurement; however the results agree fairly well with the minimum noise figure designed to be around the radar center frequency of 94.5GHz. The 10GHz bandwidth over which the measurement can be made is limited, especially compared to the VNA measurements which can span 110GHz. This is due to the frequency range over which the external multiplier and mixer have adequate conversion gain. The lowest noise figure measured is 7.5dB, and the noise figure generally lies between the original and the MIM-variation simulation results for frequencies around the operating center frequency.

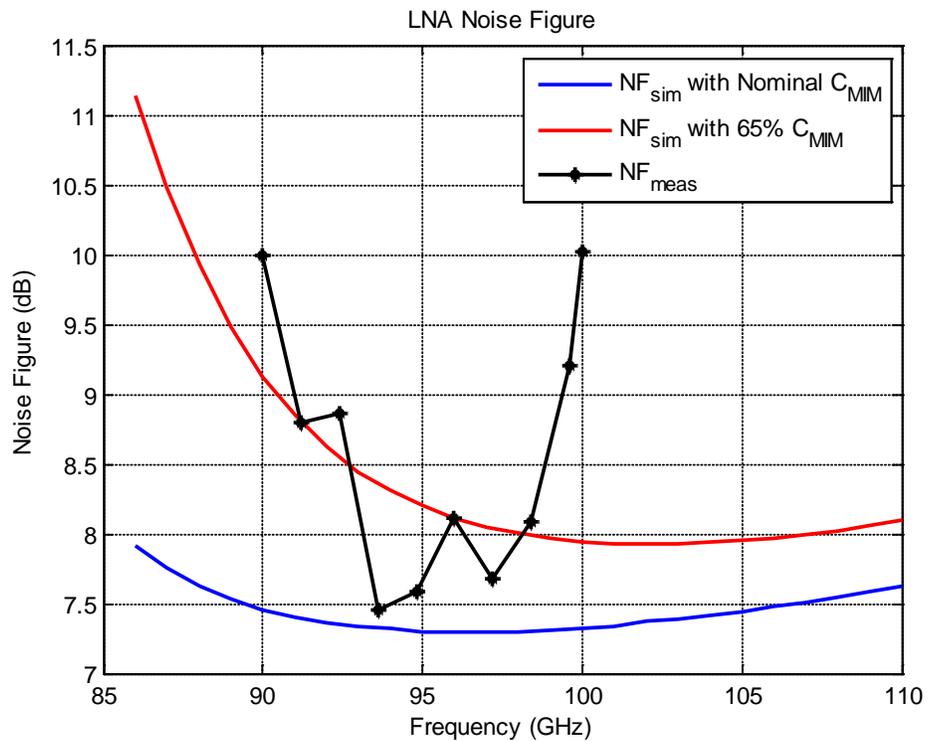


Figure 2-5: LNA noise figure measurement and simulations

Mixer Design

The mixer was designed to be as simple as possible due to tape-out deadlines, while achieving a high input compression point when combined with the LNA. It is balanced with respect to the LO port, but uses only a single RF input. The structure is essentially half of a Gilbert (current-commutating) micro-mixer, although the upper “switch” devices do not act like MOS switches at this frequency. Early designs attempted to implement a full Gilbert micro-mixer, which is somewhere between balanced- and un-balanced with respect to its single RF port; however aligning the delay of the two RF current paths to be 180-degrees out of phase proved to be difficult at the millimeter-wave operating frequency. As a result, the mixer design uses half the micro-mixer structure and simply removes the diode connection on the tail device, as shown in the schematic of Figure 2-6. The result is similar to a standard single-balanced mixer except that the transconductor that drives the RF port is the final stage of the LNA, and the tail current in the mixer is used to pull a higher current through the switch devices than in the fourth LNA stage. The output impedance of the tail device is increased slightly by resistive emitter degeneration, which also helps stabilize the bias current. A matching network and transformer balun drive the bases of the switching pair and a relatively low load resistance allows a high compression point. Differential and single-ended capacitors are used to implement a low frequency pole. A single pole can effectively filter the 94GHz LO leakage out of the IF signal since the latter is an extremely low bandwidth signal (on the order of 1-50kHz).

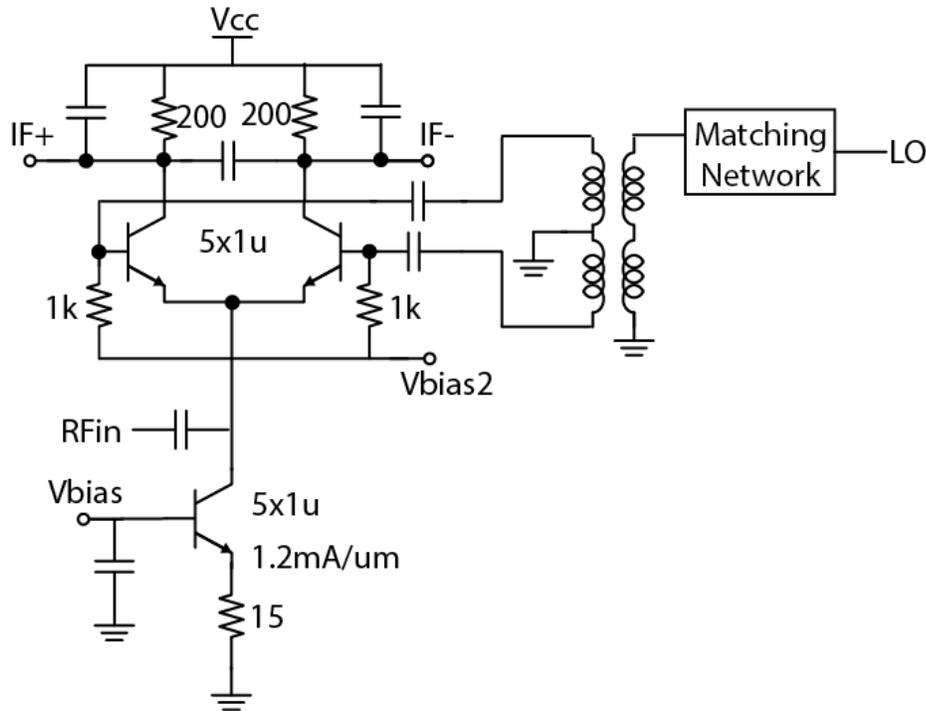


Figure 2-6: Mixer schematic

There are a variety of limitations of this mixer architecture. In particular, the balun does not provide perfectly balanced signals to the switching devices' bases, which results in different DC currents in the two paths and therefore a DC offset at the IF output. This DC offset causes problems in measurement since an amplifier is used to make the differential output single-ended in order to observe it on a spectrum analyzer and oscilloscope. In the transceiver architecture as implemented (described in Chapter 4), the LO swing is also not controllable separately from the power-amplifier output, which means that the mixer offset, conversion gain, and compression point are directly linked to the transmitter characteristics and operating point. Most importantly, the termination that the balun and matching network present to the broadside couplers of the transceiver architecture is inevitably not 50Ω , which has serious negative implications for the transmitter isolation of the transceiver as discussed in later chapters. Still, the LNA and mixer combination work well enough to demonstrate basic range-sensitive behavior in the unit-cell transceiver. An optimized mixer (and/or LNA) should be redesigned and tested separately.

Chapter 3 Asymmetric Transformer

Balun Design

Baluns and their Applications

Baluns are frequently used in RF, microwave, and millimeter-wave front end circuits to transform single-ended signals to differential signals or vice versa. In particular, antenna signals arriving at the chip or RF module boundary are frequently single-ended and are often carried by a planar waveguide structure such as a microstrip transmission line or coplanar waveguide transmission line. Later we will examine coupling of these high-frequency signals onto the chip and the methods for simulating transition structures between on-chip and on-board waveguides.

High-frequency circuitry on-chip (such as low-noise amplifiers, mixers, and IF amplifiers) will frequently prefer differential architectures for the same reasons they are preferred at lower frequencies. Among other benefits, the main advantages of differential circuits are that interferers can be rejected if they couple to both sides of the circuit in common-mode; a larger (differential) signal swing can be represented in the same voltage headroom; and the negative effects of certain parasitic elements and device nonidealities can be diminished by making them appear in common-mode but not differential-mode: for example gate-drain Miller capacitance can be neutralized in the differential mode of a differential pair; and unwanted emitter

inductance (due to a wire bond) can be made to appear only in common mode in a differential pair. Aside from the initial transition from single-ended on-board signals to differential on-chip signals, the high-frequency circuits operating within the transceiver architecture may require baluns as well: for example to make a single-ended LO signal differential in order to drive a balanced mixer. This application in particular is part of the radar transceiver architecture described in later chapters.

The method by which the single ended signals are transformed into differential signals can play a big role in determining the overall system performance. For example, consider a receiver with a differential LNA input and an off-chip balun to convert the antenna signal to a differential input the LNA can accept. In this case, the off-chip balun may be matched to 50Ω to allow the receiver performance to be insensitive to its position on-board (assuming it is connected to the chip with on-board 50Ω transmission lines). This is an important consideration for very high frequency designs where the wavelength on-board may be even smaller than the chip dimension: on the order of 1 millimeter for W-band designs. If we assume all components are matched to a 50Ω environment then the well-known Friis equation for noise factor shows that balun loss not only adds its loss to the overall system noise factor, but also increases the contribution from the LNA term by the inverse of the balun gain (which is less than 1 for a passive balun) [10].

$$F_{SYS} = F_{BALUN} + \frac{F_{LNA} - 1}{G_{BALUN}}, \quad G_{BALUN} (dB) = F_{BALUN} (dB)$$

For on-chip baluns, there is typically very little distance between the balun output and the LNA input, so transmitting the signal between the two elements via transmission line may not be necessary. In that case the balun's differential output would typically not be matched to 50ohms but instead conjugately matched to the LNA input impedance. Then the calculation above is

complicated by dependence of gain and noise factor on each block's apparent termination impedance, and Friis' simple equation is only approximately correct. In addition to balun optimization for the input of a differential circuit, at the output of an RF or millimeter-wave circuit, the design of a high power handling, low-loss balun may be critical to efficiently transfer power from a differential power amplifier to a single-ended antenna, T/R switch, or transmit filter.

Integrated Transformer Baluns

A variety of active and passive baluns have been designed and optimized for RF and microwave circuit designs, including planar waveguide versions such as the Marchand balun, and transformer versions which couple signals using the magnetic flux in a coil of wire. At millimeter-wave frequencies both are feasible on-chip, although transformer baluns are particularly attractive since they can achieve fairly wide bandwidth and comparable transmission loss while being very compact; whereas many waveguide baluns (as well as other waveguide couplers) require lengths on the order of $\lambda/4$ in their constituent waveguides, and the implemented electrical length changes proportionally with frequency. For the purposes of design, a transformer balun is just a transformer whose primary coil is driven at one end and grounded at the other end through some path which should be low-impedance. This primary coil couples magnetically (and capacitively) to a secondary coil in such a way that a differential signal appears at the leads of the secondary coil.

A transformer balun is often designed with the addition of a center-tap on the differential coil to "equalize" or make-differential the output signal by creating a low impedance at that center-tap. To allow biasing of subsequent circuitry through the secondary coil, the low impedance at the secondary coil's center tap is created with a (possibly large) capacitor. If the signals on the secondary coil and throughout the balun structure are purely differential with respect to the axis of the differential leads, then the secondary coil's center tap is a virtual ground and adding a low

impedance to the center-tap will not affect the differential signal. However, it will appear in the common mode of the output network, and therefore shunts common mode energy to ground. Such a configuration is shown in the schematic below, where Port2 is setup to analyze the differential mode output of the balun and the common mode output is singly terminated by the center-tap capacitance. (Ferrite cores are not actually used at this frequency; this report reused existing circuit element drawings).

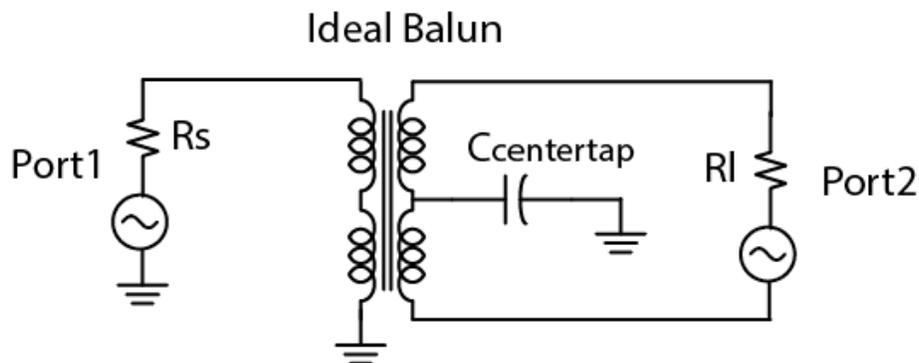


Figure 3-1: Schematic of ideal transformer balun with port connections

Maximizing the power transfer from the single ended input to the differential output is desirable. Inevitably, some of the signal power is lost due to resistive parasitics in the coils of wire, the substrate, and any nearby structures which couple electrically or magnetically to the balun. The loss in the inter-layer dielectric material is also a mechanism that dissipates power, but it is typically smaller than the loss in the nearby conducting or semi-conducting substrate because the oxide materials used to separate metal layers in an integrated circuit process can be very pure. In addition to power dissipation mechanisms, some of the signal power will appear in common-mode due to non-ideal coupling to the secondary coil. This signal power is typically not useful to the subsequent differential stage and can cause common-mode oscillation problems if the common mode is not terminated properly.

Together, the transmissive properties of the balun can be summarized in a reduced form using two metrics which can be easily simulated. The first is G_{max} , which is the gain from the single-ended input to the differential output that would be realized if those two ports were bi-conjugately matched using ideal, lossless matching networks. The value for this maximum gain will depend on the center-tap capacitance used and the common-mode termination impedance. It is typical to assume the common-mode impedance is one quarter of the differential mode impedance, as would be the case if each differential lead were terminated separately in a 50Ω single-ended environment. In practice both the common mode and differential mode termination impedances would depend on the circuit connected to the balun's differential output. The second metric of interest is common-mode rejection ratio, or CMRR, which is defined to be the ratio of the differential-mode signal power to the common-mode signal power at the output. As with G_{max} , CMRR depends on the center-tap capacitance selected and the assumed (or known) common-mode termination impedance.

At millimeter-wave frequencies, where the signal's wavelength approaches the size of the structures drawn, the design of passive elements such as transformer baluns can be tedious because an electromagnetic field solver is required to get good agreement between designs and their fabricated realizations. Such field solvers can be analytic in nature—understanding the shape of the structure being solved and the equations that describe its electromagnetic fields—or they can be geometry-agnostic in nature. The latter case is called a finite-element method and uses an adaptive meshing algorithm to subdivide arbitrary geometries into small elements with known interactions between neighbors defined by Maxwell's equations. In either case, the size, shape, and relative position of wires becomes very important to the overall functioning of the passive device.

The structures must also be made fairly simple so that, in most cases, they are designed to operate below their self-resonant frequency. For a transformer, the self-resonant frequency

seen at one port depends on the load impedance connected to the other port, but intuitively the definition is similar to that for an inductor. The self-resonant frequency is the frequency at which the secondary coil (for example) transitions from appearing inductive to capacitive. Above this frequency, a capacitive load cannot be made to resonate with the inductance of the secondary coil, since the coil presents a negative reactance (appears capacitive) to the load. Since transistors at this frequency appear predominantly capacitive, the cases when we want the balun to present capacitive impedance rather than inductive impedance to its load are less common.

There is no fundamental reason why a passive device such as a transformer balun cannot be operated above its self-resonant frequency: for example a matching network could transform a capacitive load to an inductive one which is then presented to the overall-capacitive output impedance of the transformer. This, however, requires a matching network which introduces its own transmission loss, increases area, and can potentially require a high quality factor due to a large impedance conversion ratio. The last of these concerns in particular is a barrier against reliably designing for a certain center frequency amidst process variation and modeling errors. High quality factor impedance transformations also increase the loss in any real matching network, which is itself composed of lossy elements such as transmission lines, capacitors and/or inductors. Furthermore, the desirable properties of the transformer balun (such as its CMRR) may be severely degraded above the self-resonant frequency: it may be that the majority of the signal power appears in common-mode at the secondary coil, which is undesirable.

For these reasons, transformer baluns operating in the W-band, such as those used in the 94GHz circuits described herein, typically have primary and secondary coils with few turns (often just one). In this frequency range, attempts to design 2:1 transformers using multi-turn coils or multiple coils connected in parallel or series typically resulted in self resonant

frequencies close to or below the operating frequency, so the turn-ratios are typically made 1:1. The effect on self-resonant frequency is due to parasitic capacitance to the substrate, to other parts of the coils themselves (including underpasses), and to ground nodes which act as current return paths. Even 1:1 single-turn transformer baluns can suffer from high parasitic capacitive coupling which can cause much of the output signal power to appear in common mode rather than differential mode; this is particularly the case when the primary and secondary are “stacked”, meaning they are composed of two upper metals with vertically overlaid paths. Although not a strict characterization, a general observation of EM simulations at this frequency is that as the operating frequency approaches and then exceeds the passive’s self-resonant frequency, the desirable magnetic coupling to produce a differential output signal diminishes and is overtaken by purely capacitive coupling which is both common-mode and differential-mode, degrading the CMRR of the balun.

Asymmetric Transformer Balun

A typical design of a transformer balun will simply use a symmetric transformer, with the primary and secondary leads oriented collinearly, as shown in the single-turn transformer in the lower right drawing of Figure 3-4. This is typically acceptable in low-frequency applications where transformers and transformer baluns use many turns and large center-tap capacitors can be used to balance the secondary. As described previously, W-band designs typically require single-turn coils to achieve sufficiently high self-resonant frequencies. In this case, the 180-degree orientation results in half of the secondary coil coupling to the grounded half of the primary, and the other half of the secondary coil coupling to the driven half of the primary. Such an asymmetry from the secondary coil’s perspective limits the achievable CMRR for 180-degree oriented transformer baluns if center-tap capacitance cannot be made infinite.

The concept of the asymmetric or rotated transformer balun is that it attempts to equalize the coupling between the primary and secondary coils such that—on the net—each symmetric half of the secondary coil couples to substantially similar parts of the primary coil. Here “on the net,” means that if you consider some average apparent coupling and impedance between the primary and secondary coils (by considering, for example, the coupling between many small parallel segments of wire), then the calculation should give the same result for both symmetric halves of the secondary coil. In practice, a numerical field solver is used, and this and the following argument are provided only for intuition, which can help inform what to try modeling in the simulator.

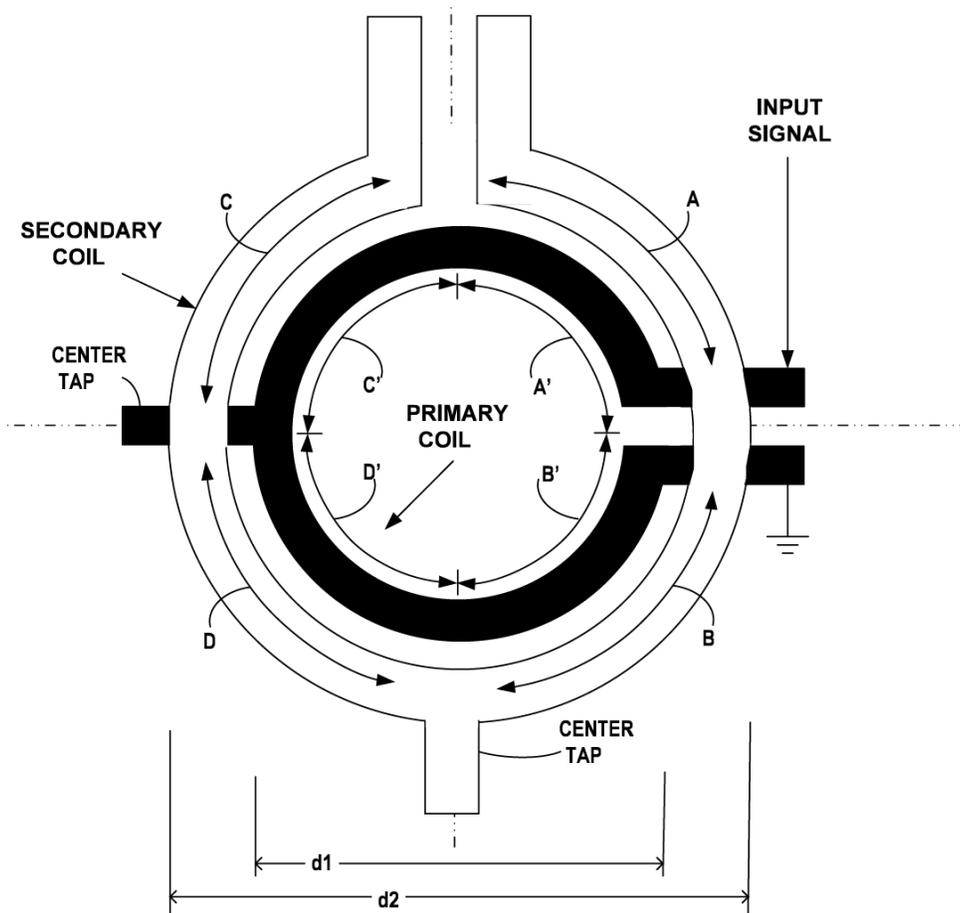


Figure 3-2: Diagram of single-turn asymmetric concentric single turn transformer balun

To see why rotating the two coils with respect to one another can improve CMRR, consider the drawing of a single-turn asymmetric (90-degree) transformer balun in Figure 3-2. The primary coil is grounded on one lead and is driven by some higher impedance on its other lead (for example a 50Ω transmission line). Then due to approximately uniform shunt capacitance and series inductance along the coil, the impedance to ground of any point along the primary coil is a decreasing function as the coil is traversed from the driven lead to the grounded lead. To simplify this complicated and unknown impedance relationship, suppose we segment it into just two parts: the upper half (A', C' segments) of the primary coil is substantially higher impedance than the lower half (D', B') due to the asymmetric lead impedances. Also, due only to proximity, the majority of the coupling between primary and secondary occurs between corresponding lettered segments: A couples to A', B to B', and so on. Coupling is both inductive and capacitive and can be adjusted by changing, for example, the widths of the traces and their relative position or overlap. This implies that the left half of the secondary coil (C, D) couples to one higher impedance segment (C') and to one lower impedance segment (D'). Similarly, the right half of the secondary coil (A, B) couples to one higher impedance segment (A') and one lower impedance segment (B'). This similarity is an intuitive explanation of the improvement in CMRR that can be achieved with a rotated structure compared to a symmetric (180-degree) transformer balun. In the symmetric balun, one half of the secondary is coupled to two low-impedance segments of the primary, and the other symmetric half of the secondary is coupled to two higher-impedance segments of the primary. Such a difference will require more center-tap capacitance to balance the differential output, as will be shown in simulations described in the next section.

A corollary of the explanation above (which simplified the primary coil into just two “zones”) is that the 90-degree orientation is not necessarily the best. In practice, as shown in Figure 3-3, the rotational orientation which maximizes CMRR will be some angle strictly between 0 and 180

degrees, assuming certain restrictions described later on the center-tap capacitors, and for a known ground impedance on the primary coil. In a real integrated circuit process, metal layers are often restricted to have 45-degree angles and edges, so the circular coils may be implemented as octagons, and their relative rotational orientation restricted to multiples of 45-degrees. For this reason the 90-degree configuration may be the best (in terms of CMRR) that can actually be implemented in single-turn transformer baluns. It also has the benefit that all its leads (the differential secondary, singled-ended primary and both center-taps) are on four distinct sides of the device. This allows connections to center-tap capacitors to be much more straightforward than in the symmetric (180-degree) configuration. As will be shown in the next section, the structure is intrinsically more balanced and requires less center-tap capacitance to maximize transmission and CMRR.

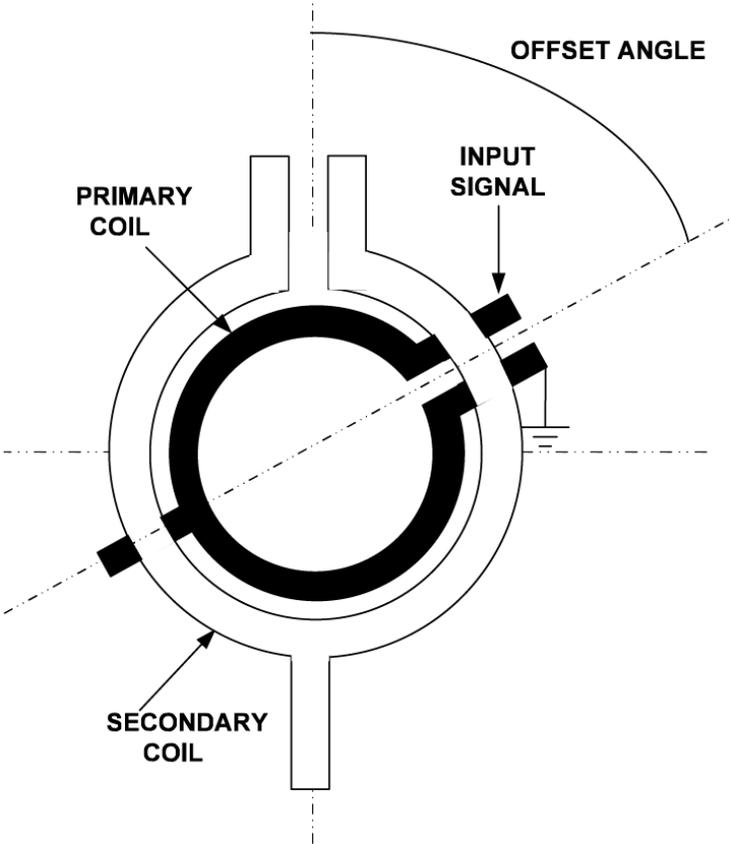


Figure 3-3: Asymmetric transformer balun with arbitrary offset angle

Simulation and Optimization

To analyze the benefits in CMRR that an asymmetric transformer balun can provide, four 50 μm diameter transformers modeled in a commercial silicon process were first simulated in an electromagnetic field solver. These are shown in Figure 3-4 with the encapsulating dielectric layers, the substrate, and top air-box hidden for clarity. The left column shows concentric coils and the right column shows “stacked” coils which overlap completely. The top row shows 90-degree orientations and the bottom row shows 180-degree orientations between the two coils.

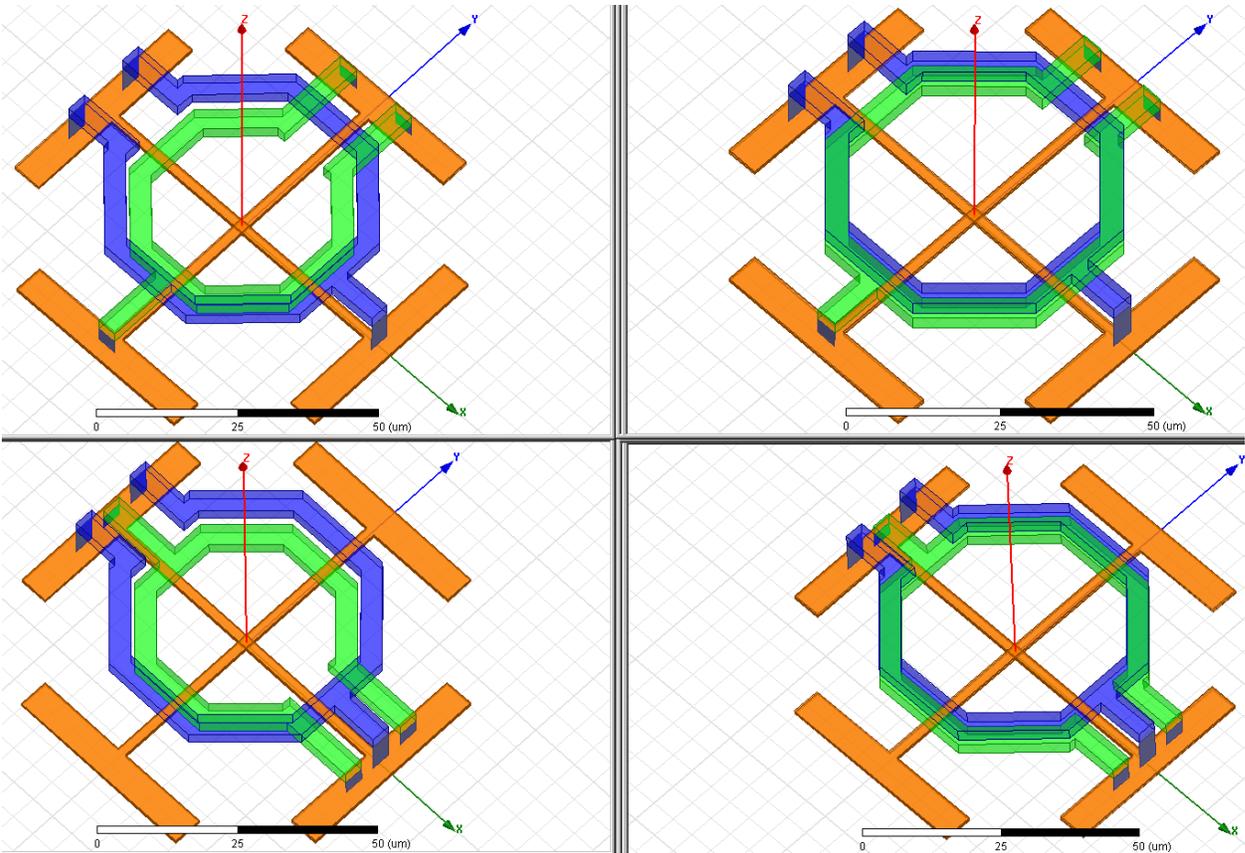


Figure 3-4: Four rotational and concentric variations of transformer baluns

Since connections to MIM capacitors are made on the topmost metal layer, and ground planes are often implemented in lower metal layers, it is most convenient to use the upper (blue) coil as the differential secondary, and drive one of the lower coil’s leads (green) with the single-

ended input, grounding its other lead. An orange grounding “cross” is used in all simulations to allow return currents to flow and to provide a common reference to the lumped ports exciting the transformer. The grounding cross is identical for all four configurations.

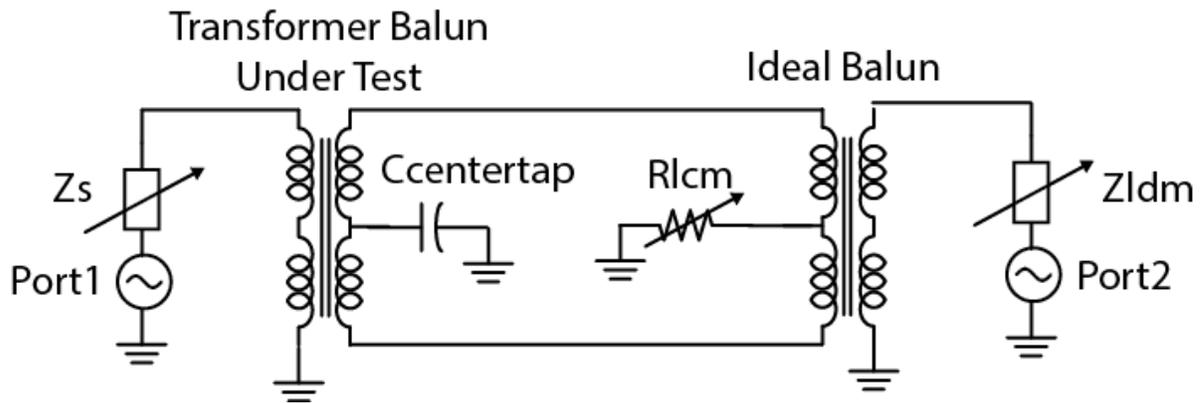


Figure 3-5: Schematic of optimization test-bench for transformer baluns

The simulation test-bench, shown in Figure 3-5 is setup to analyze $CMRR$ and G_{max} of the transformer baluns under certain on-chip assumptions, and around a nominal operating point of 95GHz. In particular, restrictions are placed on the quality factors of the input and output matching networks that would be required (based on resistive transformation ratio). For the balun diameters tested, the limits of port resistance and port reactance are never reached, and all baluns require reasonable source and load impedances under a bi-conjugate match: typically around $(30 - j40)\Omega$, which is equivalent to an optimum source impedance of 83Ω in parallel with $27fF$ of capacitance, a suitable range for connecting to transistors in this technology.

A limited amount of center-tap capacitance is allowed; in this case 500fF. According to MIM capacitor models in our process, MIMs around or above this value (depending on their geometry) have self-resonant frequencies approaching the operating frequency. This is not necessarily a problem, since MIM capacitors operated above their self-resonant frequency still exhibit a low (inductive) impedance, and would therefore still function effectively as a low center-

tap impedance to ground. The issue is that if the self-resonant frequency from the model does not match the fabricated device's, the designer could get unlucky and the self-resonant frequency could align with the operating frequency, where the center-tap capacitor could be either extremely capacitive or extremely inductive, depending on whether the operating frequency is slightly below or slightly above the self-resonant frequency respectively. Additionally, very large capacitors could cause common-mode oscillation at a frequency much lower than the operating frequency if the circuit has gain around the capacitor's self-resonant frequency. Issues such as these can be concerning, although not impossible to mitigate, if avoiding common-mode oscillation is a priority. So it is best to use a capacitor at an operating frequency far away from its self-resonant frequency.

The optimization bi-conjugately matches the transformer balun for peak gain at 95GHz by adjusting the port impedances such that $|S_{21}|$ is within 1m dB of the maximum gain G_{max} . This is simply a numerical way to compute the optimum port impedances using the optimizer. Since there is a unique bi-conjugate match solution for 2-port networks, the optimizer typically solves this part of the problem after only a couple iterations. In a real circuit, this modification of the port impedances corresponds to sources and loads (or matching networks) placed on the inputs and outputs of the balun that present to it termination impedances that maximize the balun's gain. In practice, the designer would typically do the reverse: selecting a suitable balun given their source devices and loading devices. The losses of the matching networks—if they were required—are not included, although the restrictions described above on port resistance and reactance ensure that the matching network Q would not be very high, and therefore the matching networks would not be very lossy. Finally, the common-mode is terminated with one-quarter of the differential mode resistance. This is an approximation and would depend on the actual differential transmission line connected to the transformer secondary. The results of the optimizations are summarized in the table of Figure 3-6.

Orientation	Coil Positions	$S_{21}=G_{\max}$ (dB)	CMRR (dB)	Center-tap Capacitance (fF)
90-degree	Concentric	-0.99	37.8	98
90-degree	Stacked	-0.85	37.7	86
180-degree	Concentric	-0.91	31.7	500
180-degree	Stacked	-0.75	29.6	500

Figure 3-6: Table of optimization results for four transformer balun structures

One of the main reasons that the asymmetric balun achieves higher CMRR is due to the maximum restriction placed on the center-tap capacitor. We expect that the symmetric transformer balun will achieve its maximum CMRR with infinite center-tap capacitance, since it shunts more and more common-mode energy to ground as the center-tap impedance decreases. The optimization uses ideal capacitors connected to the center-tap, so the center-tap capacitor does not become inductive or self-resonate. Indeed, the optimizer shows that both 180-degree configurations of the balun reach their maximum CMRR with the center-tap capacitance at its boundary value, which is the maximum allowed center-tap capacitance of 500fF. By contrast, to maximize CMRR in the 90-degree configurations, a much smaller capacitance is required on the secondary: only around 100fF. This saves area and potentially reduces the risk of common-mode oscillation due to a self-resonating center-tap capacitor.

The simulation results show that at millimeter-wave frequencies, the asymmetric (90-degree) configuration output performs the collinear (180-degree) configuration regardless of whether the coils are concentric or stacked, and always uses less center-tap capacitance to achieve a certain level of CMRR. The collinear configuration benefits more from offsetting the coils' diameters, since the capacitive coupling between the coils is reduced. The asymmetric orientation is intrinsically more balanced and so it makes little difference whether the coils are stacked or concentric in diameter. Finally, a small benefit in transmission is derived by using stacked coils, as is expected since the mutual inductance of the coils increases (or equivalently, less magnetic flux is "lost" by the lower coil's smaller diameter).

In applications such as driving a mixer's LO port differentially, CMRR may be important to minimize distortion. In other applications, such as when driving a differential pair which has good common mode rejection, common mode variation at the output of the balun is less important, and the improvement in CMRR may not be needed. The asymmetric transformer balun also has the disadvantage that a particular value of center-tap capacitance is needed to maximize CMRR, whereas its symmetric counterpart requires only "a lot" or "enough" center-tap capacitance. The latter may be more useful if there is significant variation in capacitor values, as there can be when using MIM capacitors at the center-taps.

Chapter 4 Radar Transceiver Front-End

Radar Unit Cell Design

The conceptual block diagram of the FMCW radar was previously shown in Figure 1-2, but differs in several ways from the architecture actually implemented for this work. The architecture implemented can be described as a polarimetric radar system in that its functionality relies on polarization-diverse signaling to achieve transmission and detection of the frequency-modulated W-band carrier. In particular, having multiple axes of polarization—as is the case with circular polarization—can improve the detection ability of radar systems by eliminating polarization “blind-spots,” which occur when targets reflect preferentially along one axis. For autonomous aerial vehicles, for example, reliable power-line detection is essential, but the radar cross section of a power line differs dramatically depending on the polarization axis of illumination and the incident angle of that illumination provided by the radar transmitter [19]. The work described in this report was to design and implement the circuit blocks for such a circularly-polarized 94GHz radar and better understand the system-level tradeoffs of the transmitter-leakage cancelling architecture. The architecture is depicted in Figure 4-1 and was devised prior to this work by Nokia Research Center, Berkeley.

The main distinguishing feature of the transceiver is that instead of physically separating transmitters and receivers to reduce the transmitter leakage, this architecture seeks to cancel

To see how the transceiver unit-cell could cancel transmitter leakage, it is necessary to sequentially follow the signal from its origin at the PA through its path to the receiver. The PA emits a differential signal, at phase references 0 and 180 degrees. The two quadrature couplers have their isolated ports connected to a Wilkinson combiner so that the transmitter leakage is cancelled by the differential rejection of the Wilkinson combiner, which can be high due to its symmetry, after first being attenuated by the isolation of the couplers. The quadrature couplers also drive two mixers: an in-phase and quadrature-phase mixer. The transmitted signal couples to a dual-polarization antenna, which could be the dual-lead patch antenna depicted in Figure 1-5. One of the feeds is delayed by 90 degrees (suppose the pair are now at 0 and 270 degrees) so the antenna emits circularly polarized waves. Objects reflect the two polarizations with the same time delay that they arrive, so the returning circularly polarized wave still has the same 0- and 270-degree relative phases, and the same polarization also passes through the 90-degree delay line again, producing in-phase received signals. These are passed through the couplers and the Wilkinson combiner where they are received and amplified by the LNA and the rest of the receiver circuitry.

This architecture is elegant in that it directly addresses the transmitter leakage problem rather than minimizing it using physical separation or a circulator made in a non-silicon process. It also supports both transmit-side beam-forming (through the LO phase shifting preceding the PA output stage) and receive-side beam-forming (using IF phase shifting, or weighted recombination of the two IF mixer outputs). Although two antenna leads are still required, the architecture may also reduce the size of the phased array since only a single antenna is needed per transceiver and that antenna is used to simultaneously transmit and receive. The architecture is also attractive for its scalability, since larger arrays can be made by tiling additional transceiver/antenna pairs with minimal additional effort, once a small-scale array has been successfully implemented.

The architecture has many important drawbacks which limit its performance and introduce design difficulties not found in fully-separated transmit / receive radars. The first and most important is that the directional couplers are passive devices, which means that whatever ratio of the power from the PA gets split between the antenna and the mixers, that same fraction of received power incident on the antenna feeds also gets split between the LNA and the PA. This adds at least 3dB to the receiver noise floor, and more in practice because the couplers are lossy elements. A related issue is that the architecture spends half of the PA output power driving the mixers and half driving the antennas. In many cases, we would prefer to transmit more power and use less of it driving the mixers, so an unequal split coupler could be designed. Unfortunately, adjusting the quadrature couplers to preferentially transmit power from the PA to the antenna also means that the received signal will preferentially couple to the PA output and not to the LNA input. Therefore, modifying the coupling ratio does nothing for the overall link budget and the couplers may as well be made to have equal split. In an array, each element's transmit power is kept to a moderate level, so it might not be a big concern that half the PA power is used to drive the mixers. Although the mixers may need substantial power to drive their LO ports, the architecture shown does not allow separate control of the LO path gain, meaning that the transmitter setting completely determines the LO-drive operating point of the mixers.

Another major limitation is that if the received signal is not circularly polarized but elliptically polarized (meaning the target reflected one polarization preferentially), then the received signals incident on the Wilkinson combiner inputs will not be purely common-mode but will have some differential component that the Wilkinson combiner will absorb. In the worst case, the reflected signal is along only one polarization axis, so an additional 3dB of power loss occurs in the Wilkinson combiner, in addition to its intrinsic loss which is in the range of 0.5-1.0dB (another insertion loss which is not required in separated transmitter / receiver architectures). Finally, the transmit leakage cancellation only works if the structure is purely differential, meaning that the

Wilkinson, broadside couplers, and PA terminals all need to be symmetrically loaded across the horizontal axis of Figure 4-1. While this can be made symmetric on-chip, the overall system will not be symmetric since the 90-degree delay line is required on one of the antenna leads but not the other. Due to its length, that delay line does not only add loss, but it also inverts the normalized input impedance of the antennas, which will typically not be 50Ω , and will certainly not be 50Ω in a real operating environment.

A final major limitation of the architecture implemented is that the broadside coupler has its isolated port diagonally opposite their driven port. This is different from rectangular waveguide or Lange couplers, whose isolated ports are on the same end as the input port. The result is that the architecture is non-planar, and the LNA must be completely enclosed in millimeter-wave signal paths. It also means that to get all signals to the chip periphery, two sets of signals must cross. In order to minimize loss and coupling along millimeter-wave signal paths, the IF signals were routed beneath a microstrip ground plane to their pads. Crossing low- and high-frequency paths seemed like less of a concern than crossing two 94GHz signals, which could experience tremendous amounts of loss at the crossing point. LNA and mixer supplies and bias currents also had to be routed beneath point of underpass, significantly increasing the complexity of the layout and potentially reducing the robustness of the transceiver by introducing cross-talk.

Due to decisions made at layout time, the implemented unit cell shown in Figure 4-2 differed from the system level design shown in Figure 4-1. In particular, the I/Q hybrids shown driving the mixers were not used, but instead a 90-degree delay was introduced on one of the two LO paths. This is the only way to avoid crossing of millimeter-wave LO signals used to drive the two mixers. Although attempts were made at every component interface to impedance match components, this choice causes mismatch on the loading of the couplers because the mixer input impedance is not exactly 50Ω and the 90-degree delay line reflects the two mixer loads

unequally. Terminating the couplers with nonstandard loads degrades their isolation and allows more transmitter leakage to reach the LNA.

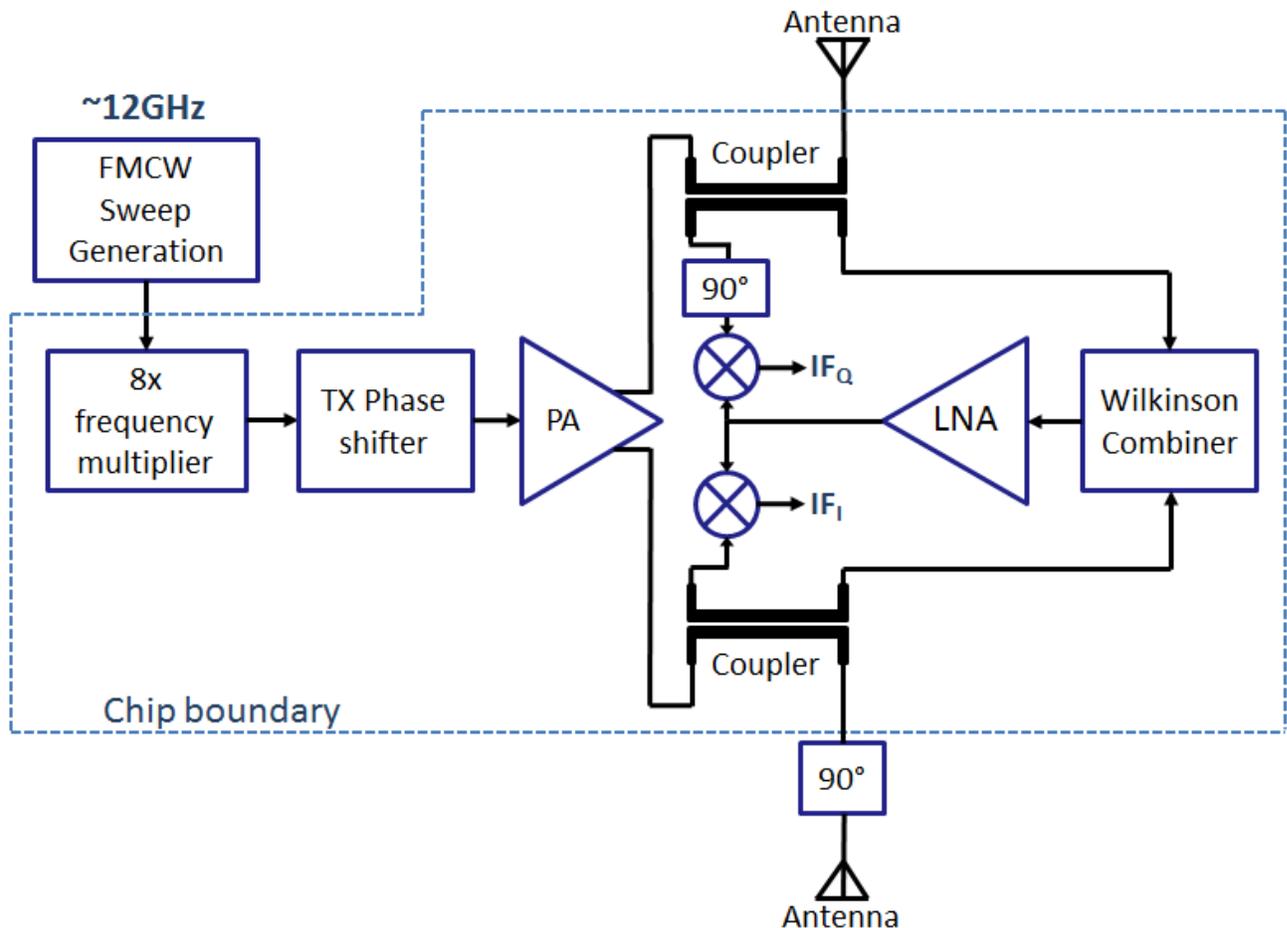


Figure 4-2: Implemented radar front-end showing chip boundary

Additionally, in order to fit the unit cell as well as several break-away components for testing, the antenna feeds were not made symmetric. This choice is in fact not that significant since in the real system the antenna feeds would be asymmetric anyway. In fact, if only the symmetric unit cell were designed and probed assuming the 90-degree antenna-feed delay line would be implemented added later, then the apparent transmitter leakage cancellation would appear better than it would be in practice. Still, the fact that the ground-signal-ground structure (GSG)

described in Chapter 5 was not well matched to 50Ω meant that this also degraded the transmitter isolation somewhat, and was analogous to the effect that mismatched antenna feeds will have on the fully integrated transceiver/antenna system.

Transmitter Leakage Analysis

The original intention of the first iteration was to develop the necessary passive structures and separated active structures (such as the LNA, PA, etc.) that would be needed to implement a full transceiver unit cell. Although this would have allowed more time to properly design the unit elements, optimism that a unit cell might work on the first try pressed the author and the co-designer to implement the full transceiver as well. This section quantifies some of the mistakes that were described above when implementing the architecture of Figure 4-1. With the help of post-silicon simulations, a more thorough understanding of how implementation details impact the transmitter leakage is developed. In particular, the mistakes made and simulations herein help inform the design and architecture changes that are proposed in Chapter 6.

As alluded to previously, the mixer-drive did not use additional I/Q hybrids to generate the in-phase and quadrature-phase mixer drive signals. This would have allowed the chip to remain perfectly symmetric as depicted in Figure 4-3. The figure also depicts perfectly balanced antenna feeds, which are the exterior leads of the broadside couplers. Due to the non-planar structure and the fact that the signal lines need to enclose the Wilkinson combiner, the LNA, and the mixers, there is a lot of lengthy routing connecting the broadside couplers to the two adjacent Wilkinson inputs shown at the top right of Figure 4-3. In addition to this loss on the receive-side, there are also a few hundred microns of routing to connect the PA output to the broadside couplers. All these excesses, in addition to the complications of the non-planar layout, can be eliminated by swapping the broadside couplers for “backwards” couplers, which could be implemented just using an underpass or using two metals layers coupling vertically.

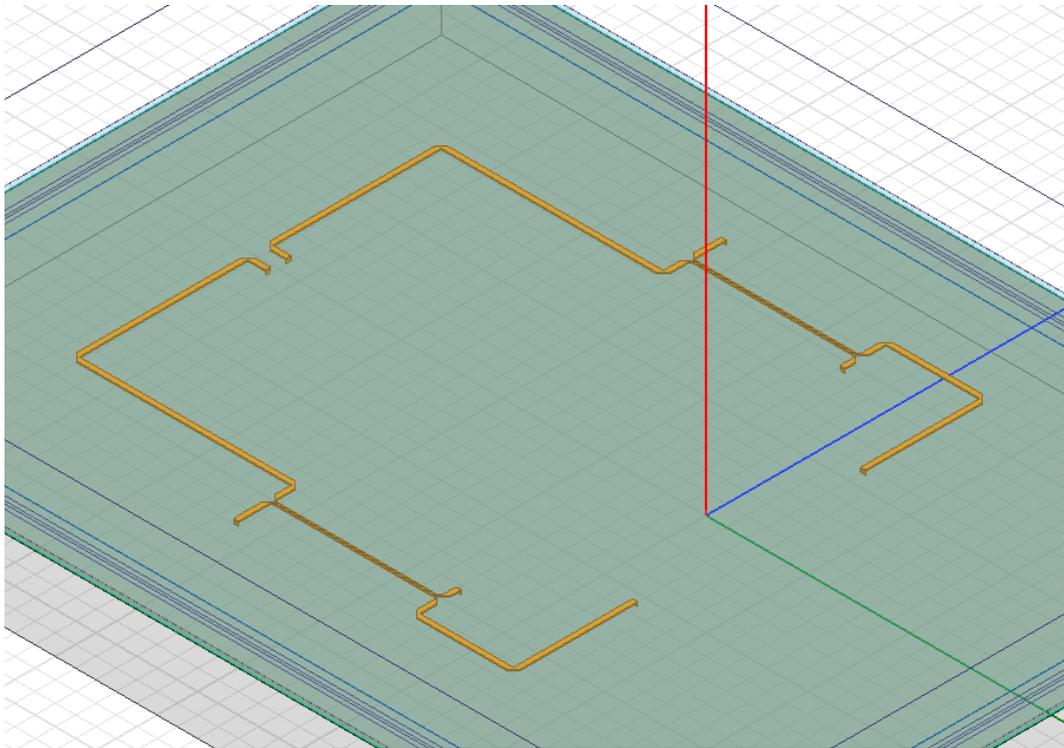


Figure 4-3: Idealized signal path showing perfect symmetry

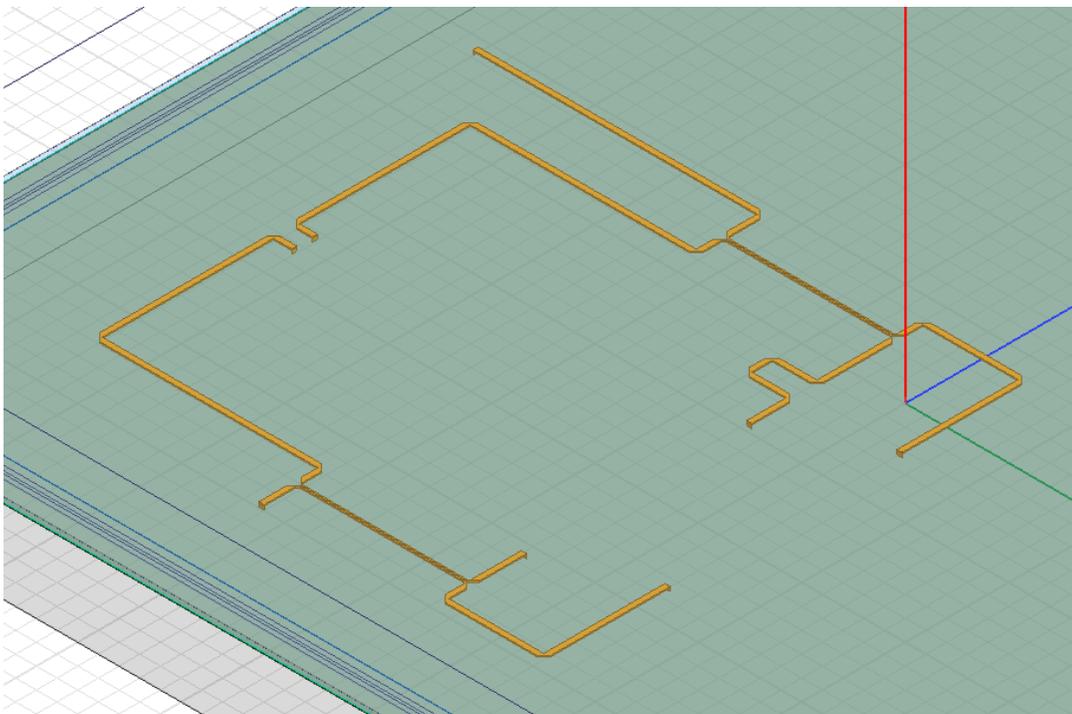


Figure 4-4: Actual signal path showing failure to maintain symmetry

Figure 4-4 depicts the signal path that was actually used. It adds, in addition to the excessive routing of Figure 4-3, two asymmetries to the signal path: the extended antenna port on the upper antenna feed (which takes the signal to pads at the chip periphery), as well as the meandered 90-degree delay feeding one of the mixers. The result of all these lines and couplers (including, now, the Wilkinson coupler as well) is that the insertion loss between the upper antenna and the LNA exceeds 8dB from pads to LNA input, as shown in Figure 4-5. Although the loss is significant, most of it is actually due to the architecture: we expect 6dB loss due to the broadside and Wilkinson couplers, the other 2-3dB comes from the GSG pad structure (not shown in Figure 4-4), the microstrip transmission lines, and the insertion loss of the broadside coupler and the Wilkinson coupler. If the two antenna receive signals arrive in phase and with equal amplitude (an idealized case), then 3dB of this loss due to the Wilkinson disappears.

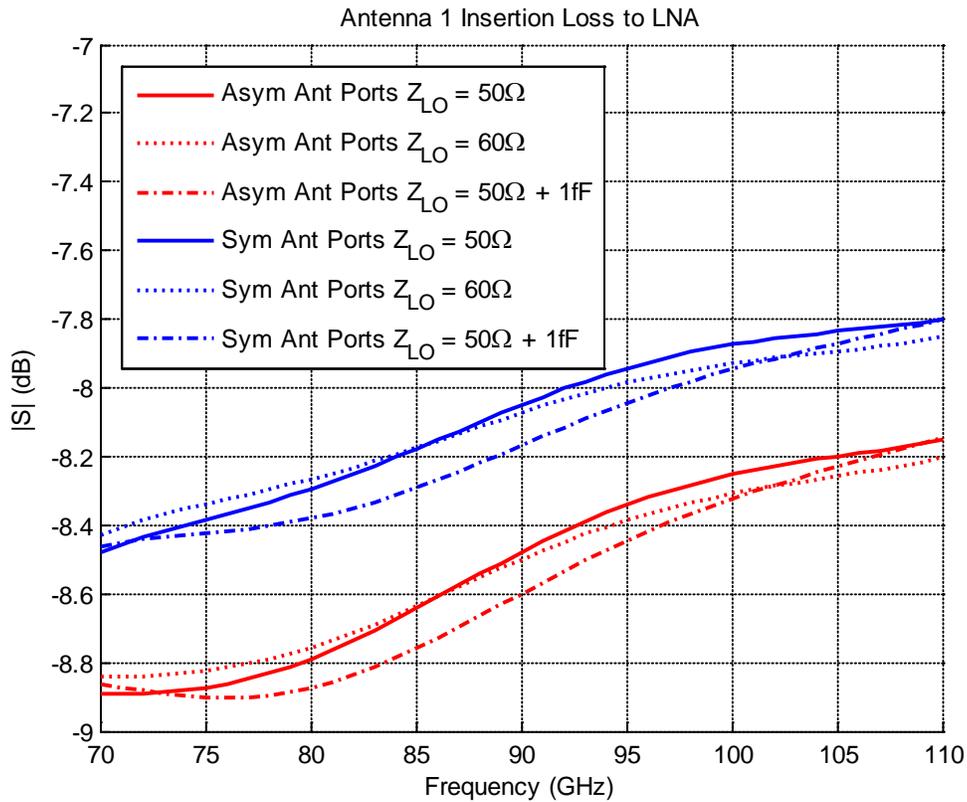


Figure 4-5: Upper Antenna Insertion Loss to LNA

Figure 4-5 also demonstrates the effect that mismatched loading has on the insertion loss both with and without the antenna-feed length mismatch. The effect is fairly minimal as the LO port impedance is adjusted (solid, dotted, dashed lines), and adds about 0.5dB due to the addition of the extra antenna feed line shown in Figure 4-4. We will see later that although the insertion loss is only slightly affected, the transmitter leakage is rapidly degraded as the LO input impedances deviate from 50Ω. Finally, Figure 4-6 shows similarly bleak results: the worst-case loss from one of the two PA outputs to its antenna is significant, although not very dependent on the LO termination impedance that the mixer presents to the passive network.

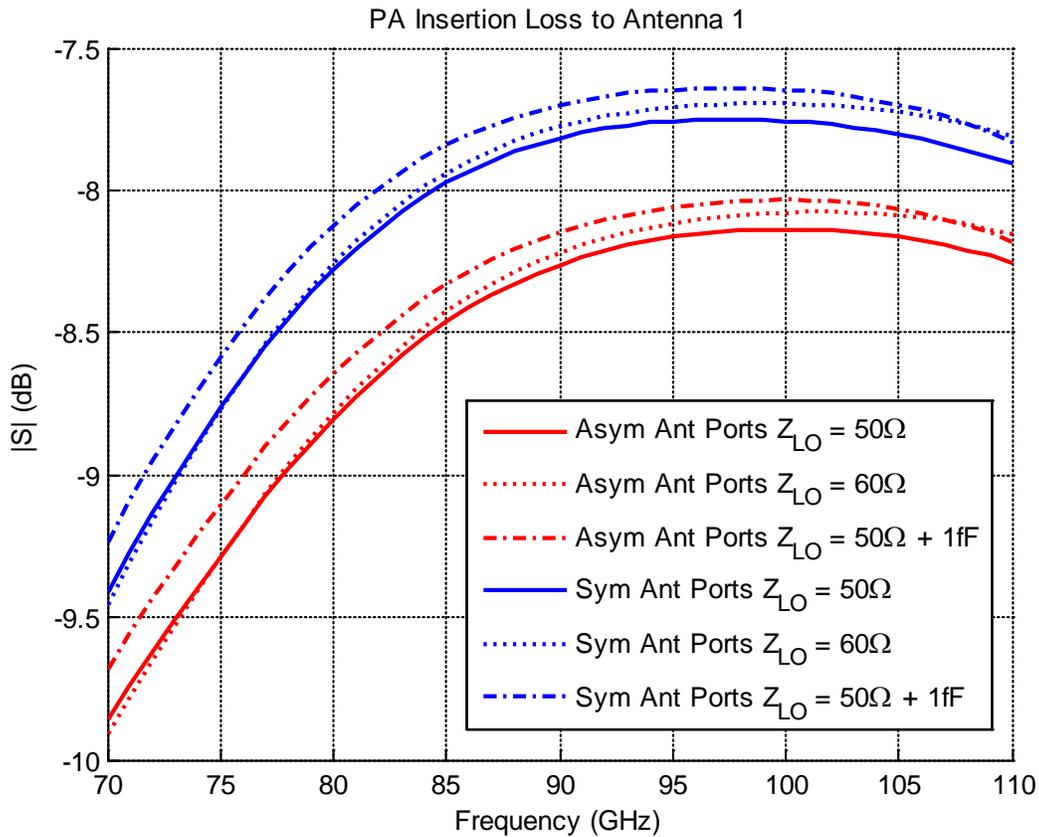


Figure 4-6: Upper PA Insertion Loss to Antenna

The major issue with introducing the asymmetries described above is that the transmitter isolation degrades severely, particularly as the port impedances become even slightly mismatched. This is illustrated in Figure 4-7, which shows what happens to the isolation in both the symmetric and asymmetric antenna feeds of Figure 4-3 and Figure 4-4 as the two LO port impedances are adjusted. The 90-degree meander line on one LO port is included in both sets of data, so variations in the LO port impedance affect each half of the circuit differently. From Figure 4-7 we see that the degradation is severe, regardless of whether the antenna ports are mismatched. However, with asymmetric antenna ports (as will actually be the case with the 90-degree antenna feed), the transmitter leakage is even worse, as shown in the red-curves.

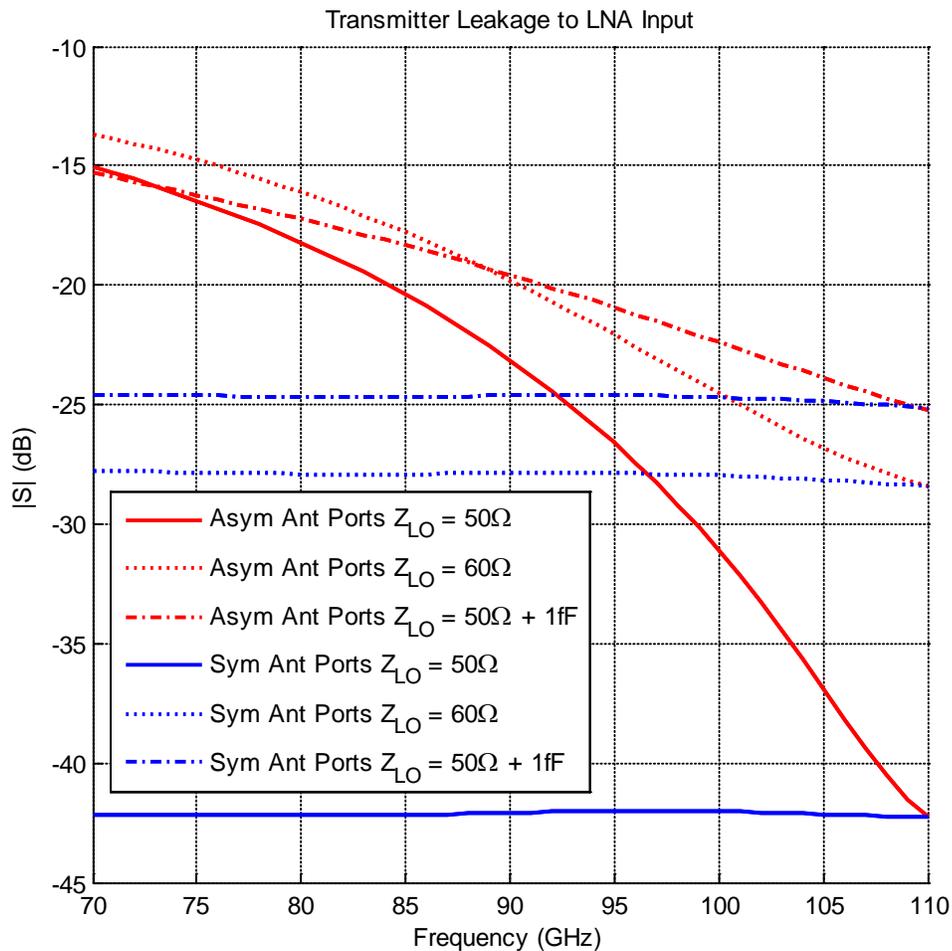


Figure 4-7: Transmitter leakage analysis under LO port impedance mismatch

Fabricated Chip

The fabricated chip showing the individual components for testing and the fully implemented unit-cell transceiver is shown in Figure 4-8. The chip was mounted and wire-bonded to a printed circuit board in order to perform on-chip probing and characterize some of the test components as well as for probing the antenna ports to examine the transceiver unit-cell. A picture of the chip mounted on the board is shown in Figure 4-9, next to 0201-size decoupling capacitors on a variety of power supplies. A power regulator and bias-current setting board was designed and assembled to provide a number of supply voltages and analog current biases to the chip. The chip was designed to allow it to be flipped onto a package or board and satisfies all pad-related design rules using the smallest pad-class. The design rules for this class of pad require large numbers of pads (two rows) to be placed at each corner of the die, as seen in Figure 4-8.

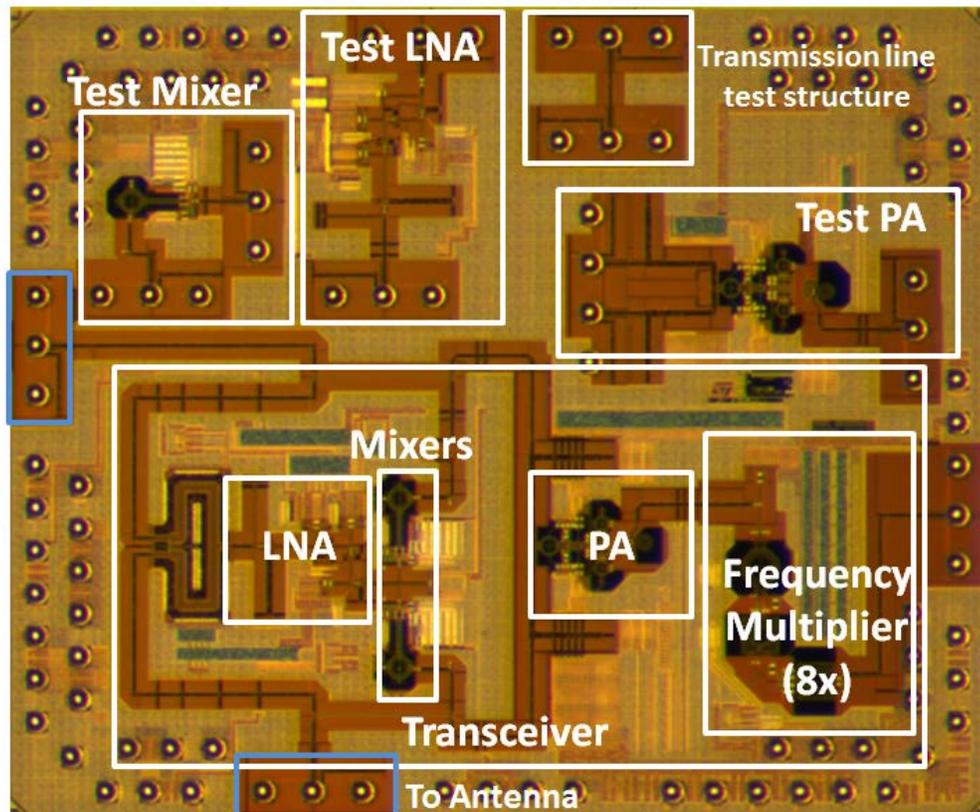


Figure 4-8: Fabricated 3x2.5mm chip showing test-components transceiver unit-cell

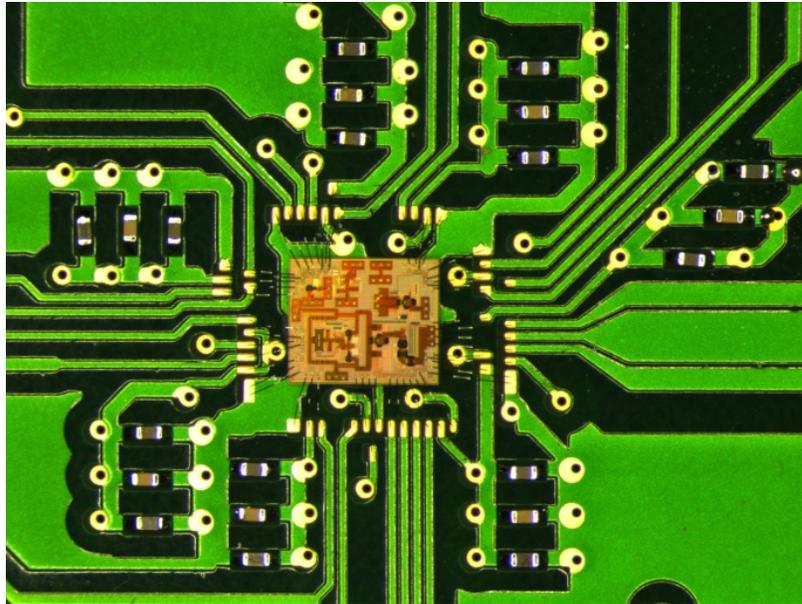


Figure 4-9: Wire-bonded chip-on-board used for probing

The transceiver unit cell exhibits strong transmitter leakage as a result of the mixer-feed and antenna-feed asymmetries, as the simulations in the previous section demonstrate. This results in a number of undesirable features of the output IF waveforms. The first is that the transmitter and receiver paths have non-constant gain versus frequency, as well as nonlinear phase versus frequency. This results in frequency-dependent leakage- and LO-paths, meaning that the output DC operating point of the mixer changes somewhat throughout the frequency ramp. Even if this did not occur, the strong transmitter leakage causes a large number of spectral components at harmonics of the ramp period to show up in the IF output spectrum, making discrimination of the ranging-signal versus those interferers a difficult task, even with post-processing. Additionally, since the mixer was not balanced with respect to the RF input, LO phase imbalance at the mixer input (meaning the nominally differential LO signals do not have perfect 0- and 180-degree phase relationship) results in static DC offset between the differential IF output leads. If an instrumentation amplifier following the mixer is used in lab measurements to take the differential IF signal and make it single-ended, the result is a large DC offset at that amplifier's output.

The leakage at the IF waveform is well-predicted by the analysis done in the first chapter of this report, and is also demonstrated by the measurements shown in Figure 4-10. The transceiver unit cell is excited around 12GHz using a demonstration FMCW synthesizer board from Hittite. One of the two antenna ports was probed using W-band GSG probes, and then coupled through a short section of cable and a coaxial-to-rectangular-waveguide-coupler to a W-band horn antenna. The data is captured for a set of target distances, showing the evolution of the waveform as the target distance changes. In fact, this is the expected behavior if the leakage term is strong, and although the transmitter leakage adds significant interference at the IF output, the radar still operates as a very sensitive interferometer, cycling through these repetitive waveforms repeatedly as the target is moved each wavelength (3mm) away from the antenna. The result is an entertaining lab demonstration, since targets can be moved very carefully using a probe manipulator, and the phase relationship of the output IF waveform evolves very smoothly and deterministically at the micron- and millimeter-levels of movement. The waveforms of Figure 4-10 are offset vertically for clarity, and their DC component—which is due to the mixer DC offset described previously—is removed. Due to high losses in the cables, connectors, and probes; as well as the fact that only half of a single unit-cell is tested at once, only a short target range is possible before the variation in the IF waveform with target distance is not readily apparent on an oscilloscope. The distance implies that less than a full waveform of the IF range data appears during half the modulation period. Therefore, extraction of the range information from the IF waveforms is difficult, even if the waveforms are post-processed knowing the time points of a frequency ramp initiation and termination. Unfortunately, the transmitter leakage of the first transceiver unit-cell implemented seemed too high to warrant an array and on-board antenna design. A second iteration of the transceiver unit cell (including the redesign of several sub-components which worked acceptably well during this first iteration) will fix the asymmetries and hopefully demonstrate significantly lower transmitter leakage. For details on the proposed changes to fix the transceiver implementation, refer to Chapter 6.

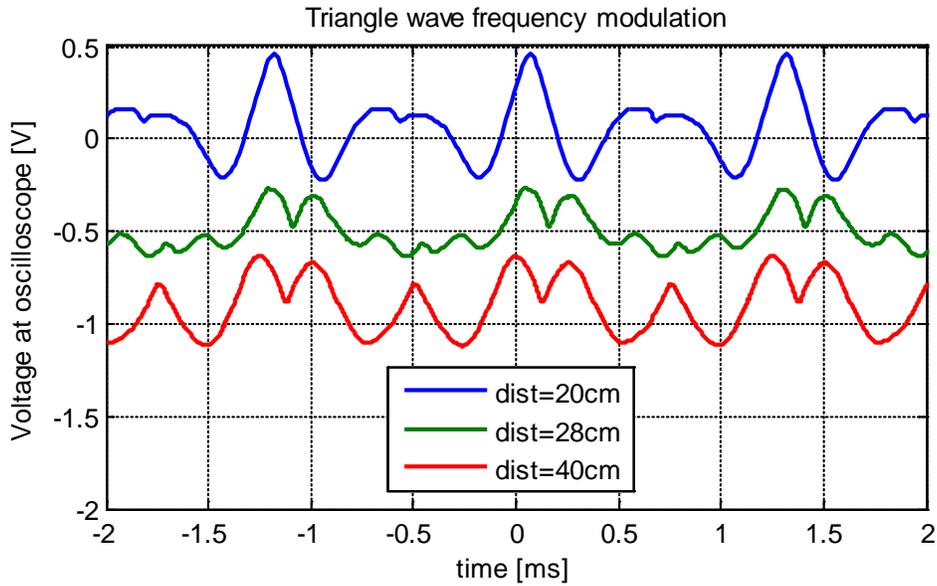


Figure 4-10: IF waveforms versus conducting target distance

For a comparison of the IF waveform shapes predicted by the simplified leakage calculation refer to Figure 4-11, which demonstrates the down-converted FMCW modulation under significant transmitter leakage. The calculation assumes ideal circuits and mixing. Qualitatively, the evolution of the IF waveforms with time delay (or target distance) agrees with calculations.

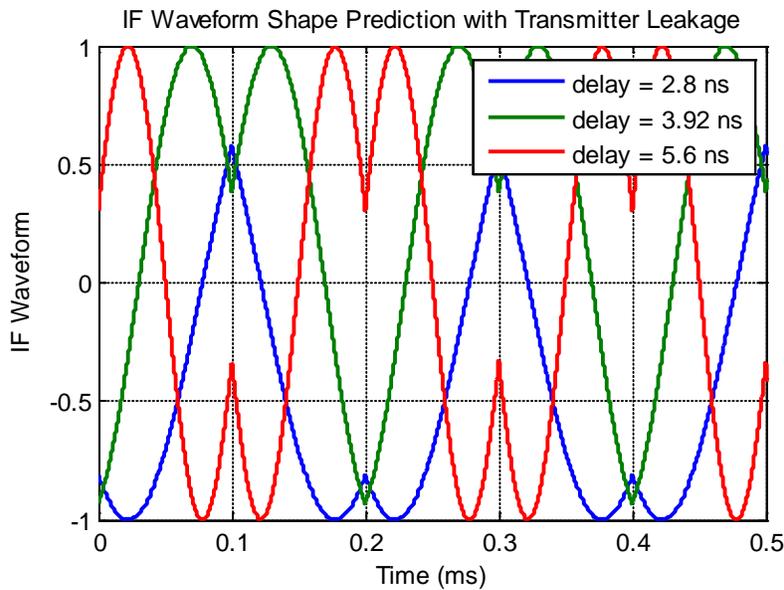


Figure 4-11: Example IF waveforms with TX leakage assuming ideal circuits

Chapter 5 High Frequency Modeling

Getting Signals off the Chip

FMCW radar systems typically use low fractional bandwidth signaling, which can allow simpler matching networks to be used: often first-order (simple “LC”) matching networks can achieve sufficient bandwidth for FMCW applications. Unfortunately, implementing matching networks compactly, aligning their implementations with the designed center-frequency, and simulating large passive structures in an electromagnetic field solver can complicate the design process of even first-order matching networks. Although a variety of more complicated methods have been devised to couple millimeter-wave signals off-chip (including aperture coupling or using on-chip antennas) this chapter seeks to show pad structures that can be used for many narrow-band millimeter-wave applications. This allows transmission line and antenna structures to be implemented on a board, which is typically much cheaper. As alluded to previously, one of the mistakes made in the design of the chip was the mismatched center frequency of the ground-signal-ground (GSG) pad structure, so this chapter clarifies the proper way to model these structures in the popular commercial simulator HFSS, and the mistakes made during the first iteration of the design.

Port Excitations in Electromagnetic Simulators

There are two basic ways to excite a ground-signal ground pad structure in an electromagnetic simulator. The air-coplanar probes (and all coplanar waveguides) that are used to measure millimeter-wave devices exhibit predominantly lateral electric fields, so one excitation that mimics this field shape is the method of exciting the center pad using lateral lumped ports extending from the two ground pads, as shown in Figure 5-1. The figure also shows a shunt transmission line to ground (which acts as an inductor) to resonate with the pad capacitance of the signal line, and a microstrip transmission line extending away from the GSG structure that connects to the circuits on-chip. The other method is shown in Figure 5-2, which shows a vertical lumped port beneath the signal pad, exciting it from the lower-metal-layer ground-plane. Additionally, the figure shows the use of a wave-port to excite the microstrip lead on-chip, although that alteration makes little difference in the results.

Ansoft training for HFSS erroneously explained that very wide lumped ports can have their own capacitance to nearby metal structures, and very long lumped ports can have their own partial inductance, since the lumped port models a sheet-current applied along its length. Using this assumption, the inductance of the lumped ports of Figure 5-1 can be calculated; a closed form expression exists for partial inductance of sheet currents given their dimensions. Then the two ports of Figure 5-1 could be connected in common-mode to model equal electrical excitations between the signal pad and each surrounding ground pad, and a capacitance or negative inductance could be added in schematic to eliminate the partial inductance contribution of the lumped ports. This turns out to be incorrect, and—if lumped ports do have any intrinsic capacitance or inductance—it is already accounted for by the HFSS simulator. Note that if the same mistake were made using the model in Figure 5-2, the calculated lumped port inductance would be smaller, and therefore the designer (erroneously) would add more capacitance to remove its effect, and the change in apparent resonant frequency would be small.

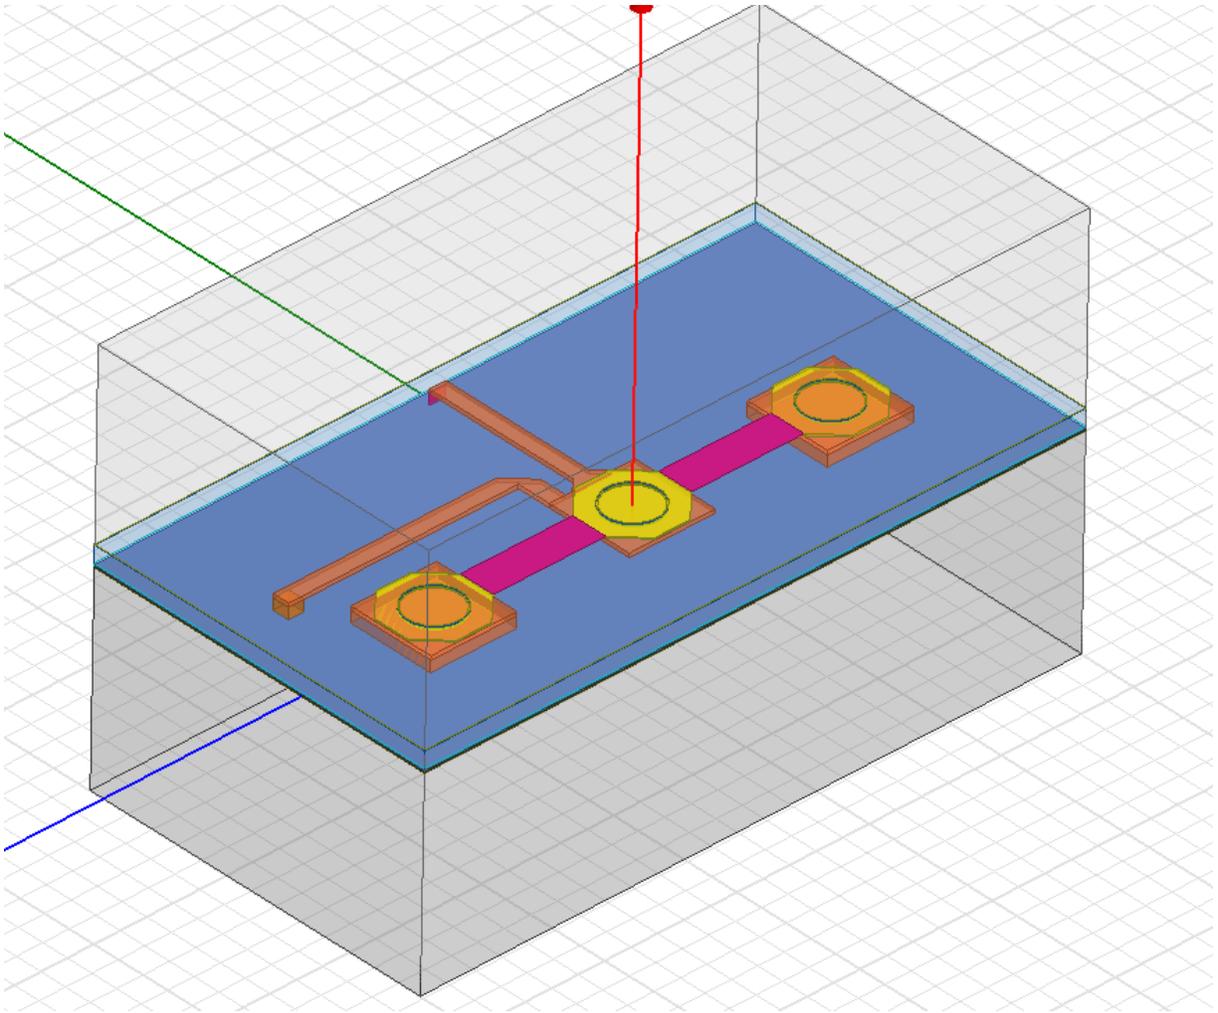


Figure 5-1: Excitation of on-chip GSG with lateral lumped ports

It turns out the correct way to model these structures is not to overthink the lumped ports, and to just use the results directly out the simulator, without accounting for any parasitic characteristics of the lumped ports. The three cases are illustrated in Figure 5-3, where we see that the raw simulation results from Figure 5-1 and Figure 5-2 are approximately equivalent (blue and black curves), whereas the method erroneously considering port parasitics (red) shows a resonance shifted 10GHz lower in frequency. Since the GSG structure was optimized to match at 94.5GHz using the latter method, the red curve is well-aligned to the intended frequency; however the actual chip shows that the 110GHz resonance is actually what resulted.

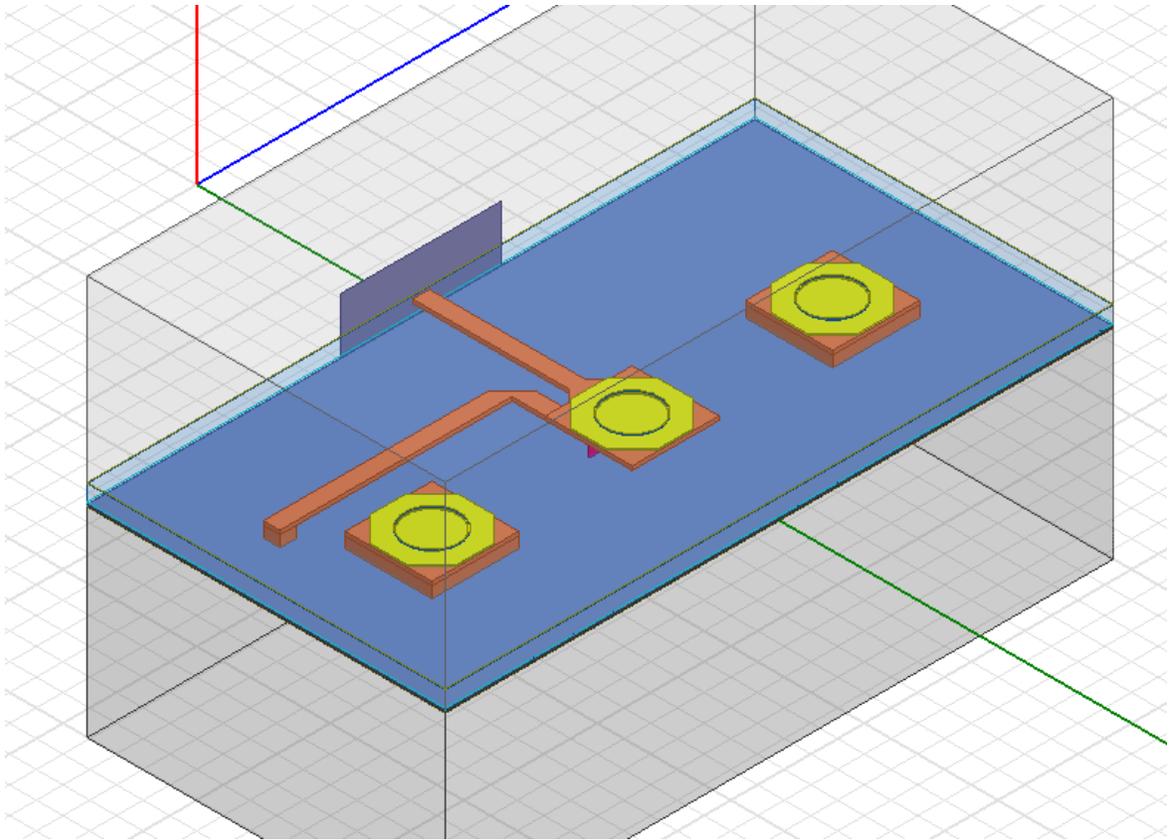


Figure 5-2: Excitation of on-chip GSG with vertical lumped port

Despite the frequency mismatch, the simulations in Figure 5-3 depict fairly low loss and wide bandwidth: the loss can be as low as 0.3dB, although realistically it will depend on the probe position and VNA calibration. Unfortunately, and as a direct result of the modeling error, the input impedance seen by the on-chip circuit from the microstrip port of the GSG structure is no longer 50Ω , and this has a negative effect on the isolation of the broadside couplers used in the radar transceiver since they were designed to be loaded with 50Ω ports. The change in impedance seen at the microstrip port is plotted in Figure 5-4, where we see that it is actually approximately $(54 + j14)\Omega$ rather than $(49 + j2)\Omega$. It also introduces mismatch in the transceiver unit-cell since the antenna leads were not made symmetric on-chip. In a full implementation of the transceiver architecture, the antenna ports will typically be mismatched anyway, due to the unfortunate need for a 90-degree delay on one antenna port but not the other.

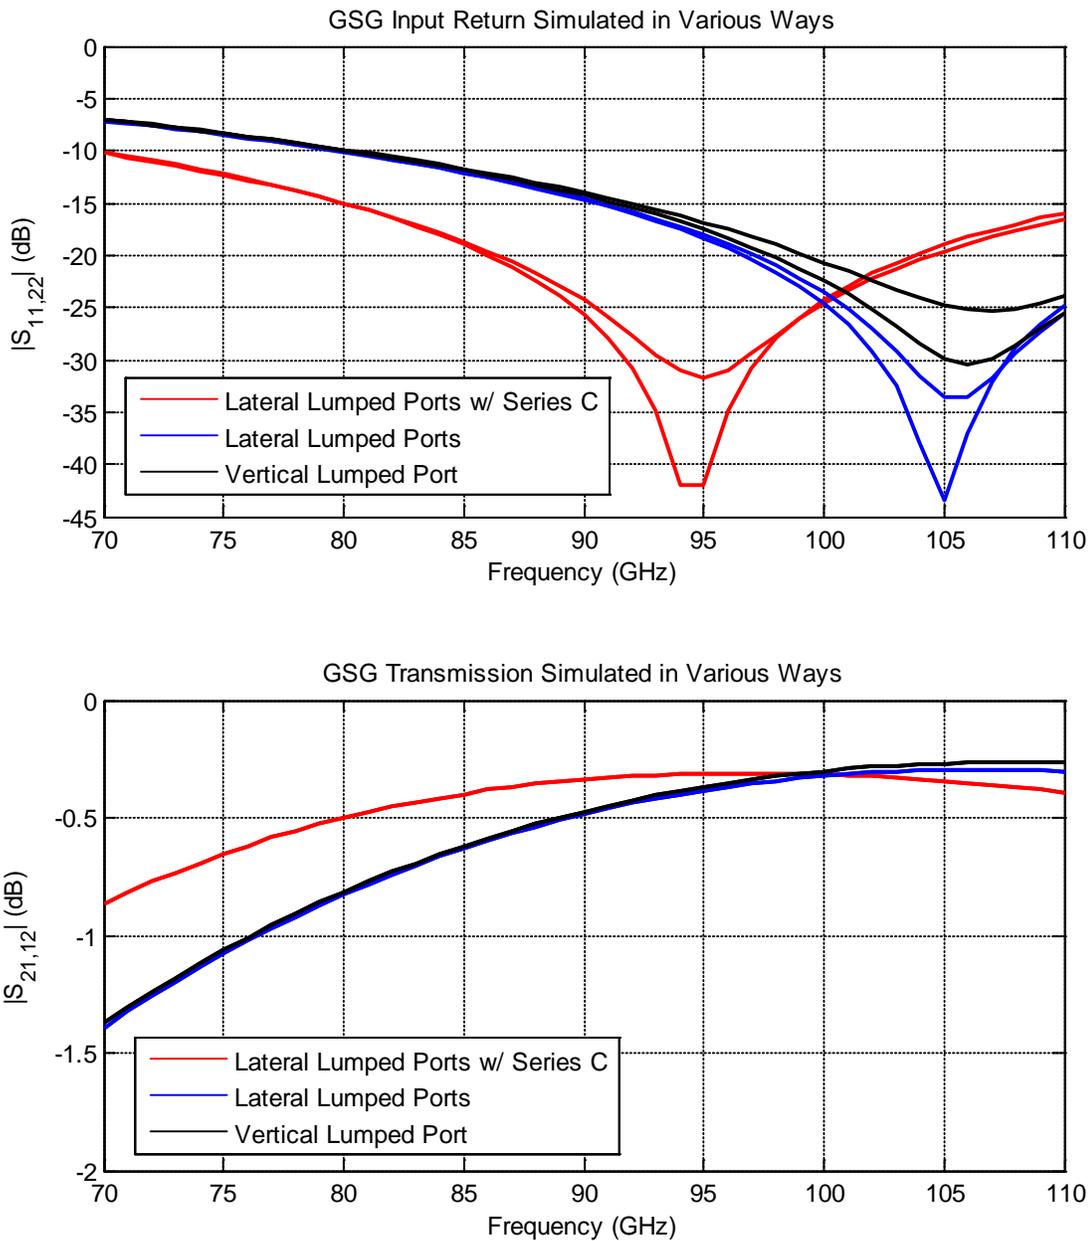


Figure 5-3: Simulated GSG structures with original and corrected modeling methods

Also of interest in Figure 5-3 is that both excitation methods of Figure 5-1 and Figure 5-2 produce very similar results (blue and black curves); in particular they produce almost identical resonant frequencies and transmission losses. This indicates that simulation results are not extremely sensitive to the excitation method used in the electromagnetic simulator.

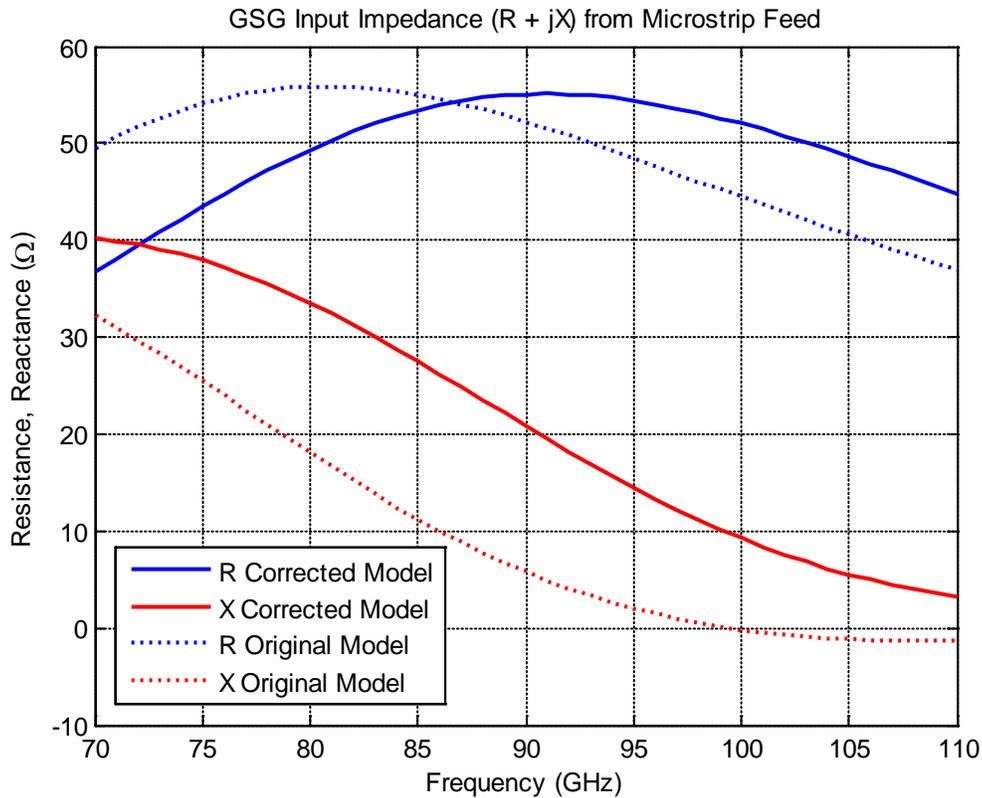


Figure 5-4: Comparison of GSG input impedance with original and corrected modeling methods

Measured Passive Structure

A connection of two GSG structures was included on-chip to form a through-line structure on which probe calibration could be checked, probe leveling could be performed, and some estimation of the passive structures' loss could be made. The measurement data is illustrated in Figure 5-5 and is the justification for why the lumped ports have no partial inductance, described in the previous section. By placing two of the "raw" GSG elements back to back, fairly good agreement in magnitude is achieved between measurement and simulation; in particular the peak resonance is closer to 110GHz than the 95GHz center frequency designed using the incorrect assumptions of the previous section. The phase agreement is less ideal; although the VNA phase measurement may be suspect at these frequencies, as illustrated above 95GHz.

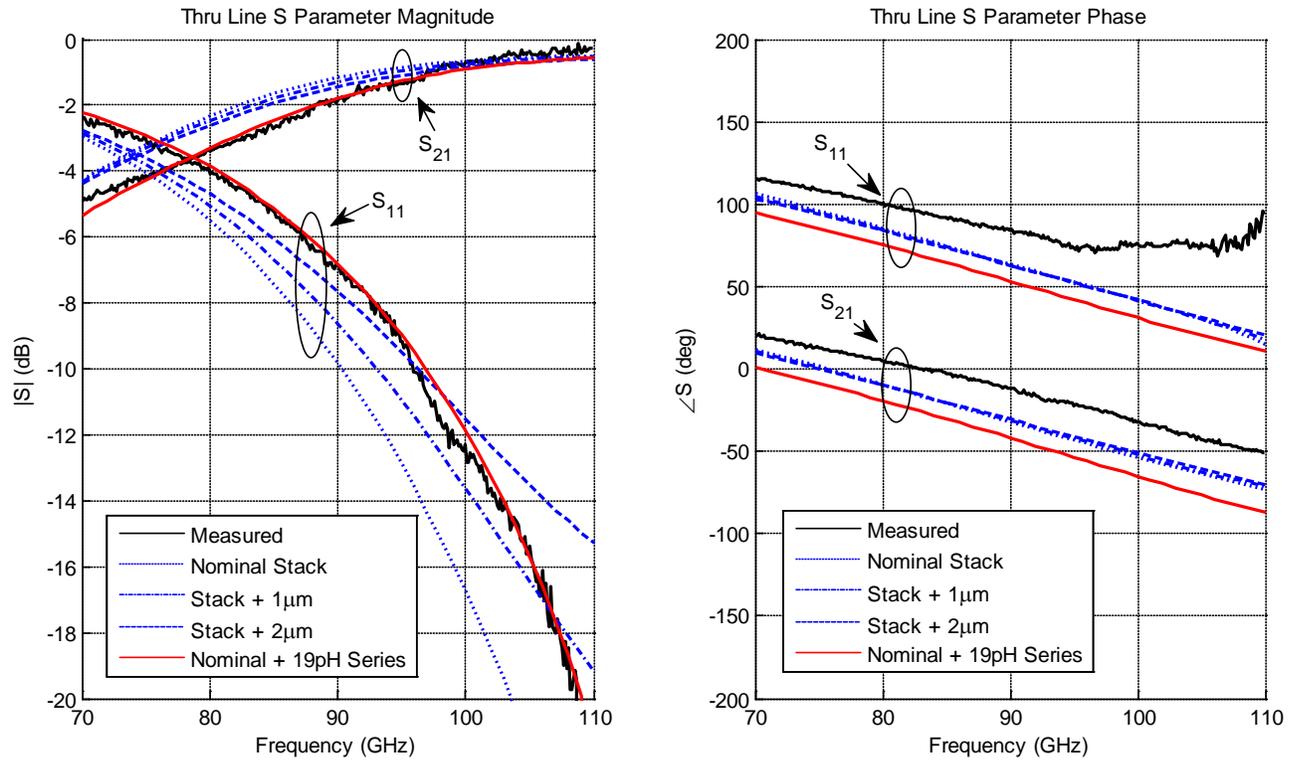


Figure 5-5: Comparison of simulated and measured thru-line GSG structures

To align the magnitude measurement and simulation perfectly, a variety of additional variations were simulated, including raising the dielectric stack thickness (the distance between the top metal and the ground plane). This is shown in blue curves of Figure 5-5, as is the effect of schematically adding series inductance to the signal lines. The latter adjustment makes the measured and simulated data agree very well when 19pH is added to each GSG structure. The explanation for this could be imperfect probe calibration, but it could also be that the simulation underestimates the inductance between the signal and ground lines. The latter makes sense since the model of the ground plane is composed of the lowest two metals and a solid metal that matches the via vertical conductivity. In the simulation the ground plane is a solid whereas to satisfy design rules, the actual ground plane has an overlapping square “checkered” grid of the lowest two metals with vias connecting them periodically. In the simulation, lateral currents can

flow in the via layer (since it is a solid metal object), whereas in reality the vias are small and vertical, and lateral currents can only flow in metal layers themselves. More investigation should first be done, but the required additional inductance to match the measurements and data suggests that an improved model of the ground plane inductance should be developed; perhaps this could be as simple as thinning the metal layer used in simulation. The actual ground plane structure is impractical to simulate in HFSS due to its high detail and the concomitant long simulation time.

Based on the frequency alignment mistakes made in the GSG structure of the first chip iteration, it has been re-tuned using the correct methodology. An alternative structure that uses capacitive stub tuning has also been designed and is shown in Figure 5-6. Here the pad capacitance is rotated around the smith chart such that a capacitive stub can be used to tune it out rather than the inductive (shorted) transmission lines used previously. Here the pad dimensions have been increased to see the effect of the larger pad-class available in the same process, and the pitch also been increased to allow low-cost boards to serve as flip-chip hosts.

The redesigned structure has the benefit that its transmission is reasonably high across a broad bandwidth, whereas the shunt inductor of the previous design increasingly shorted out signals at lower frequencies. The redesigned structure might not pass “antenna rules” if the GSG is connected directly to a device, a transistor gate, or even a MIM capacitor, whose dielectrics could break down during processing if the antenna rules are not satisfied. Still, Figure 5-7 shows that the larger pad lowers the bandwidth: the -15dB matching is only across 85-102GHz, whereas the shunt-inductor structure can achieve a -15dB matching from 80-110GHz+, as shown in Figure 5-3. There are other factors to consider when using a smaller pad-class, however, such as the bumps that will be applied (for flip-chip on-board packaging), the cost of a fine-pitch board, and the number of pads and their positioning which will be required in order to provide a robust physical connection between the chip and the board.

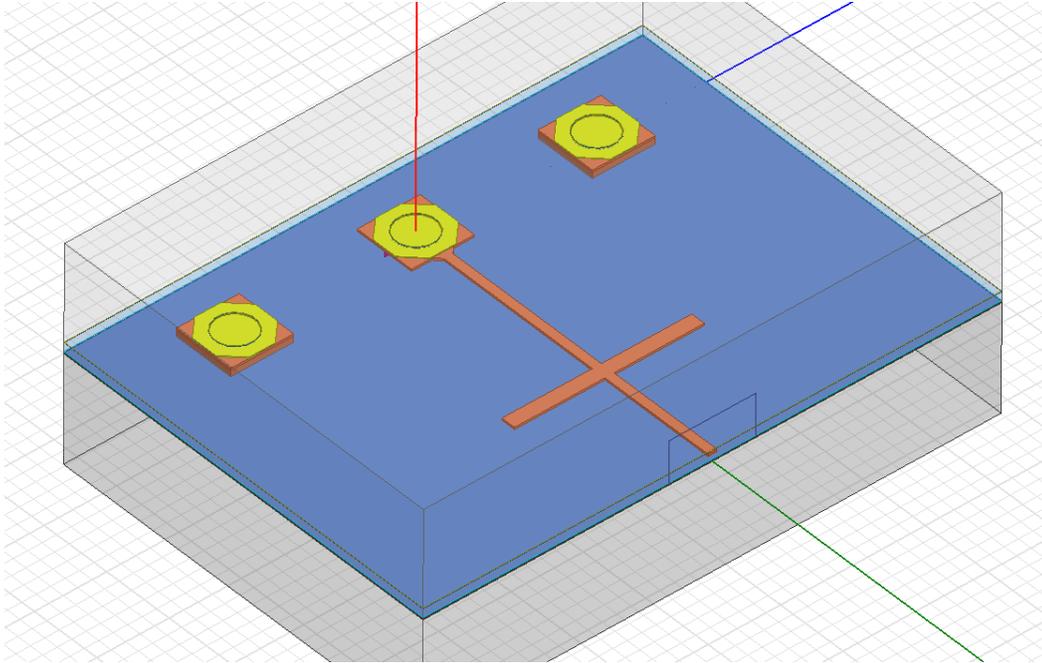


Figure 5-6: Redesigned GSG with capacitive stub tuning

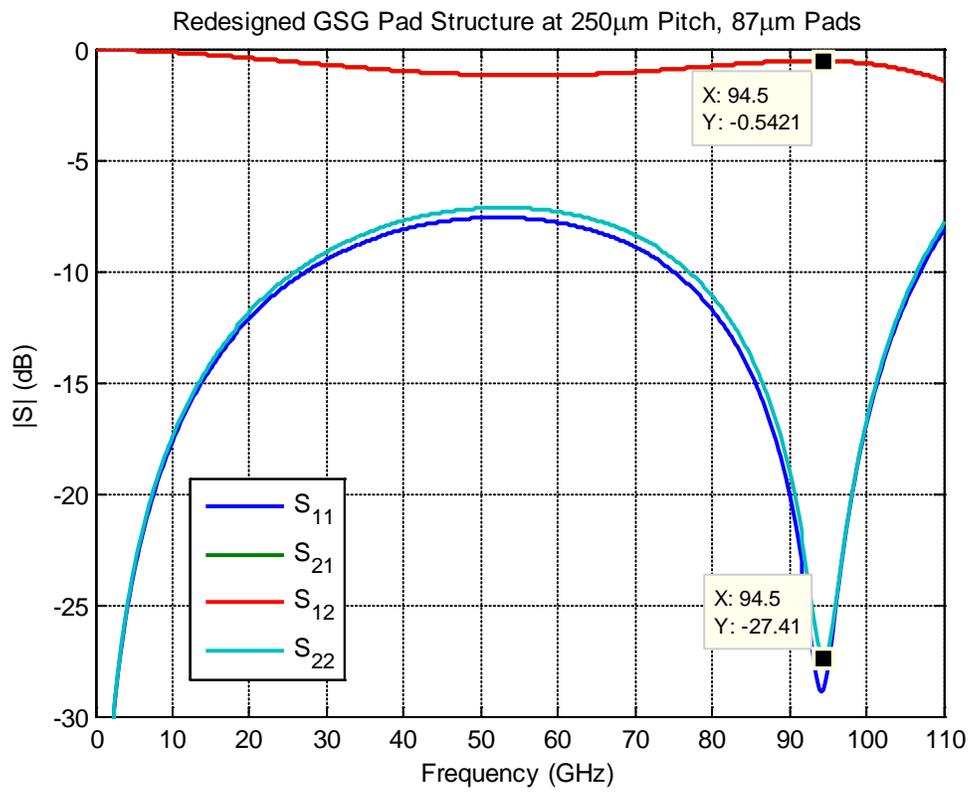


Figure 5-7: Simulation of redesigned GSG structure

Chapter 6 Conclusion

Proposed Changes

Based on the millimeter-wave design experience gained and analysis of the measurement results of the standalone components and the unit cell, a variety of suggestions for improving this architecture can be made. A large portion of the negative features of this single-antenna leakage-cancelling architecture described in Chapter 4 will still remain; in particular the system noise figure will still be very high compared to other architectures which place less loss in front of the receiver. This section summarizes proposed changes to be implemented on the second iteration of the design, briefly reiterating the issues that occurred in the first iteration and why they should be corrected.

The majority of proposed changes are architectural, and a revised architecture is depicted in Figure 6-1. To solve the asymmetry problems, the LO ports of the broadside couplers need to be symmetrically loaded through symmetric lengths of transmission line. This is impossible in the existing architecture while still maintaining in-phase and quadrature-mixers and without crossing transmission lines carrying LO/LO or LO/RF W-band signals. A better solution is to switch the quadrature generation to the RF side using either a 90-degree delay on one RF path or a quadrature hybrid at the LNA output. Additionally, to improve the termination accuracy of the LO ports of the broadside couplers, they should be terminated with buffer amplifiers rather than the mixer LO port. This allows both passive- and active-device common-mode rejection to

drive the mixers fully differentially, and it also provides symmetric termination impedances to help maintain transmitter leakage cancellation via balanced signal paths from the PA to the Wilkinson combiner. Using buffer amplifiers to drive the mixers also helps reject power-level variation in the PA output across the frequency ramp, and it allows control of the mixer LO drive strength that is independent of the transmitter drive strength and the loss of the passive structures. The buffers are actually shown as two buffers driving the LO ports of the mixers in Figure 6-1; this could in reality be implemented as a dual-input single-output tuned amplifier stage where only one signal path is biased at one time. Such a configuration would allow the receiver to accept auxiliary LO inputs from a PLL or other source other than the transmitted signal traveling through the couplers. Then if the transmitter isolation is still inadequate under perfect layout symmetry, the chip can still be used as a dedicated transmitter or receiver in a phased array simply by activating or deactivating the PA. This will allow investigation into the antenna design and packaging requirements, an area that received little attention since the first iteration was only probed on-chip but not flipped onto a board.

Most importantly for layout convenience and reducing the routing loss on the signal path, the couplers need to be changed to “backward” couplers, such as vertically coupled lines rather than horizontally coupled lines, as shown in Figure 6-1. This allows the architecture to be implemented without any of the RF, LO, or IF signals crossing, and also unwraps the signal path that completely encloses the LNA, obviating the need for power-supply and bias-current underpasses beneath the microstrip ground plane, and reducing the routing loss preceding the LNA and that between the PA and the couplers. The couplers can then be placed closer together since the LNA does not need to be surrounded by them, and therefore the matching between the two differential paths can be improved by virtue of their proximity, potentially reducing transmitter leakage. In general, the number of bends in the high-frequency lines can also be reduced, and to improve the coupler port matching, the architecture can replace the

Wilkinson coupler with a Gysel coupler, whose port impedance-matching depends less on the phase shift of resistors and their leads and more on resistor matching, which is typically good in silicon.

Other system-level changes include transitioning to a lower-frequency modulation input, since 12GHz FMCW signals are still fairly difficult (and power hungry) to generate and distribute. The 12GHz Hittite FMCW synthesizer uses at least an order of magnitude more power than the entire test chip designed for this work. Lab equipment that can implement linear frequency sweeps around 12GHz is significantly more costly and less accessible than that which operates below 3GHz. It is natural then to implement the frequency multiplication as a phased-locked loop (PLL) and not as a nonlinear frequency multiplier circuit. This alone is a significant area of future work since PLLs operating around 100GHz are nontrivial to design.

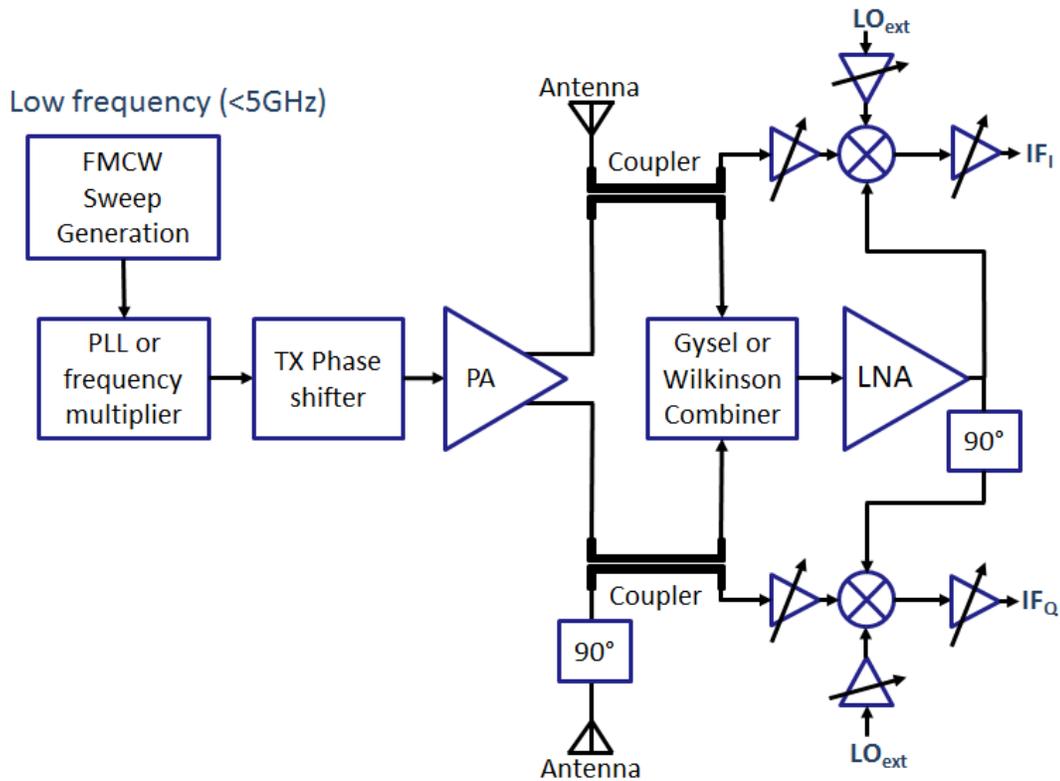


Figure 6-1: Proposed revised unit cell architecture

The IF amplifiers that implement signal summation from several array elements will need to be integrated on-chip in order to easily tile the transceiver into a phased array. The first iteration unit-cell transceiver used external amplifiers for measurements. More thorough analysis of the phase shifting resolution required should be done; the phased array analysis presented in Chapter 1 was performed for RF phase shifting in an antenna array, but it may be that a large enough array also requires only coarse IF phase-rotator resolution as well. The receiver should be redesigned with double-balanced mixers to help reduce the output offset voltage due to imperfectly differential LO drive signals. This will allow higher gain IF amplifiers to be used before their outputs saturate, and ensure that the IF amplifiers see nominally zero differential inputs due to the possibly imperfect LO drive of the mixers. Adding buffers preceding the LO ports of the mixers will also help solve the mixer output offset voltage problem by driving the mixer LO ports more differentially.

The LNA can also have slightly higher gain and lower area, and it would be preferable at the array level if all sub-circuits in the transceiver operated off of the same supply voltage. The current method for connecting the LNA to the mixers uses impedance matching which is unnecessary when they are in close proximity. Designs in progress now indicate that excess gain in the LNA can be traded off for gain flatness if the mixer presents a fixed load impedance to the LNA but the LNA is not conjugately matched at its output to that impedance. The LNA's passive power splitting also allows LO leakage signals to interfere with one another, since the two mixers are directly coupled to one another via their RF port. While this is common in many low-frequency passive mixers, it may reduce self-mixing of the frequency-modulated LO signal if each mixer is driven by separate, parallel final stages of the LNA, which are in turn driven as parallel loads to the preceding stage of the LNA. This is especially important with the current single-balanced mixer architecture, which suffers from poor LO-to-RF isolation due to the capacitive emitter coupling rather than excitation of the switching pair through a transconductor.

The GSG structure has been retuned based on the results presented in Chapter 5. The pad structure can be simplified to ease board-level routing. In particular, pad groups in the corners of the chip can correspond all to ground or all one supply voltage so that complex two-row “escape” board-level routing is unnecessary. The pad pitch should also be increased to $150\mu\text{m}$ or more for all pads (not just the high-frequency GSG structures that are probed) so that the board manufacturing for flip-chip packaging is more feasible and less expensive. Current biasing (particularly in the transmitter) should be made all-digital to reduce the number of analog pins required. Finally, beam-steering at both the transmitter and the receiver (i.e., adjustment of the IF amplifiers) needs to be digitally-controlled even if a small number of elements are placed in a phased array.

Summary

In summary, an uncommon architecture of a direct-coupled transmit/receive FMCW radar transceiver was investigated and implemented. The architecture uses a single dual polarization antenna to transmit and receive circularly polarized waves and implements a form of on-chip circulator by directing energy between common- and differential-modes on a symmetric waveguide structure. The polarimetric radar architecture has several limitations that were described in detail, and the implemented version suffers from poor transmitter isolation due to layout asymmetries of this particular implementation, and also to a lesser extent inherent to antenna asymmetries which form the basis of operation of the polarimetric radar. Suggestions for architecture improvements were described and design work has begun on a second iteration that addresses the challenges encountered with the first iteration and described herein.

Circuit building blocks operating at millimeter-wave frequencies—nominally 94GHz—were designed, laid out, measured, and analyzed. These included a low power LNA with fairly good input referred compression point; mixers and their interface with the LNA; several millimeter-

wave passive components; as well as—to a lesser extent—collaboration on the transmitter which was designed by another student. The LNA achieves around 14dB of gain, more than 22GHz of bandwidth, a lowest noise figure of 7.5dB while consuming 20mW. This makes it comparable to published results in the 77-120GHz frequency range, and extremely competitive in bandwidth. An asymmetric balun structure that achieves higher CMRR while using less center-tap capacitance was also presented, simulated, and implemented in the LO distribution network of the radar transceiver.

The design process served as a wonderfully immersive learning experience for designing high-frequency integrated circuits. The experience allowed development of the knowledge necessary to design circuits that operate at millimeter-wave frequencies, to synthesize new high-frequency passive structures and interface them with modern transistors, and to perform high frequency measurements to feed back and update the design techniques that should be used on future designs.

Bibliography

- [1] A. Natarajan and S. Reynolds, "A fully-integrated 16-element phased-array receiver in SiGe BiCMOS for 60-GHz communications," *Solid-State Circuits, IEEE Journal of*, vol. 46, no. 5, pp. 1059–1075, 2011.
- [2] S. Kim and G. Rebeiz, "A low-power BiCMOS 4-element phased array receiver for 76–84 GHz radars and communication systems," *Solid-State Circuits, IEEE Journal of*, vol. 47, no. 2, pp. 359–367, 2012.
- [3] L. Wang and S. Glisic, "A single-ended fully integrated SiGe 77/79 GHz receiver for automotive radar," *Solid-State Circuits, IEEE Journal of*, vol. 43, no. 9, pp. 1897–1908, 2008.
- [4] S. Reynolds and B. Floyd, "A silicon 60-GHz receiver and transmitter chipset for broadband communications," *Solid-State Circuits, IEEE Journal of*, vol. 41, no. 12, pp. 2820–2831, 2006.
- [5] S. T. Nicolson, S. Member, K. H. K. Yau, S. Pruvost, V. Danelon, P. Chevalier, P. Garcia, A. Chantre, S. Member, B. Sautreuil, and S. P. Voinigescu, "A low-voltage SiGe BiCMOS 77-GHz automotive radar chipset," in *Microwave Theory and Techniques, IEEE Transactions on*, 2008, vol. 56, no. 5, pp. 1092–1104.
- [6] N. Demirel and R. Severino, "Millimeter-wave chip set for 77–81 GHz automotive radar application," in *New Circuits and Systems Conference (NEWCAS), 2011 IEEE 9th International*, 2011, pp. 253–256.
- [7] J. Lee and Y. Li, "A fully-integrated 77-GHz FMCW radar transceiver in 65-nm CMOS technology," *Solid-State Circuits, IEEE Journal of*, vol. 45, no. 12, pp. 2746–2756, 2010.
- [8] K. Schmalz and W. Winkler, "A subharmonic receiver in SiGe technology for 122 GHz sensor applications," *Solid-State Circuits, IEEE Journal of*, vol. 45, no. 9, pp. 1644–1656, 2010.
- [9] R. Pilard and D. Gloria, "94 GHz silicon co-integrated LNA and Antenna in a mm-wave dedicated BiCMOS technology," in *Radio Frequency Integrated Circuits Symposium (RFIC), 2010 IEEE*, 2010, pp. 83–86.
- [10] A. Niknejad and H. Hashemi, *mm-Wave silicon technology: 60GHz and beyond*. 2008.
- [11] G. M. Brooker, "Mutual Interference of Millimeter-Wave Radar Systems," *IEEE Transactions on Electromagnetic Compatibility*, vol. 49, no. 1, pp. 170–181, Feb. 2007.
- [12] M. Schneider, "Automotive radar—status and trends," *German microwave conference*, pp. 3–6, 2005.

- [13] J. May and G. Rebeiz, "A W-band SiGe 1.5 V LNA for imaging applications," in *Radio Frequency Integrated Circuits Symposium, 2008. RFIC 2008. IEEE*, 2008, pp. 241–244.
- [14] B. Floyd, "V-band and W-band SiGe bipolar low-noise amplifiers and voltage-controlled oscillators," in *Radio Frequency Integrated Circuits (RFIC) Symposium, 2004, IEEE*, 2004, pp. 295–298.
- [15] V. Subramanian and G. Boeck, "60 GHz SiGe LNA," *2007 14th IEEE International Conference on Electronics, Circuits and Systems*, pp. 1209–1212, Dec. 2007.
- [16] B. Dehlink, H. Wohlmuth, and K. Aufinger, "A low-noise amplifier at 77 GHz in SiGe: C bipolar technology," in *Compound Semiconductor Integrated Circuit Symposium, 2005. CSIC '05. IEEE*, 2005, pp. 287–290.
- [17] S. Nicolson and S. Voinigescu, "Methodology for Simultaneous Noise and Impedance Matching in W-Band LNAs," *2006 IEEE Compound Semiconductor Integrated Circuit Symposium*, no. 5, pp. 279–282, Nov. 2006.
- [18] S. P. Voinigescu, M. C. Maliepaard, J. L. Showell, G. E. Babcock, D. Marchesan, M. Schroter, P. Schvan, and D. L. Hareme, "A scalable high-frequency noise model for bipolar transistors with application to optimal transistor sizing for low-noise amplifier design," *IEEE Journal of Solid-State Circuits*, vol. 32, no. 9, pp. 1430–1439, 1997.
- [19] K. Sarabandi, "A radar cross-section model for power lines at millimeter-wave frequencies," *IEEE Transactions on Antennas and Propagation*, vol. 51, no. 9, pp. 2353–2360, Sep. 2003.
- [20] C. Kim, J. Kim, D. Baek, and S. Hong, "A circularly polarized balanced radar front-end with a single antenna for 24-GHz radar applications," *Microwave Theory and Techniques, IEEE Transactions on*, vol. 57, no. 2, pp. 293–297, 2009.