Dynamically Controllable Integrated Radiation and Self-Correcting Power Generation in mm-Wave Circuits and Systems

Thesis by

Steven M. Bowers

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2014

(Submitted September 27, 2013)

For Mom, Dad, Chris, and Baby Bowers

Acknowledgments

"Success is not a random act. It arises out of a predictable and powerful set of circumstances and opportunities."

-Malcolm Gladwell

This is a quote from one of my favorite books, *Outliers: The Story of Success*, by Malcolm Gladwell. The basic theme of the book is that great things are done by intelligent people who are presented with great opportunities. To be sure, people need to take those opportunities, but in almost all cases, there are great people supporting those that end up getting the credit. I believe that often success stories are as much about the people achieving success as about those that support them, and in educational success stories, it is the teachers who may not always get the credit they deserve. For me personally, the multitude of teachers in my life have shaped my intellect, my ideologies and my personality, and I am grateful for everything they have done for me. I have worked hard to achieve what I have, but an equal amount of credit goes out to those that gave me the opportunity to do something great. There are more of them than I can possibly hope to thank personally here, but I hope to acknowledge those that had the greatest impact on my life in the paragraphs to follow.

It may seem trite, but the person who had the single biggest contribution to my personal ideologies, what I consider to be the core of who I am, outside of my immediate family was my kindergarten teacher Mrs. Munemo. She taught me the value of true empathy; taking a moment to stop and consider exactly what it might be like to be in the position of someone else. This ability to empathize with those I disagree with has been invaluable to me on many occasions, and has enabled me to identify mistakes in judgement I have made early, and to appreciate all of the opportunities that I have been lucky enough to receive. She may have given me an unsatisfactory mark in skipping (I have practiced, and truly excel at skipping now), but I remember her for her firm insistence that if a judgement that is to be made about a person, it must be based upon their merits and not the circumstances of their upbringing, cultural background, or socioeconomic status.

I was privileged growing up to be able to interact with fantastic teachers not just at school, but also through family friends. Two of the key people who fostered my interest in science were Drs. Evelyn Hu and David Clark. At the time, they were colleagues of my parents at UCSB, and some of the kindest and most genuine people I have ever met. They would give me science-related gifts for my birthday and after I completed the

normal course of whatever the experiment was, would ask what else I could do with the components from the kit. If there was a beginning to my research career, this was it. They helped foster a curiosity that would not accept any answer "because that's the way it is," but rather required a logical explanation for anything in my world that didn't fit the reality my mind expected. There were many other teachers and mentors in my personal life growing up, but I would like to personally thank: Ann Blanton and Gary Maxwell, Colleen Kegg and Bob Simon, Karen Kegg and Tim Johnson, and my best friends in grade school, Andrew Maxwell, Jeff Simon, and Andy Johnson.

Since I was little, I knew I wanted to be an electrical engineer like my Dad (more on that later), but I first knew I wanted to be a teacher my junior year in High School in my AP physics class with Mr. Abo-Shaer. It was his first year at Dos Pueblos High School, and I have never seen such a passion for teaching others before or since then. We had the usual discussion about the curriculum of the class, but I will remember him most for our discussions about the way he taught the first year physics course. AP physics was the first class I took where the teacher also taught a lower level class. Being in the AP class, I was able to discuss with Mr. Abo-Shaer the challenges of being a first-year teacher, and teaching the lower-level class for the first time. It was the first time that I was able to see the impact that a teacher could have on his students from the perspective of someone who had already taken the class. A great teacher will always find a way of showing their students their full potential, and motivating them to achieve it, rather than focusing on their shortcomings and failures.

Between my junior and senior year of High School, I had the privilege of being part of the apprentice researcher's program at UCSB run by Wendy Ibson, a program where local high school students worked with graduate students on cutting edge research. I worked in the Physics Department on onset voltages of chaotic patterns in nematic liquid crystals. Wendy provided the necessary glue that enabled the inexperienced high school students to interact with the generous graduate students. I had the opportunity to present my research at the 2002 Southern California Conference for Undergraduate Research at Caltech, so in the end the program introduced me to three things: academic research, conference publications and presentations, and Caltech, which years later would become my home.

I went to UCSD for my undergrad and loved it there. The professors seemed truly eager to teach, and cared about the amount their students learned, not just how they performed on a test. There were three professors there that really made a lasting impression on me: Drs. Joseph Ford, Sadik Esner, and Gabriel Rebiez. Dr. Ford had an optics lab class, and as far as I can gather, the point of the class was to allow advanced students to use incredibly expensive optical equipment normally reserved for research labs in a classroom setting to give them a feel for the types of things they could accomplish through advanced research. Dr. Esner gave me my first opportunity at the undergraduate level to do independent research. I worked with graduate student Dr. Bing Shao on a portable microscope setup with optical chopping that could be loaded onto an expedition boat to enable in-situ imaging to measure angularly-resolved light scattering (ARLS) to characterize submicron cells in three dimensions. Drs. Esner and Shao gave me my first experience at independent research, and

convinced me that pursuing a Ph.D. in graduate school was the path I wanted to pursue. One might take notice that the first two of my undergraduate mentors are in the field of photonics, but my graduate research has mainly focused on mm-wave circuit design. While I loved the work I was doing in photonics, a microwave circuit class I took with Dr. Rebeiz my junior year at UCSD peaked my interest, and changed my research trajectory. Dr. Rebeiz cared passionately for his students, and would always make time for those interested to ask questions or learn about subjects from the class or from his group's research.

After college I started my graduate studies at Caltech, and was fortunate enough to step into a lab with an amazing group of students. My sincerest gratitude goes out to all of the students of the CHIC and MICS lab including Aydin, Yu-Jiu, Hua, Ed, Jay, Jenny, Joe, Amir, Alex, Behrooz, Firooz, Costis, Matt, Meisam, Manuel, Saman, Mayank, and Krishna, who early on helped me get up to speed on the state of the art in IC design, and later on who helped me reach deadlines that often seemed impossible. I need to thank both Kaushik (K1) and Kaushik (K2) for always being willing to grab a cup of coffee on the balcony, and Juhwan and Florian for always having words of encouragement if research wasn't going quite as planned. I was also lucky enough to work with both Kaushiks, Amir, Alex, Behrooz and Firooz as collaborators on parts of the research presented in this thesis, and thank them for all the long hours we put in to get results. I was able to TA for Dr. Behnam Analui when he was a guest lecturer at Caltech, and learned a great deal from our discussions about lecturing at the college and graduate levels.

When I first was talking to my advisor, Prof. Ali Hajimiri, about coming to Caltech, and also deciding about pursuing microwave circuits verses optical electronics, I remember asking what I was worried would be a dangerous question. There is a large amount of iterative work being done in circuits, where increased performance is gained by redesigning an old well-known architecture in a new advance process. I asked Ali if he thought there was still room for truly innovative circuit design, and after some careful consideration, he replied that he had no interest in iterative designs, and that in order to succeed in the group, novel designs were going to be critical. He would be willing to take on the risks of a student learning an entirely new area of research that the group was not an expert on in order to produce truly innovative work, but had no time for porting old designs to new processes. This is a mentality that I really agree with, though it is also very risky. In a field where world-record performance specifications are what get noticed (sometimes even above innovative design), passing up a sure thing that is less interesting to swing for the fences on new designs can backfire, but graduate school is exactly the time to take large risks, especially in your research, and to strive to achieve lasting contributions to the field. I also owe a large debt of gratitude to him for giving me the freedom to pursue the path of research that interested me the most, and to allow me to take projects I was working on in new and interesting directions, rather than sticking to a prearranged project plan set out early on.

My thanks also go out to the rest of my thesis committee: Prof. David Rutledge, Prof. Azita Emami, Dr. Sandy Weinreb, Prof. Hyuck Choo, and Prof. Kerry Vahala. My discussions with Prof. Rutledge always reveal a new insightful way to think about the problem I am facing. Prof. Emami always has shown interest in my research and her encouragement throughout my graduate studies has been incredibly valuable. Dr.

Weinreb's passion for research is inspiring, and his measurement expertise has been beneficial on several occasions. I thank Prof. Choo for his friendly encouragement and advice, and Prof. Vahala for his invaluable discussions as I took on a new field of research in silicon photonics toward the end of my studies. I would also like to thank Prof. Volnei Pedroni, who has always showed great interest and enthusiasm in my research.

A portion of this work was done as part of a project with IBM T.J. Watson laboratory, and I want to thank everyone there that we worked with, including José Tierno, Dan Friedman, Ben Parker, and Herschel Ainspain for their efforts throughout the project.

There are a great many people who are required to keep a research lab operating, many of whom likely don't get nearly the credit they deserve, but many many thanks go out to everyone who enabled me to do the research presented in this thesis, including Carol, Tanya, Estella and Michelle. I think that our lab would have imploded several times over in the time I have been there if not for Michelle's efforts to keep everything up and running.

Finally, I need to thank my family for all the support they have given me, which has always been unwavering. I am incredibly fortunate to have grown up with two excellent role models as my parents. They've always done the perfect job of striking the balance of encouraging me to succeed while letting me learn and even fail on my own. My Dad stoked my interest in technology at a young age. We always were building something, whether it was a train set, or a robotics kit or a go kart or a club house. Building things was something that I took somewhat for granted, as it had always been a part of my life, but after coming to Caltech, I have realized what a valuable skill set it is to be able to look at an object, figure out its functionality, and if broken, be able to fix it. My Dad is an engineer, and from early on I knew that I wanted to be one as well. Now that I have seen the inner workings of academia, I marvel that he was able to spend so much time with us as kids, but I am incredibly grateful for every moment of it. More than anything else though, I am grateful that he has taught me what it means to be a great father, and I hope to do at least half as good a job in that role as he did for me.

My Mom has always shown me what the value of hard work is. At a young age, she taught me to be a musician, and I was able to see first hand the benefits of practice and hard work. She gave me the choice of what I wanted to learn, but instilled within me the work ethic to ensure that I was successful at it. She has always been someone who I can talk to about anything, and has always been there to listen. She has also shown me that there is more to life than just work, and that striking a balance between different aspects of life is a key to staying happy. I thank her most for her unconditional love. My brother Eric has also always been a role model I can look up to. As the older brother, he would always be the trailblazer of what we could get away with as kids, and he made staying on the good side of our parents much easier than it would have been otherwise. I know I can always count on him, and that he will be there for me if I need him. I also want to express my gratitude to my late Grandpa Bowers, who had a policy of always trying to learn at least one thing every day. He taught me to have a passion not just for knowledge, but to take pleasure in the discovering of new knowledge and learning new things, and for that I am grateful. Both of my Grandmothers

and my late Grandpa Durbahn also deserve thanks for helping me grow and mature as a person. Though not strictly family by blood, Jason Wiese has been like a brother to me since college, and his support has been instrumental to my ability to complete this degree, and for that I am grateful.

My wife Christine has kept me sane throughout my studies at Caltech. If there is anyone who truly knows the hours that were put in to produce this work it is her. There were endless long nights, early mornings, and weekend plans that were canceled to meet the dreaded tapeout deadlines, and yet she was always there supporting me regardless of the hour of the night or the wrenches I threw into her plans. One of my single lab mates once asked me how I could still get all my work done while being married, and to this day I wonder how any of them can get all of their work done without the support I have gotten every night when I come home. Her understanding seems to know no bounds, though I always seem to find ways of pushing it further. She has taught me that it is ok to fail, if you learn from that failure and use that knowledge to help you succeed in the future. She is now carrying our first child, and I cannot think of anyone I would rather be the mother of my children than her.

There are many more people who have enabled me to succeed in this endeavor, and though there is not room to thank every one of them individually, I am deeply thankful to everyone who has helped me along the way.

Life has meaning when you assign meaning to it. For me, I hope to help those in my community, provide a life of opportunities for my family, leave a lasting impact on the scientific world though my research, and to pay forward what was given to me by teaching the next generation of scientists and engineers. I am truly thankful for the army of support I have received every step along the way, and am grateful for the opportunities that support has granted me.

Abstract

This thesis presents novel design methodologies for integrated radiators and power generation at mm-wave frequencies that are enabled by the continued integration of various electronic and electromagnetic (EM) structures onto the same substrate. Beginning with the observation that transistors and their connections to EM radiating structures on an integrated substrate are essentially free, the concept of multi-port driven (MPD) radiators is introduced, which opens a vast design space that has been generally ignored due to the cost structure associated with discrete components that favors fewer transistors connected to antennas through a single port.

From Maxwell's equations, a new antenna architecture, the radial MPD antennas based on the concept of MPD radiators, is analyzed to gain intuition as to the important design parameters that explain the wide-band nature of the antenna itself. The radiator is then designed and implemented at 160 GHz in a 0.13 μm SiGe BiCMOS process, and the single element design has a measured effective isotropic radiated power (EIRP) of +4.6 dBm with a total radiated power of 0.63 mW.

Next, the radial MPD radiator is adapted to enable dynamic polarization control (DPC). A DPC antenna is capable of controlling its radiated polarization dynamically, and entirely electronically, with no mechanical reconfiguration required. This can be done by having multiple antennas with different polarizations, or within a single antenna that has multiple drive points, as in the case of the MPD radiator with DPC. This radiator changes its polarization by adjusting the relative phase and amplitude of its multiple ports to produce polarizations with any polarization angle, and a wide range of axial ratios. A 2x1 MPD radiator array with DPC at 105 GHz is presented whose measurements show control of the polarization angle throughout the entire 0° through 180° range while in the linear polarization mode and maintaining axial ratios above 10 dB in all cases. Control of the axial ratio is also demonstrated with a measured range from 2.4 dB through 14 dB, while maintaining a fixed polarization angle. The radiator itself has a measured maximum EIRP of +7.8 dBm, with a total radiated power of 0.9 mW, and is capable of beam steering.

MPD radiators were also applied in the domain of integrated silicon photonics. For these designs, the driver transistor circuitry was replaced with silicon optical waveguides and photodiodes to produce a 350 GHz signal. Three of these optical MPD radiator designs have been implemented as 2x2 arrays at 350 GHz. The first is a beam forming array that has a simulated gain of 12.1 dBi with a simulated EIRP of -2 dBm. The second has the same simulated performance, but includes optical phase modulators that enable

two-dimensional beam steering. Finally, a third design incorporates multi-antenna DPC by combining the outputs of both left and right handed circularly polarized MPD antennas to produce a linear polarization with controllable polarization angle, and has a simulated gain of 11.9 dBi and EIRP of -3 dBm. In simulation, it can tune the polarization from 0° through 180° while maintaining a radiated power that has a 0.35 dB maximum deviation from the mean.

The reliability of mm-wave radiators and power amplifiers was also investigated, and two self-healing systems have been proposed. Self-healing is a global feedback method where integrated sensors detect the performance of the circuit after fabrication and report that data to a digital control algorithm. The algorithm then is capable of setting actuators that can control the performance of the mm-wave circuit and counteract any performance degradation that is observed by the sensors. The first system is for a MPD radiator array with a partially integrated self-healing system. The self-healing MPD radiator senses substrate modes through substrate mode pickup sensors and infers the far-field radiated pattern from those sensors. DC current sensors are also included to determine the DC power consumption of the system. Actuators are implemented in the form of phase and amplitude control of the multiple drive points.

The second self-healing system is a fully integrated self-healing power amplifier (PA) at 28 GHz. This system measures the output power, gain and efficiency of the PA using radio frequency (RF) power sensors, DC current sensors and junction temperature sensors. The digital block is synthesized from VHDL code on-chip and it can actuate the output power combining matching network using tunable transmission line stubs, as well as the DC operating point of the amplifying transistors through bias control. Measurements of 20 chips confirm self-healing for two different algorithms for process variation and transistor mismatch, while measurements from 10 chips show healing for load impedance mismatch, and linearity healing. Laser induced partial and total transistor failure show the benefit of self-healing in the case of catastrophic failure, with improvements of up to 3.9 dB over the default case. An exemplary yield specification shows self-healing improving the yield from 0% up through 80%.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Wireless electromagnetic (EM) radiation affects almost all aspects of human life. From low frequency radio waves used to broadcast content to remote locations, to optical light that allows us to see the word around us, to the high end of the spectrum and waves like X-rays that can be used to see through biological tissue and help diagnose injuries, EM radiation is used in innumerable applications. At low radio wave frequencies, large passive antennas are used to radiate EM energy. The size of these antennas is proportional to the wavelength of the radiated wave, so as the frequency of the radiation increases, the size of the antenna decreases. In the high mm-wave and THz frequency range, the size of these antennas become on the order of millimeters, and even microns, and the integration of these antennas onto an integrated circuit substrate becomes feasible. As the dimensions of the antenna became lower than that of the substrate, the tendency has been to take existing antenna designs that were optimized under the constraints of discreet design and to implement them on-chip. However, integration of antennas on chip is not trivial, as the affect of having a lossy substrate with high dielectric constant forces design tradeoffs that until recently resulted in antenna efficiencies that were too low to be widely utilized, and many of the assumptions used to optimize discreet antennas are no longer valid in an integrated setting.

Our main research interest lies in increasing the level of integration of antennas and developing new architectures and design methodologies that take advantage of the new possibilities that arise when an antenna is integrated onto the same substrate as its driver circuitry. However, as the frequency increases toward the f_{max} of the transistor, or the maximum frequency where small signal power gain can occur, the ability of the driver circuitry to create power, and the ability of the antennas to radiate that power diminishes. These two limitations, efficient antennas and power generation, limit the upper frequency that can be produced by mm-wave devices. This thesis advances the state of the art in both categories, by utilizing available integration techniques with a holistic design strategy that considers antenna, electromagnetic and circuit designs together in a single design space, rather than isolating them and designing their components separately.

For the purposes of this work, we will refer to an antenna as the passive structure that when driven by one or more driving sources, radiates electromagnetic fields. We will refer to a radiator as the combined unit of an antenna along with the driving sources. A key insight into the benefit of integrating the entire radiator on chip

comes from an examination of the cost structure of an integrated circuit, and how it differs from a discreet design. In discreet designs, every single component, and especially the transistors, are expensive. Connections between components are likewise expensive, both in terms of monetary cost as well as performance cost, and thus, the total number of transistors and connections between those transistors should be kept to a minimum to keep performance high and costs low. However, in an integrated environment, this entire design paradigm is completely false, as the cost of the chip is determined solely by its area, and because for mm-wave circuits, the passive structures dominate the area, so the incremental cost of a transistor is, for all intents and purposes, free. Connections between the transistors are also free, so a single connection between the driver circuits and the radiating antennas is no longer necessary or optimal. This realization opens up an entirely new design space that is largely unexplored. Much of the design intuition about antenna design inherently relies on the antenna having a single port, so an examination of this new design space from basic principles is required.

Another benefit of having essentially free transistors available on-chip is the vast digital computational power that can be utilized. As the industry moves to smaller and smaller device nodes, the minimum dimensions approach atomic levels, and thus a few atoms of variation can cause severe performance degradation to the analog circuit. By using this digital computational power, process variations that cause the performance of different transistors to act differently can be combated by implementing a type of immune system on the chip. This system, known as self-healing, creates a feedback path that first senses the performance on the analog circuit, converts it to the digital domain, calculates a better state for the circuit to operate in, and then sends that state to actuators that are capable of adjusting the variables of importance after the chip has been fabricated. While this method can be used for any analog circuits in any process capable of digital computation, it really begins to become necessary for designs high in the mm-wave range on small feature size CMOS processes.

This thesis explores several designs that take advantage of both of these observations to increase the performance and reliability of integrated radiators and power generation blocks.

1.1 Contributions

- The design methodology of Multi-Port Driven (MPD) radiators as a way to improve the performance of integrated radiators and several novel radiating structures designed for use on an integrated circuit as examples of this approach. The radiators are designed using a holistic approach that combines the antenna, electromagnetic, and circuit design into a single process. This enables the removal of several lossy blocks that degrade performance including impedance matching, power transfer, and power combining, by designing the current patterns on the antenna to, not only radiate efficiently, but also to present the optimal impedance directly to the input drivers at multiple ports. Several proof-of-concept versions of the MPD were designed and fabricated, adding functionality to such self-healing, as well as an optically driven version.

- Dynamic Polarization Control (DPC) as a technique to control the polarization of the field radiated from an antenna or an array of antennas entirely electronically, with an exemplary DPC radiator offered as a proof-of-concept.
- Self-healing mm-wave circuits for increased reliability and performance, with examples shown both as a means to improve the radiation of an MPD radiator, and as a means to improve the performance and yield of a mm-wave power amplifier.
- Specifically, some of the highlights that resulted from designs created and demonstrated while researching these three concepts include:
  - A 160 GHz radial MPD radiator with a measured 4.6 dBm effective isotropic radiated power (EIRP) and -2 dBm total radiated power,
  - The first integrated radiator with dynamic polarization control, a 105.5 GHz 2x1 MPD radiator array with dynamic polarization control, that in linear polarization mode achieves a measured polarization angle tuning from 0° through 180° while maintaining an axial ratio of greater than 10 dB, and when tuning the axial ratio, can maintain a specified polarization angle while tuning the axial ratio from 2.4 dB to 14dB with beam steering while supplying a measured 7.8 dBm EIRP and a total radiated power of 0.9 mW, the highest demonstrated total radiated power and EIRP per element of any integrated radiator in silicon without external modifications,
  - The design and simulation of a self-healing MPD radiator that is capable of detecting its performance on chip, and controlling actuators to correct for any performance degradation,
  - The design and simulation of an optically driven MPD radiator array at 350 GHz capable of beam steering or dynamic polarization control with a simulated EIRP of -2 dBm,
  - A fully integrated self-healing power amplifier at 28 GHz that is capable of healing for process variation, transistor mismatch, load impedance mismatch, and even partial and total transistor failure, with measurements from 20 chips showing the robustness of the system, and a yield improvement from 0% to 80%.

1.2 Organization

This thesis is organized as follows. A review of the achievements and challenges associated with mm-wave radiation and power generation is presented in Chapter 2. After that, Chapter 3 introduces the concept of MPD radiators, where numerous driver stages work in unison to excite a MPD antenna and combine impedance matching, power transfer, and power combining into a single radiating element. An analysis of the operation of a radial MPD is given and the design, implementation and measurement of an MPD radiator operating at 160 GHz is presented. Chapter 4 expands the utility of MPD antennas by incorporating dynamic polarization

control (DPC) and modulation into the MPD radiator. With DPC, the polarization of the MPD antenna is dynamically controlled, which means that it can change its polarization to deliver a polarization matched wave to any receive antenna, regardless of its own polarization or orientation in space.

A version of the MPD radiator that utilizes self-healing as a way to improve the robustness of the radiator is discussed in Chapter 6. The self-healing system uses substrate mode pickup sensors to infer the far field radiation, and DC current sensors to sense the DC power consumption to determine the performance of the radiator after it has been fabricated. This data is sent to an integrated micro-processor that in turn can control amplitude and phase actuators to correct for any performance degradation caused by process or environmental variations that is within the actuation space of the radiator. Next, Chapter 5 takes the idea of the MPD antenna, and replaces its driver circuitry with photonics components, using the beat frequency of two lasers to produce THz power that is then radiated from the MPD antenna on an integrated photonics process.

Chapters 3 through 5 deal with integrated radiators; specifically different implementations and applications of the MPD Radiator. Chapter 7 instead looks specifically at power generation, and the degradation due to process and environmental variation that is observed at mm-wave frequencies. An in-depth look at self-healing as a means to correct for that degradation is given, and a fully integrated self-healing power amplifier at 28 GHz is presented, including measurements of up to 20 chips showing performance improvement across the lot.

Chapter 2

On-Chip Radiation and Power Generation from mm-Wave Integrated Circuits

2.1 Efficient Radiation from an On-Chip Antenna

Efficient radiation is a key component of an effective wireless system, whether it is being used for communication and data transfer, or for imaging or ranging applications. Traditionally at frequencies below mm-wave, after the final transistor stages of a radio frequency (RF) transmitter, the generated RF power is taken off chip and fed into a discreet antenna. As antennas are required to be at least around one half of a wavelength (λ) for efficient radiation [1], the size of the antenna scales proportionally to the wavelength, and at mm-wave frequencies, the antenna size becomes comparable to the size of the integrated circuit (IC) driving it. This opens up the possibility of further increasing the integration of the system by integrating the antenna itself on-chip.

On-chip antennas have been investigated for several decades as a way to increase the integration level of transmitter and receiver systems in ICs. [2–11].

One of the main reasons that the efficiency of on-chip antennas suffers is due to coupling to surface modes in the substrate, a phenomenon that has been studied extensively [12–17]. These on-chip antennas are sitting at the interface between the air on the top, and a lossy dielectric substrate with dielectric constant higher than air, which draws most of the power initially down into the substrate (Figure 2.1).

The antenna will excite different modes in the substrate depending on the thickness of the substrate, the frequency of radiation, and the configuration of the antenna itself. The plot on the right of Figure 2.1 shows the various modes that get excited by a dipole antenna as the ratio of the substrate thickness to the wavelength is increased. It is important to note that while these values are specifically for the dipole antenna, they show a general trend of exciting modes once the thickness of the substrate is above the cutoff for that mode in the substrate.

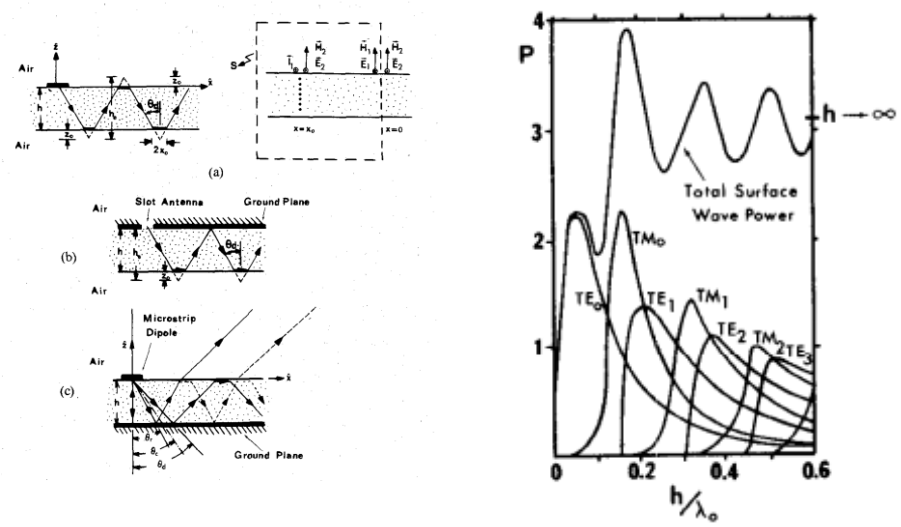

Figure 2.1: Depiction of electromagnetic rays reflecting and refracting within a substrate, and the surface modes excited by a dipole for different substrate thicknesses from [12].

Various solutions that involve expensive post-processing have been offered to disrupt these substrate modes. One way to combat this loss is to mount the chip on top of a silicon lens [5, 18, 19]. The silicon lens disrupts the substrate modes in the area of the chip, and causes the EM waves to hit a normal surface regardless of the angle they are traveling within the lens, minimizing the amount of reflection and also power lost to substrate modes, as seen in Figure 2.2. While these lenses are effective in coupling radiation out from the substrate into free space, they are bulky and expensive, and a solution that does not require this type of external modification is desired.

Another method of minimizing the substrate mode loss is through the use of a supersaturate [9]. This supersaturate is a high dielectric constant slab that is placed on top of the IC, (Figure 2.3). Some versions of this design [20] print an antenna on top of the superstrate as well. The radio frequency (RF) power is coupled to the antennas through feeds on the chip, and because of the high dielectric constant of the substrate, and the near-field nature of the coupling, more of the power is coupled up and radiated out from there. This is also an effective way of improving radiation efficiency, but again uses the expensive post-processing techniques that are involved with printing the antenna onto the dielectric supersatrate, aligning the superstrate on the IC, and attaching it, and thus a solution without post-processing is still desired.

Both of these solutions require the addition of a secondary dielectric that either must be machined precisely in the case of the silicon lens, or requires careful alignment and another lithographic step to print the antennas on top in the case of the superstrate. An integrated solution, that utilizes the options available directly on a standard IC process, with minimal, if any, post-processing would significantly reduce the cost of mm-wave radiating systems.

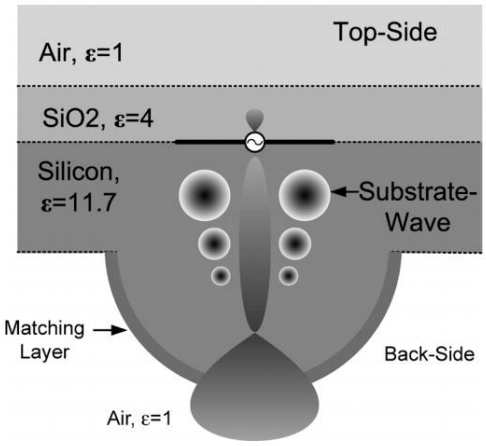

Figure 2.2: Operation of a silicon lens to minimize substrate waves and to radiate the highest power down through the bottom of the substrate from [5].

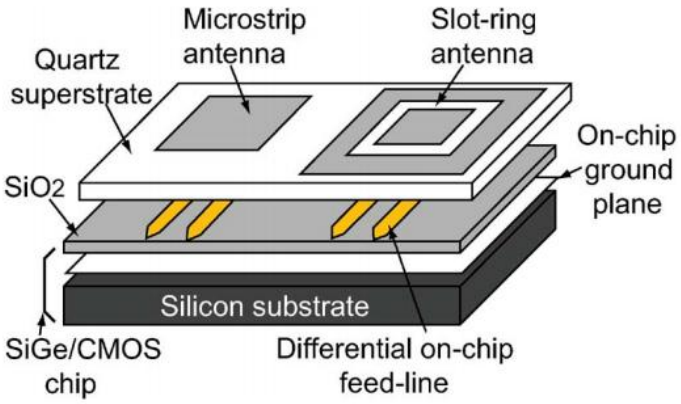

Figure 2.3: 3-dimensional depiction of the use of a dielectric superstrate (quartz in this case) to couple electromagnetic power up from the metal stack to antennas printed on top of the superstrate from [9].

In some cases, as antennas have been integrated onto the substrate, the types of antennas that were optimized in the discreet domain were ported, often with little modification to the general operation of the antenna. Dipole antennas are common [4, 5], as shown in Figure 2.4, as are patch antennas [6], seen in Figure 2.5, as they offer simple design and implementation. Slot based antennas [3, 7] are implemented to take advantage of the planer fabrication technologies that these processes offer. However these antennas usually suffer from low radiation efficiency, and still often require matching networks to match the input impedance of the antennas to the impedance, lowering the efficiency further.

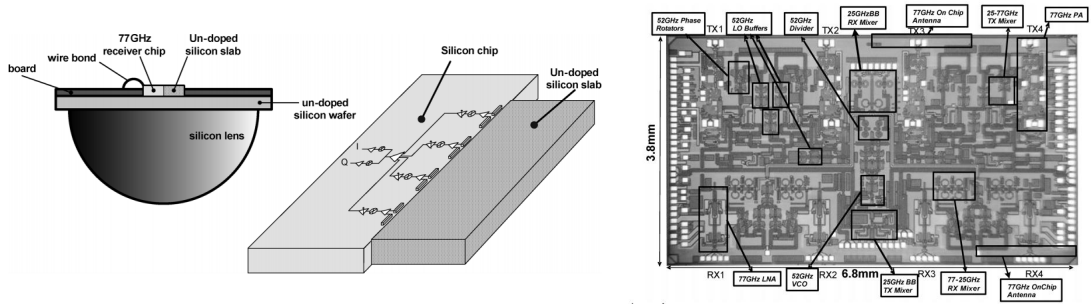

Figure 2.4: Radiation scheme showing silicon lens, block diagram and die photo from [5].

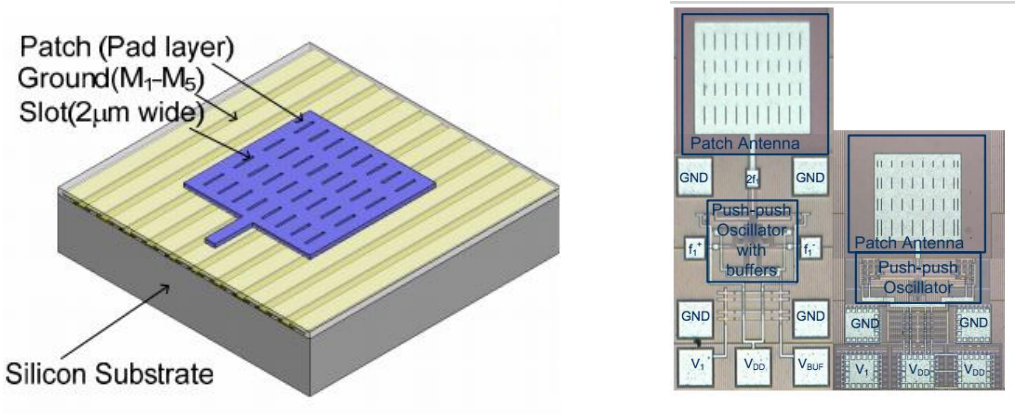

Figure 2.5: 3-dimensional rendering of the patch antenna, and the die photo from [6].

One must note that there is nothing fundamental at the levels of substrate modes excited in Figure 2.1, and it has been shown that careful engineering of the currents on the surface of the substrate can cancel out most of the substrate modes to the first order [21]. This means that extreme care must be given to both the ability of the on-chip antenna to radiate, but also to not couple heavily into substrate modes, decreasing the efficiency.

An analysis of current waves on a plane in free space is given in Appendix A, where the relationship between the currents on that plane and the radiated electric field in the direction broadside to the current

plane in the far field is derived from Maxwell's equations to be:

$$\mathbf{E} = \frac{-j\omega\mu}{4\pi}\frac{e^{-jkr}}{r}\int_C \mathbf{I}dl. \tag{2.1}$$

This equation shows that to maximize the broadside radiation from a set of currents on a plane, as is the case of an on-chip antenna, the integration of the currents on that plane should be maximized. This means that the currents should be engineered in such a way that at any instant, most of the current is facing in relatively the same direction, and thus when integrated, add constructively. Signal currents will form mirror currents when near a ground plane [22], so to create effective radiation from the on-chip antenna, the ground metals as well as the signal lines must be engineered so that both the signal current paths as well as the return ground current paths integrate constructively.

The mm-wave frequency range is where on-chip antennas begin to become viable, as efficient antennas require an aperture roughly half a wavelength in diameter [23, 24], which means that antennas at frequencies lower than mm-wave are too large to integrate on all but the largest IC substrates. If the size of an IC is on the order of millimeters, then the frequency where integrating antennas onto an IC substrate becomes feasible is on the order of tens of GHz. This means that in order for these integrated antennas to be truly useful, the process needs to be able to support efficient power generation at these frequencies.

Due to advances in fabrication processes, and the continual reduction in minimum feature size of transistors in silicon integrated processes driven by Moore's law, power generation at mm-wave frequencies in standard silicon processes has been shown to be reliable and efficient [4, 5, 25–32]. The ability to create RF power at mm-wave frequencies in standard, reliable, and relatively inexpensive processes has led to investigations into integrating antennas onto IC substrates.

2.2 State of the Art in On-Chip Integrated Radiators

This section will briefly look into five of the state of the art integrated radiators with on-chip antennas and without external modification.

In [33–36], the distributed active radiator (DAR) is presented. The DAR consists of a traveling wave oscillator, where the signal line is arranged in a figure-8 fashion that keeps differential pairs of transistors near each other (Figure 2.6). At the fundamental frequency of oscillation, this also means that differential signals are always traveling next to each other on the two signal lines, acting as a differential transmission line, and thus the fundamental frequency has minimal radiation. If these transistors are driven hard into their nonlinear regions, a second harmonic signal is also produced, which is in phase for each differential pair. This means that currents on the two signal lines will be going in the same direction at any instant and thus will radiate effectively. Coupling to substrate modes is also kept to a minimum by the placement of several of the DAR structures in an array in such a way that the substrate modes from various radiators cancel each

other out. Several versions of the DAR have been published, culminating in a 4x4 array at 280 GHz. This array achieves a +9.5 dBm effective isotropic radiated power (EIRP) with total radiated power of 190 μW, and demonstrates beam steering of 80°. The DAR is a self-oscillating radiating structure, which means that the reactive elements that determine the frequency of oscillation (the tank) are also the same elements used to radiate. This means that there is a danger of the DAR picking up signals that are near in frequency and injecting those signals back into the oscillator, which might pull, or in the worst case lock, the radiator to an undesired frequency. This concern can be alleviated by designing radiators that are driven unilaterally, where the reactive elements that determine the oscillation are separated from the radiating elements by unilateral amplification stages, ensuring that any external signals do not get coupled into the tank of the oscillators.

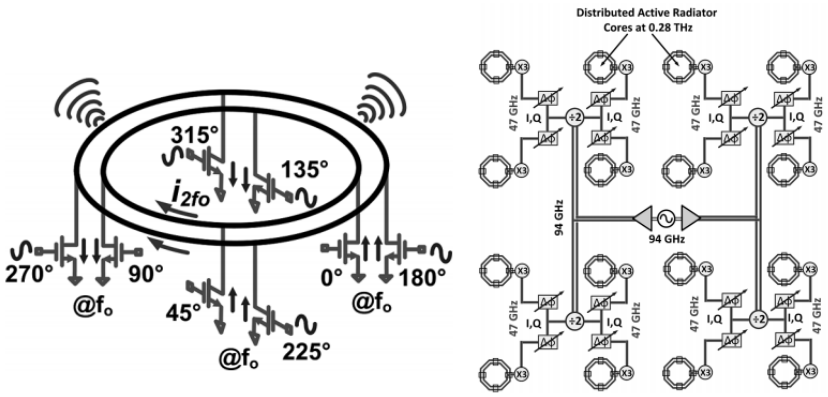

Figure 2.6: Conceptual depiction of a single distributed active radiator (DAR), and the block diagram of the 4x4 DAR phased array from [34].

In [37], a 210 transceiver is implemented using a 2x2 array of dipole antennas (Figure 2.7) in 32 nm CMOS SOI. The transmitter is made up of a 210 GHz VCO whose output can be modulated with on-off keying (OOK). The signal is amplified and split into 4 and is amplified again before being sent into the antennas. This work achieves a 5.13 dBm EIRP, with a saturated power amplifier (PA) power of 4.6 dBm per PA, or 10.6 dBm for the system. This leads to an antenna gain from the output of the PA to the far field electric field of -5.47 dBi, which means there is still significant room for improvement in the performance of the antennas.

In [8], an array of 4 leaky wave antennas are used both as receive as well as transmit antennas at 260 GHz (Figure 2.8). The leaky wave antennas are implemented as microstrip transmission lines with a bottom ground plane that is still above the silicon substrate, shielding the signal from the lossy substrate and preventing surface modes from being excited. However this type of arrangement means that mirror return currents will appear in the bottom ground plane and will cancel out a significant portion of the radiation. This means that very large currents are required to get a reasonable total radiated power, which often leads to lower antenna efficiency, as the metal losses at 260 GHz on integrated processes are often quite high. This can be seen in

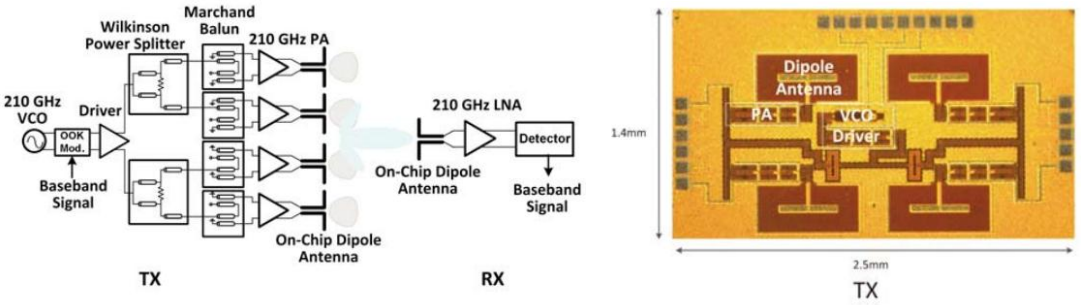

Figure 2.7: Block diagram of the transmitter and receiver, and the die photo of the transmitter showing the radiator array from [37].

this work, as the PA provides +13 dBm of output power, and the reported EIRP is +5 dBm at 260 GHz, which means that the antenna gain from the output of the PA to the far field is -8 dBi. The total radiated power is not reported. This work also focuses on non-coherent (OOK) modulation to create a wireless fatal-link over a 40 mm range at 14 Gb/s. The leaky wave antennas are also 1250 μm long, which at 260 GHz is more than one wavelength in air, which means that compared to their wavelength, these antennas require a large footprint, and thus become very costly.

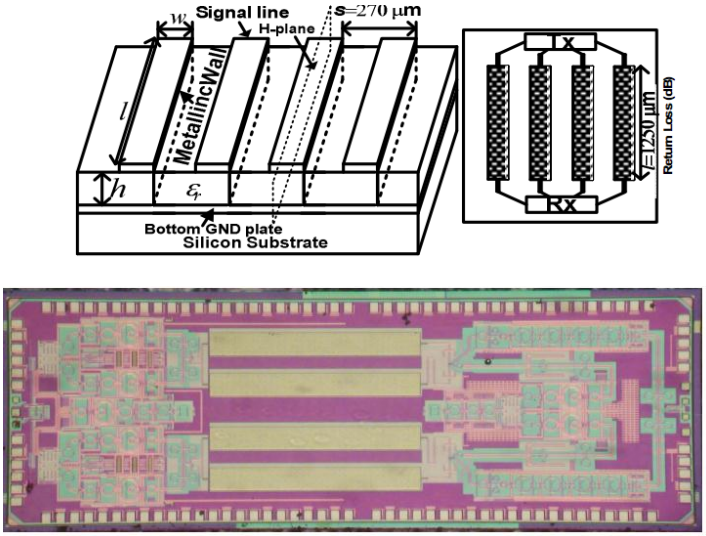

Figure 2.8: 3-dimensional depiction of the four leaky wave transmission line antennas, their use for both transmitter and receiver, and a die photo showing the transceiver from [8].

In [38], tapered dipoles are used for both the transmit and receive antennas in a transceiver at 164 GHz (Figure 2.9). The radiation was picked up by a standard RF probe hovering over the top of the chip. Because the antenna properties of the probe were never measured, no standardized metrics of the radiator performance are presented. The polarization of the probe being used as an antenna is also not discussed, which means that

comparisons between different antennas measured in this setup may be due to the performance of the antenna itself, or possibly due to better polarization matching between some of the antennas and the probe. Another concern is that the probe is likely located in the near field of the on-chip dipole, which means that the near-field coupling between the probe and dipole may not only be interfering with the radiation of the dipole, but also may be changing its input impedance, which affects the output power of the driver PA. Still this setup is able to confirm that the chip is capable of radiating some power at 164 GHz.

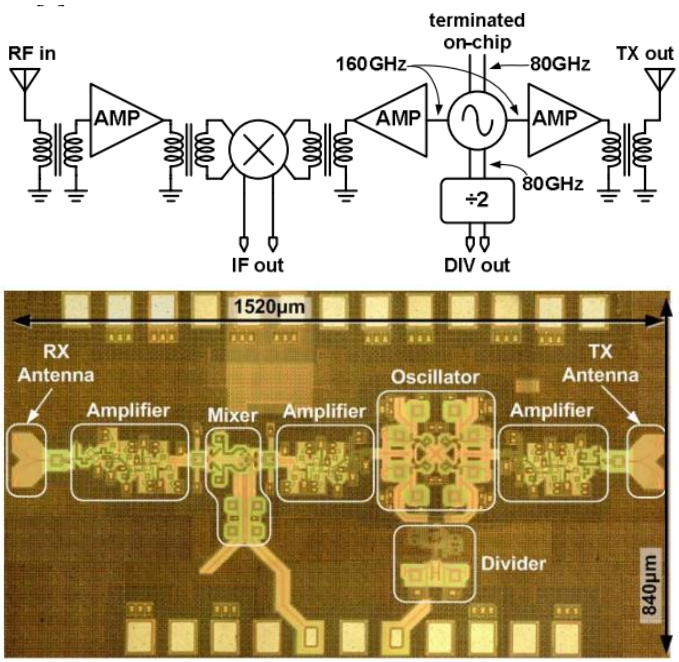

Figure 2.9: Block diagram of the transceiver and die photo showing transmit and receive antennas from [38].

2.3 Performance Degradation due to Process and Environmental Variation

Another important consideration that is critical for robust design of mm-wave circuits is process and environmental variation. The rise of digital computation and personal computing has led to continual advances in semiconductor technologies at an exponential pace, following Moore's Law. In each successive processing node, the minimum feature size decreases, improving performance, but also bringing some tradeoffs in terms of variation both between chips as well as between transistors on the same chip [39, 40, 40, 41]. One major source of this variation is random dopant fluctuations (RDF) in the channel of a transistor [42, 43]. A typical 130 nm CMOS process will have several hundreds of dopant atoms in the channel region. In contrast in a 32 nm process, only a few tens of dopants control important transistor characteristics like threshold voltage, etc.

A second source of variation is line-width control in these advanced processes. Line edge roughness (LER) caused by lithographic and etching steps directly impact the overlap capacitances as well as other device parameters like drain induced barrier lowering (DIBL) and threshold voltage [44].

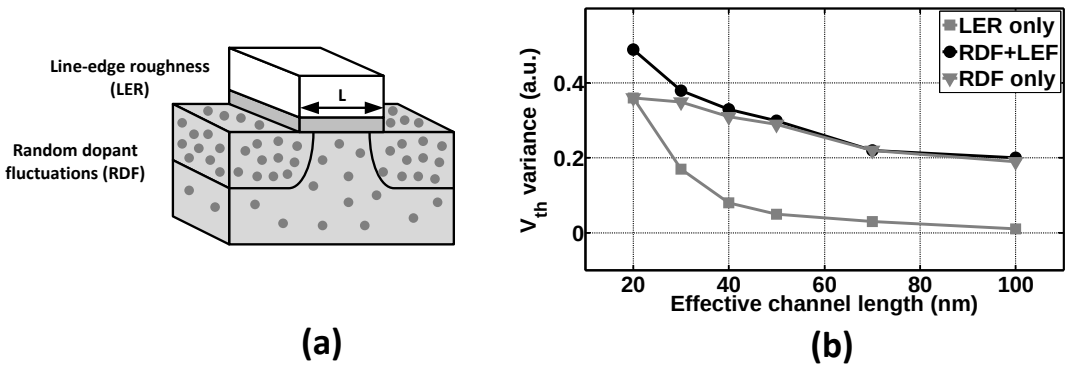

Figure 2.10: (a) Two main sources of variation in a MOSFET, and (b) threshold voltage variation over technology node [39].

Figure 2.10 shows how threshold voltage variations scale with process technology nodes [39]. As can be seen, the variation is much more manageable at larger nodes, and the variation is expected to continue to increase at smaller nodes as the total number of dopant atoms as well as the channel length reduces even further [45]. If the variation can be dealt with, however, the smaller transistors can enable new applications for mm-wave power generation, enabling transmitters and amplifiers at higher frequencies, powers, and efficiencies. Another issue that analog designers face is that due to the digital processing market being the driving force pushing the scaling, the models provided by the foundries early in the node's development stage are primarily designed for digital use, and are often not reliable at mm-wave frequencies.

In addition to these static sources of variation, dynamic temperature variations across the same die can give rise to varying sub-threshold leakage and supply voltage variations, thereby directly affecting overall system performance. Variability in operating environment of power generation systems can adversely affect their performance. This comes in the form of temperature variation, degradation due to aging [46], and in the case of power amplifiers that are driving antennas, load impedance mismatch caused by voltage standing wave ratio (VSWR) events [47, 48] that occur when objects in the environment interact in the near field of the antenna as can be seen in Figure 2.11. Dealing with this issue is critical especially when the amplifiers are in a phased array, as interactions between antennas can allow signals from other elements in the array to couple back through the antenna [4, 5, 49–52]. Power amplifiers are generally tuned to provide the optimal output power at maximum efficiency for a designed load, so when that load changes, the performance drops, and in extreme cases, can damage the chip if care is not taken to ensure breakdown voltages are not exceeded for any expected VSWR events.

Self-healing is a method that can mitigate these issues by identifying any degradation and modifying the

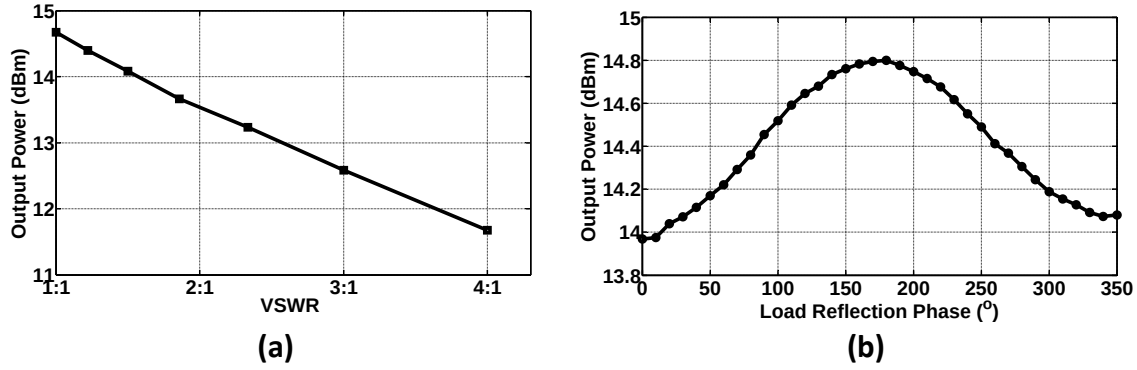

Figure 2.11: (a) Output power variation with voltage standing wave ratio (VSWR) magnitude for a fixed phase, and (b) variation with phase variation for a fixed VSWR magnitude for one design of a mm-wave power amplifier (PA) [53].

circuit to improve its performance post-fabrication with minimum overhead by incorporating a feedback loop into the system that takes advantage of the very low cost of digital processing available in these advanced processes.

2.4 Current Technology in Adaptive and Self-Healing Circuits

Self-healing mm-wave systems are just starting to become feasible, and while there is not a large body of previous work to refer to, a few other works have been published, two of which will be briefly summarized below.

In [54], partially integrated self-healing was used to detect and correct for spurious tones in an 8-12 GHZ synthesizer (Figure 2.12). The spurious tones were detected using the voltage controlled oscillator (VCO). Control voltage sensors sampled the voltage digitally using a digital signal processing (DSP) unit. The system was corrected using actuators in the form of a periodic correction signal that counteracts any deviation in the VCO control using narrow, digitally generated pulses. This work integrated the sensors and actuators on-chip, but an off-chip DSP was used to control the system. The system achieved a fundamental spur reduction of 10 dB over the frequency range, with a total spurious power reduction of 20 dB at 10.4 GHz.

In [55], a 60 GHz self-healing radio-on-a-chip is presented (Figure 2.13). This system uses envelope variation, power level, and temperature sensors to determine transmitter performance when the self-healing controller produces test tones at various programmable frequencies and amplitudes. The PA has actuators in the form of bias tuning, and there are also actuators for in-phase/quadrature (IQ) correction, allowing adjustment of IQ DC offset, IQ phase, and IQ relative amplitude. The receiver can be healed by connecting the known output of the envelope sensor to the input of the receiver chain, and adjusting the bias levels within the low noise amplifier (LNA) to optimize noise figure. [56, 57] describes the operation of the actuated PA in more detail.

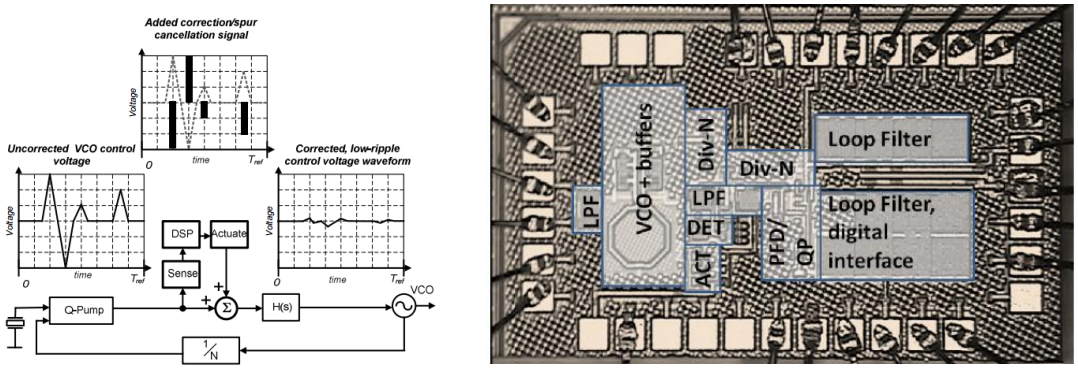

Figure 2.12: Block diagram of the self-healing synthesizer to suppress spurious tones, and the die photo of its implementation from [54].

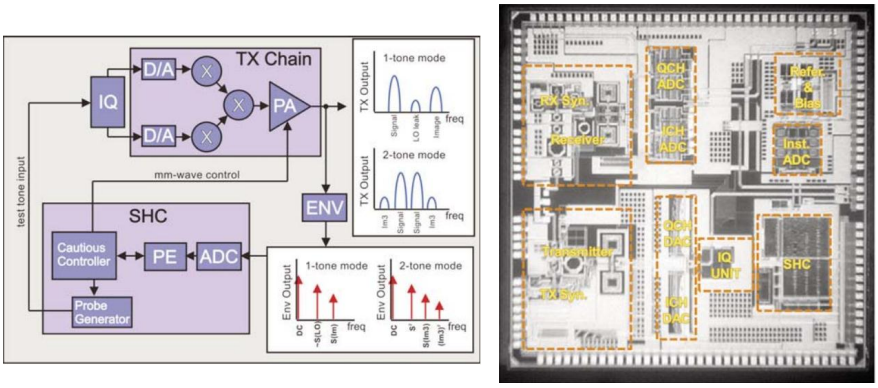

Figure 2.13: Block diagram of the transmitter half of the self-healing system, and the die photo of the entire self-healing transceiver from [55].

Similar adaptive circuits schemes include built-in self-test (BIST). BIST enables an IC to test itself to identify process variation and mismatch occurring on an IC and to correct the circuit for it in a similar way to self-healing. While there may be some semantic debate about where BIST ends and self-healing begins, for the purposes of this work we consider BIST to be a test that is run once when the chip is first turned on to deal with inherent variation from fabrication, and self-healing to be a more generic method that includes an algorithm that can be run anytime the performance is suffering and can heal not just for fabrication error, but also for environmental variation, and other time-dependent sources of performance degradation. Some examples of mm-wave BIST include [58–61].

Chapter 3

Multi-Port Driven Radiators

3.1 Introduction

Transistor scaling in advanced silicon processes continues to enable designers to push the boundaries of high frequency designs in silicon integrated circuits (ICs). Recently, advancements in mm-wave power generation have shown that these processes can be used to create power in mm-wave up to low THz frequencies [26, 29, 51–53, 62–66]. However, once the power is created on the chip, it still needs to be efficiently transferred off chip to be useful in the real world. Above 100 GHz, traditional methods of power transfer such as wire-bonding and flip chip are often unreliable and highly lossy [67, 68]. If the power can be extracted efficiently from the ICs, low cost solutions for applications ranging from imaging [11, 69] for biomedical cancer detection or security to high bandwidth point-to-point communications [8, 37] can become feasible.

Traditionally, electromagnetic radiation at radio frequency (RF) and mm-wave frequencies is produced using a transmitter, with separately designed blocks for signal generation, power amplification, impedance matching, and power transfer, and an antenna that is connected to the transmitter through a single port (Figure 3.1a) [4, 38, 70, 71]. In the nomenclature for this paper, we will refer to an antenna as the passive device that radiates electromagnetic fields when it is driven, and a radiator as the combination of an antenna and its driver circuitry. This traditional block-by-block single-port design is the result of an optimization of the costs of these radiators when the antenna and transmitter driver circuitry blocks were discreet components. In the discreet case, individual transistors are expensive, and physically connecting the blocks with more than one port is similarly expensive. As antennas began to be integrated onto the same substrate as their supporting circuitry [5–7, 9, 10, 33, 34], this traditional design methodology has often continued to be used. Integrating the antenna onto the substrate eliminates the need for other power transfer methods to get the power off chip, however, the design strategies of integrated antennas need to be reevaluated, as the cost analysis of the integrated design is vastly different than that of the discreet design. For integrated radiators the cost of the radiator is dominated by the area of the design.

At high mm-wave frequencies, the size of the antennas (often around one half of the wavelength) is small enough to affordably fit on an integrated circuit substrate, but is still large compared to the individual

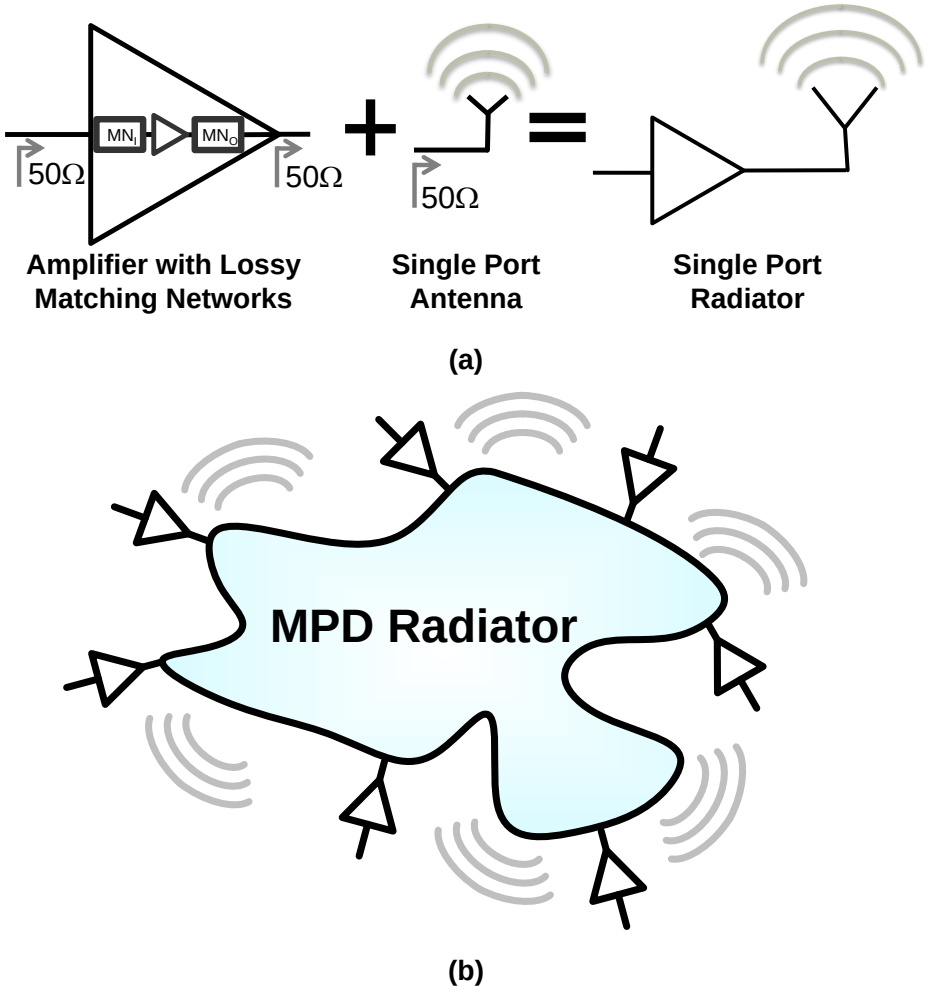

Figure 3.1: (a) Traditional single port radiator with lossy matching networks and independent design of individual blocks, and (b) proposed multi-port driven (MPD) radiator design methodology for integrated radiation where drivers directly feed multiple ports of the antenna.

transistors. The frequency for the radiators, as well as the number of antenna elements if the radiator is to be an array will then primarily determine the size of the IC and thus its cost. This means that transistors have become essentially free in these designs and can be placed anywhere they are useful [72]. Also, as long as metal density requirements are met, connections between transistors and the antenna are also free, and thus the antenna can be driven from any point where it is useful from a radiation standpoint. This opens up an entire design space of multi-port driven (MPD) radiators that is often ignored in traditional single port antenna design. By using a hybrid-design methodology of the drivers and integrated antennas with multiple drive ports, the driver circuitry can be placed closer to the radiating elements by eliminating many of the lossy blocks including power-combining networks, output-impedance matching networks, and off-chip power transfer elements such as wire bonds or flip chip (Figure 3.1b). There is still a tradeoff however, as integrated antennas can suffer from increased metal losses compared to discreet antennas, and can encounter losses associated with their proximity to the lossy substrate. As frequencies increase, and wavelengths get shorter, the losses associated with impedance matching and off-chip power transfer go up, while the losses associated with radiating on top of a lossy substrate and the area cost of that IC go down, making integrated radiators a cost effective solution at high frequencies.

Section 3.2 investigates MPD antenna design, proposes an architecture for such an MPD antenna, and presents an analysis of its free-space broadside radiation. Section 3.3 covers the electromagnetic simulation and design of the MPD antenna, and Section 3.4 covers the design of the driver circuitry. The top-level radiator design and simulations are explained in Section 3.5, with measurements of the chip presented in Section 3.6. A slot based MPD design is covered in Section 3.7, followed by a differentially driven linear radiator design in Section 3.8 and finally concluding remarks are given in Section 3.9.

3.2 Multi-Port Driven Antenna Architecture and Analysis

Much of the intuition designers have for the limits of antenna design comes from theories that assume single port design. One feature of MPD design that alleviates a common issue with single port designs is creating efficient antennas with low input impedance. While it is often true for a single port design that the radiation resistance, and thus the input impedance, needs to be high for a single port design, by utilizing multiple ports intelligently, virtual short circuits can enable low impedance designs with high efficiencies, as the currents from various ports can be engineered to add up without increasing the voltages seen at the ports. This can often be seen through a superposition of the various drive ports, and is important for integrated radiators as the voltage breakdown of the transistors within the power amplifier stages limit the peak voltage, which means that current must be increased for high power designs, and the optimal output impedance decreases [53, 73].

Another advantage of MPD design is the ability for the antenna itself to do power combining. For high power, single port power amplifiers, the signal is often divided, amplified, and then combined together [63]. The power combining stage occurs after the final output transistors, and thus any loss in that stage can signif-

icantly degrade the output power and efficiency. Using MPD design, the power combiner stage is absorbed into the antenna itself, so the currents produced by the various driver stages begin to radiate immediately after being produced by the transistors, and the power is combined quasi-optically in the air due to the integration of the currents around the antenna adding up coherently.

Taken to their limit, multi-port driven antennas can create any physically realizable current pattern on the plane by injecting currents wherever they are needed around the antenna, something that is often not feasible or even possible with a single port design. In this MPD inverse design approach, the current pattern is first determined by the desired far-field radiated pattern. The MPD antenna and its driver circuitry are then designed to produce the desired current.

For the purpose of gaining MPD design intuition, an analysis of a radial MPD antenna in free space is presented from basic electromagnetic equations. To keep the solution simple enough to gain this intuition, only the broadside radiation of the antenna is considered. The behavior of the radiation pattern in non-broadside directions is then studied through electromagnetic simulations in Section 3.3.

3.2.1 Broadside, Far-Field Radiation From Currents on a Plane

The analysis will start from the time-harmonic equations for electric field and magnetic vector potential in a uniform, isotropic, linear medium in the far field in spherical coordinates that is derived in Appendix A that follows the derivation in [1],

$$\mathbf{E} = -jw\left[A_\theta(r,\theta,\phi)\hat{\mathbf{u}}_\theta + A_\phi(\theta,\phi)\hat{\mathbf{u}}_\phi\right] \tag{3.1}$$

and

$$\mathbf{A} = \frac{\mu}{4\pi}\iiint_V \mathbf{J}\frac{e^{-jkr}}{r}dv, \tag{3.2}$$

which can be derived from Maxwell's equations.

If the current densities to be considered are all contained within wires, as is the case for the MPD antenna, the volumetric integral becomes a line integral. Also, because the currents are all in a finite area on the same plane, the observation distance $r = R_0$ is constant from all currents to the far-field radiation observation point broadside to that plane, and thus (3.2) can be simplified to

$$\mathbf{A} = \frac{\mu}{4\pi}\frac{e^{-jkR_0}}{R_0}\int_C \mathbf{I}dl. \tag{3.3}$$

The currents are all in the plane perpendicular to r, $A_r(\theta,\phi) = 0$, making

$$\mathbf{A} = A_\theta(\theta,\phi)\hat{\mathbf{u}}_\theta + A_\phi(\theta,\phi)\hat{\mathbf{u}}_\phi. \tag{3.4}$$

Plugging $\mathbf{A}$ into (3.1) gives

$$\mathbf{E} = \frac{-j\omega\mu}{4\pi}\frac{e^{-jkR_0}}{R_0}\int_C \mathbf{I}dl, \tag{3.5}$$

that can be used to solve the broadside far-field electric field produced by any given current distribution on a plane.

3.2.2 Multi-Port Driven Antenna Architectures

The result of (3.5) is that in order to maximize the far-field electric field at any instant, the integration of the currents on the plane need to be maximized. This is accomplished by having the currents on the plane pointed roughly in the same direction at any given instant, and is why a transmission line, where the ground current is opposite of the signal current, does not radiate, while a dipole, where the currents in both halves of the antenna are pointed in the same direction, does.

One example of an effective way to create currents that integrate coherently and radiate effectively in the signal lines is to create a traveling wave current around an unbroken ring that is around a wavelength (λ) in circumference, causing a phase change due to delay along the line of π as the wave travels half way around the ring, as well as a rotation of the current's direction in space of π due to the curvature of the ring. There are then two current maximums at any instant on opposite sides of the ring that point in the same direction, and rotate around the ring over time, as shown in Figure 3.2.

The maximum currents being pointed in the same direction at any instant create an effective radiating structure. However, the structure needs to be driven, and return ground currents of those drivers also need to be considered and engineered such that the integration of the return currents do not cancel out that of the signal currents. For this design, driver amplifiers inject power into the ring at evenly spaced points around the ring and at evenly spaced phases, exciting the traveling wave on the ring. The return currents are routed through ground lines that connect the input ports back to the center of the antenna. Taking the nomenclature of the bicycle wheel, these lines will be referred to as the ground 'spokes', as shown in Figure 3.3. Due to the symmetry of the problem, a virtual short circuit is expected at the origin, which creates standing current waves on the ground spokes. If a configuration can be found that makes the integration of the ring and spoke currents add up coherently, both currents will contribute to the overall radiated electromagnetic field. The goal of the rest of the analysis is to determine the conditions to make this possible.

Two assumptions are required for the analysis. The first is that the current pattern of a wire that is terminated with a virtual short circuit will have a sinusoidal current standing wave similar to a transmission line terminated with a short circuit. This is an approximation similar to the complement assumption about wires that have open circuits at one end that is often used in the analysis line antennas such as the dipole [1]. Higher accuracy can be achieved by modeling the current similar to a lossy transmission line by taking into account loss due to metal attenuation as well as power delivered through radiation, but there is minimal effect to the current patterns, and the simplicity of the equations are then lost. The second assumption is that there

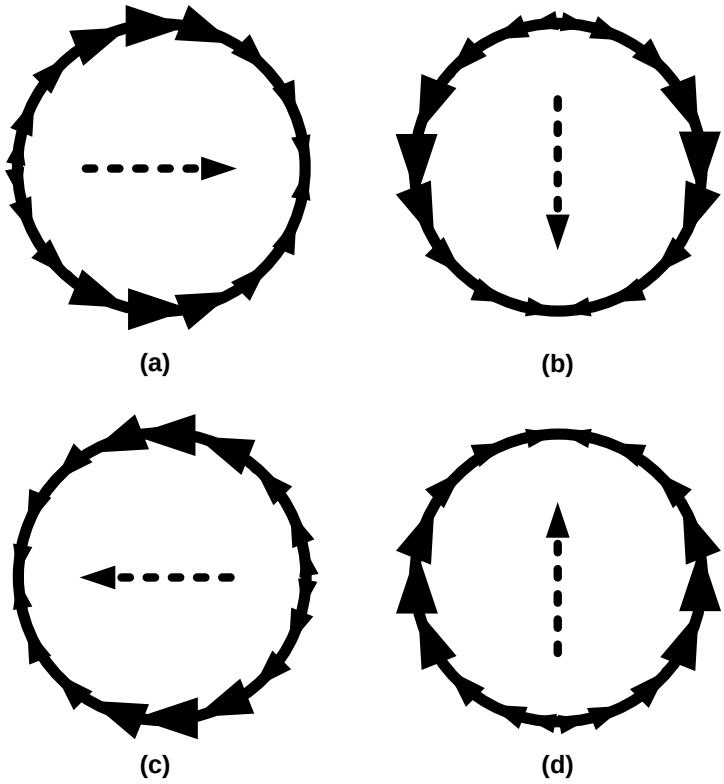

Figure 3.2: Instantaneous current distribution for the traveling current wave around the ring for four phases: (a) 0^o, (b) 90^o, (c) 180^o, and (d) 270^o, where the integrated current vector depicted as the dashed arrow shows maximum currents on ring pointing in the same direction for all phases.

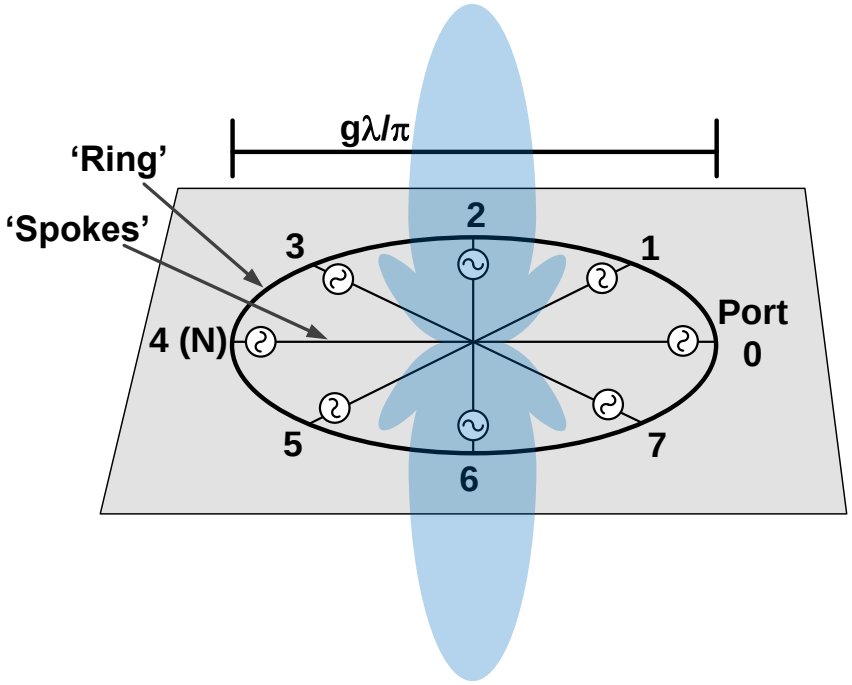

Figure 3.3: Proposed MPD antenna architecture showing ring and ground spokes with an overlay of the expected radiated beam, as well as design variables g and N. For this design there are 4 spokes so $N = 4$.

is no coupling between the different spokes or between the spokes and the ring. This assumption will enable the use of superposition, and is verified as a valid assumption through electromagnetic simulation.

To solve the electric field produced by the MPD antenna, it is broken down into its two sections, the ring and the spokes, each evaluated separately using superposition. Because of the symmetry of the design, the currents entering the ring at each port will have the same amplitude and will be phased evenly based on the position of the port around the ring. If this current magnitude (I_s) is known, the power sources that have finite source impedances at the ports can be replaced with current sources with magnitudes I_s and the appropriate phases using source substitution. This means that when superposition is applied, the nulled sources become open circuits with zero current flow.

In order to keep the analysis more general, two design parameters, g and N, are defined for the MPD antenna, where g is the circumference in units of the wavelength λ, and N is the integer number of spokes, spaced evenly with $N \geq 2$. While the analysis will keep N as a design variable, the implemented design that is shown in the figures is for a 4 spoke MPD antenna, so $N = 4$ (Figure 3.3). The superposition analysis will consider 2 ports at a time located opposite of each other on the ends of ports n and $N + n$, with differential drive phases due to symmetry, and then sum N of these pairs of ports to consider the entire system. The polar coordinate angle (ϕ) on the plane of the antenna is used for many of the calculations.

3.2.3 Solution to Currents on the Ring

The ports at $\phi = 0$ and π are considered first as shown in Figure 3.4, which are the ports on either side of spoke $n = 0$, with all other ports nulled as open circuits. With only two of the ports of the ring being driven, standing current waves are expected rather than the traveling current waves that occur once all ports are driven. Because of the differential symmetry of the system, there will be virtual short circuits of maximum current flow on the ring at $\phi = \pi/2$ and $3\pi/2$. There are two planes of symmetry which mean that the $\hat{\mathbf{u}}_\mathbf{y}$ component of the current will cancel, while the $\hat{\mathbf{u}}_\mathbf{x}$ component of the current in each quadrant will add coherently. This means that only the $\hat{\mathbf{u}}_\mathbf{x}$ component in the first quadrant from $\phi = 0$ to $\pi/2$ needs to be integrated, and then can be multiplied by 4 to get the overall integration for the ring. Also, since the current input into each port is split into two directions on the loop, the magnitude of the current going into each quadrant of the loop is $I_s/2$.

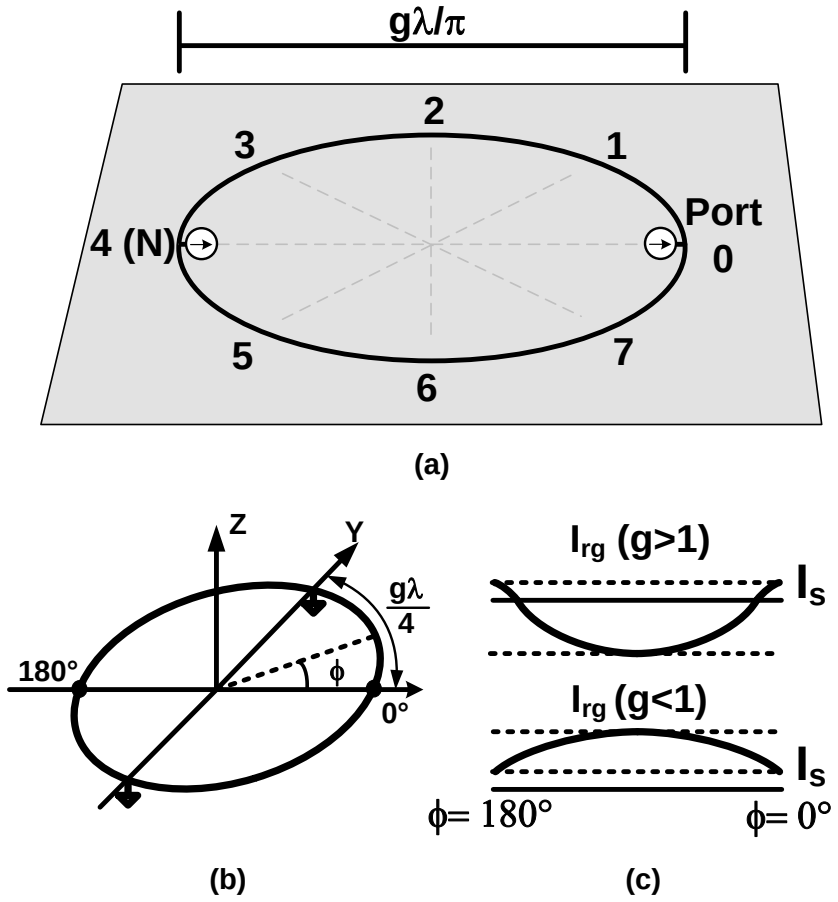

Figure 3.4: (a) Shows the ring portion of the MPD antenna being driven by only ports 0 and N as part of the overall superposition analysis, while (b) depicts the location of the virtual shorts observed under these driving conditions and (c) shows the counterclockwise oriented standing wave current along the ring, where the sign of the maximum current changes sign between $g < 1$ and $g > 1$ at 0 phase for the first two quadrants.

Using the virtual short on the positive y axis, and taking the current to be a sinusoidal standing wave from that point back to the port, the magnitude of the current along the ring in the first quadrant ($\mathbf{I_{rg}}$) can be written as

$$|\mathbf{I_{rg}}| = I_{maxr} \cos \beta l, \tag{3.6}$$

where I_{maxr}, β and l represent the maximum time harmonic current phasor on the ring that occurs at $\phi = \pi/2$, the wave propagation constant ($\beta = 2\pi/\lambda$) and the length along the line from the short circuit at $\phi = \pi/2$, respectively.

Noting that $l = (\pi/2 - \phi)g\lambda/(2\pi)$,

$$|\mathbf{I_{rg}}| = I_{maxr} \cos \left[g(\pi/2 - \phi)\right]. \tag{3.7}$$

Next, defining $\phi' = \pi/2 - \phi$, the $\hat{\mathbf{u}}_\mathbf{x}$ component of the current in the first quadrent becomes

$$\mathbf{I_{rgx}} = -I_{maxr} \cos\left(g\phi'\right) \cos\left(\phi'\right) \hat{\mathbf{u}}_\mathbf{x}. \tag{3.8}$$

Inserting this into (3.5) in polar coordinates gives

$$\mathbf{E_{rgx}} = \frac{-j\omega\mu}{4\pi} \frac{e^{-jkR_0}}{R_0} \int_0^{\pi/2} -\frac{\lambda g}{2\pi} I_{maxr} \cos\left(g\phi'\right) \cos\left(\phi'\right) d\phi' \hat{\mathbf{u}}_\mathbf{x}. \tag{3.9}$$

Pulling out the terms not dependent on ϕ' and noting that $\lambda = 2\pi c/\omega$ where c is the speed of light, yields

$$\mathbf{E_{rgx}} = I_{maxr} \frac{j\mu g c}{4\pi} \frac{e^{-jkR_0}}{R_0} \int_0^{\pi/2} \cos\left(g\phi'\right) \cos\left(\phi'\right) d\phi' \hat{\mathbf{u}}_\mathbf{x}. \tag{3.10}$$

Solving out the integral gives

$$\mathbf{E_{rgx}} = -I_{maxr} \frac{j\mu c}{4\pi} \frac{e^{-jkR_0}}{R_0} \frac{g}{g^2 - 1} \cos\left(g\frac{\pi}{2}\right) \hat{\mathbf{u}}_\mathbf{x}. \tag{3.11}$$

Because $I_s/2 = |\mathbf{I_{rg}}|$ for $\phi = 0$, the input current from the port I_s can be calculated from (3.7).

$$\frac{I_s}{2} = I_{maxr} \cos\left(g\frac{\pi}{2}\right), \tag{3.12}$$

and thus

$$\mathbf{E_{rgx}} = -I_s \frac{j\mu c}{8\pi} \frac{e^{-jkR_0}}{R_0} \frac{g}{g^2 - 1} \hat{\mathbf{u}}_\mathbf{x}. \tag{3.13}$$

Considering all 4 quadrants, the E-field is multiplied by 4, and remembering that the $\hat{\mathbf{u}}_\mathbf{y}$ components of the

currents cancel due to symmetry, making $\mathbf{E_{r0y}} = 0$,

$$\mathbf{E_{rg}} = 4\mathbf{E_{rgx}} = -I_s \frac{j\mu c}{2\pi} \frac{e^{-jkR_0}}{R_0} \frac{g}{g^2-1} \hat{\mathbf{u}}_\mathbf{x}. \tag{3.14}$$

This means that the far-field electric field will change directions when $g = 1$. This is due to the fact that when $g < 1$ the length of line from the virtual short on the y axis to the port on the x axis is less than $\lambda/4$, which means there is no node and all of the current along the line in that quadrant is always traveling the same direction, either clockwise or counterclockwise, as seen in Figure 3.4c. However, when $g > 1$, the length of the line is longer than $\lambda/4$, which will produce a node in the current standing wave, and will make the current at the virtual short point in the opposite direction along the line as the current near the port, as shown in Figure 3.4c. For the case where $g = 1$, the equation goes to infinity. This is due to the assumption that the line can be treated as a lossless transmission line, and for a $\lambda/4$ line terminated with a virtual short, the impedance is infinite at the input port. Thus for a finite input current going into an infinite impedance, you expect to get infinite power. Modeling more realistically as a lossy transmission line removes this asymptote. This is why the phase of the electric field due to the current on the ring is reversed by changing the design from $g < 1$ to $g > 1$.

3.2.4 Solution to Current on the Spokes

The current on the spokes, caused by the ports at $\phi = 0$ and π on either side of spoke $n = 0$ (Figure 3.5) is examined next. The radius of the ring, and thus the length of the spoke is $g\lambda/2\pi$. Using differential symmetry again, a virtual short circuit is observed at the origin. Thus all non-driven spokes have one end as a short circuit, while the other is an open circuit. Remembering the assumption of low coupling between spokes, the current on those spokes are taken to be zero, and thus neglected. For the driven spokes, the line is terminated with a virtual short, and results in a current standing wave that is given by

$$\mathbf{I_{sp}} = I_{maxsp} \cos\left(\frac{2\pi}{\lambda} x\right) \hat{\mathbf{u}}_\mathbf{x}, \tag{3.15}$$

where I_{maxsp} is the maximum current on the spoke that occurs at the center of the antenna. Plugging this into (3.5) gives

$$\mathbf{E_{sp}} = \frac{-j\omega\mu}{4\pi} \frac{e^{-jkR_0}}{R_0} \int_{-g\lambda/2\pi}^{g\lambda/2\pi} I_{maxsp} \cos\left(\frac{2\pi}{\lambda} x\right) dx \hat{\mathbf{u}}_\mathbf{x}, \tag{3.16}$$

which when integrated and considering $\lambda = \omega/2\pi c$ yields

$$\mathbf{E_{sp}} = -I_{maxsp} \frac{j\mu c}{2\pi} \frac{e^{-jkR_0}}{R_0} \sin(g) \hat{\mathbf{u}}_\mathbf{x}. \tag{3.17}$$

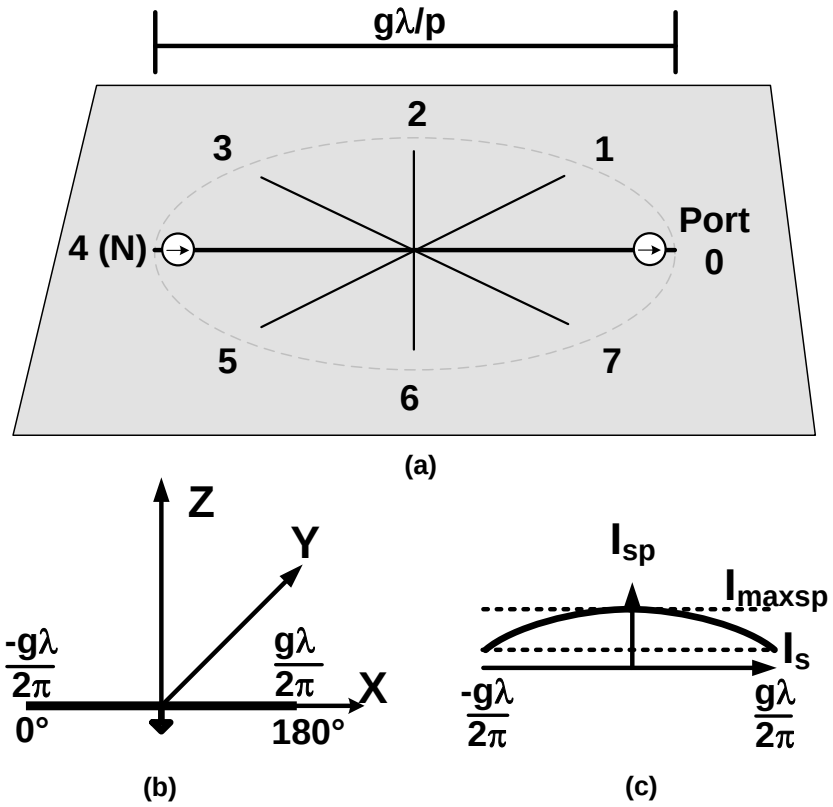

Figure 3.5: (a) Shows the spoke portion of the MPD antenna being driven by only ports 0 and N as part of the overall superposition analysis, noting that the phase of the current driven into the ground spokes is opposite of that for the ring, and (b) depicts the location of the virtual short circuit observed at the origin and (c) shows the standing wave current along the driven spoke at 0 phase.

Using (3.15) we can solve I_s in terms of I_{maxsp}, with

$$I_s = I_{maxsp} \cos(g) \tag{3.18}$$

which yields

$$\mathbf{E_{sp}} = -I_s \frac{j\mu c}{2\pi} \frac{e^{-jkR_0}}{R_0} \tan(g) \hat{\mathbf{u}}_\mathbf{x}. \tag{3.19}$$

This equation is similar to (3.14), but because the line length of the spoke is different than that of the ring for a given value of g, nodes appear on the two lines at different values of g. For the spoke there is no node for $g < \pi/2$ (Figure 3.5c), and thus the current on a driven spoke will be in phase across the entire spoke at any instant as long as $g < \pi/2$.

3.2.5 Summing up the E-fields From All Ports

With (3.14) and (3.19), the superposition of the fields caused by the currents on the rings and the spokes from ports 0 and N gives

$$\mathbf{E_0} = -I_s \frac{j\mu c}{2\pi} \frac{e^{-jkR_0}}{R_0} \left[\frac{g}{g^2 - 1} + \tan(g) \right] \hat{\mathbf{u}}_\mathbf{x}. \tag{3.20}$$

When considering the contribution of the other ports to the overall electric field, both the phase shift as well as the angular rotation around the origin must be taken into account as

$$\mathbf{E_n} = e^{j\frac{\pi n}{N}} E_0 \left[\cos\left(\frac{\pi n}{N}\right) \hat{\mathbf{u}}_\mathbf{x} + \sin\left(\frac{\pi n}{N}\right) \hat{\mathbf{u}}_\mathbf{y} \right], \tag{3.21}$$

where

$$E_0 = -I_s \frac{j\mu c}{2\pi} \frac{e^{-jkR_0}}{R_0} \left[\frac{g}{g^2 - 1} + \tan(g) \right]. \tag{3.22}$$

3.2.6 Summation of Fields of a Multi-Port Driven Antenna

After solving out the electric field produced by ports on either side of spoke $n = 0$ of an N-spoke MPD antenna $\mathbf{E_0}$ (3.20), a general form of the electric field produced by ports on either side of spoke n ($\mathbf{E_n}$) that takes into account the phase shift of I_s as well as the rotational shift in ϕ of the direction of the current is given by (3.21). The total E-field of the MPD antenna ($\mathbf{E_{MPD}}$) produced using superposition of all the ports is expressed as

$$\mathbf{E_{MPD}} = \sum_{n=0}^{N-1} \mathbf{E_n}. \tag{3.23}$$

This summation is calculated by noting that

$$\frac{\mathbf{E_n} + \mathbf{E_{N-n}}}{E_0} = e^{j\frac{\pi n}{N}}\left[\cos\left(\frac{\pi n}{N}\right)\hat{\mathbf{u}}_\mathbf{x} + \sin\left(\frac{\pi n}{N}\right)\hat{\mathbf{u}}_\mathbf{y}\right] + \\ -e^{-j\frac{\pi n}{N}}\left[-\cos\left(\frac{\pi n}{N}\right)\hat{\mathbf{u}}_\mathbf{x} + \sin\left(\frac{\pi n}{N}\right)\hat{\mathbf{u}}_\mathbf{y}\right], \tag{3.24}$$

which simplifies to

$$\mathbf{E_n} + \mathbf{E_{N-n}} = 2E_0\left[\cos^2\left(\frac{\pi n}{N}\right)\hat{\mathbf{u}}_\mathbf{x} + j\sin^2\left(\frac{\pi n}{N}\right)\hat{\mathbf{u}}_\mathbf{y}\right]. \tag{3.25}$$

(3.23) is broken into two groups based on if N is even or odd:

$$\mathbf{E_{MPD}} = \begin{cases} \mathbf{E_0} + \mathbf{E_{N/2}} + \sum_{n=1}^{N/2-1}\left[\mathbf{E_n} + \mathbf{E_{N-n}}\right] & \text{if } N \text{ is even.} \\ \mathbf{E_0} + \sum_{n=1}^{\lfloor N/2\rfloor}\left[\mathbf{E_n} + \mathbf{E_{N-n}}\right] & \text{if } N \text{ is odd.} \end{cases} \tag{3.26}$$

First, in the case where N is even,

$$\sum_{n=1}^{N/2-1}\left[\mathbf{E_n} + \mathbf{E_{N-n}}\right] = \left(\frac{N}{2} - 1\right)(\hat{\mathbf{u}}_\mathbf{x} + j\hat{\mathbf{u}}_\mathbf{y})E_0. \tag{3.27}$$

Using (3.21),

$$\mathbf{E_0} = E_0\hat{\mathbf{u}}_\mathbf{x} \quad \text{and} \quad \mathbf{E_N/2} = jE_0\hat{\mathbf{u}}_\mathbf{y}. \tag{3.28}$$

Thus,

$$\mathbf{E_{MPD}} = \frac{N}{2}(\hat{\mathbf{u}}_\mathbf{x} + j\hat{\mathbf{u}}_\mathbf{y})E_0 \quad \text{when } N \text{ is even.} \tag{3.29}$$

In the other case when N is odd,

$$\sum_{n=1}^{\lfloor N/2\rfloor}\left[\mathbf{E_n} + \mathbf{E_{N-n}}\right] = \left[\left(\frac{N}{2} - 1\right)\hat{\mathbf{u}}_\mathbf{x} + j\frac{N}{2}\hat{\mathbf{u}}_\mathbf{y}\right]E_0, \tag{3.30}$$

which when added to $\mathbf{E_0}$, yields

$$\mathbf{E_{MPD}} = \frac{N}{2}(\hat{\mathbf{u}}_\mathbf{x} + j\hat{\mathbf{u}}_\mathbf{y})E_0 \quad \text{when } N \text{ is odd.} \tag{3.31}$$

And thus with (3.29) and (3.31), we get

$$\mathbf{E_{MPD}} = \frac{N}{2}(\hat{\mathbf{u}}_\mathbf{x} + j\hat{\mathbf{u}}_\mathbf{y})E_0. \tag{3.32}$$

Finally, inserting (3.22) into (3.32) gives the final solution of $\mathbf{E_{total}}$ in terms of I_s, ω, N, g and r:

$$\mathbf{E_{MPD}} = -I_s \frac{j\mu cN}{4\pi} \frac{e^{-jkR_0}}{R_0} \left[\frac{g}{g^2-1} + \tan(g) \right] (\hat{\mathbf{u}}_\mathbf{x} + j\hat{\mathbf{u}}_\mathbf{y}). \tag{3.33}$$

The critical design insight of the entire analysis comes from this equation which suggests that to have the fields caused by the spokes and the ring add up coherently in the far field, $g/(g^2-1)$ and $\tan(g)$ should have the same sign, which means that $1 < g < \pi/2$. This is interesting as one might suspect that a ring with a traveling current wave on it might be a fairly resonant structure, but in reality, it is quite wide-band, with functional operation anytime the wavelength is less than the circumference of the ring, and greater than the circumference multiplied by $2/\pi$. Also, the constraints on the circumference (g) are independent of the number of spokes N. It is important to note that with this model, the fields go from being completely destructive in their interference to completely constructive at $g = 1$. Using the model of a lossy transmission line, the phase difference between the two circularly polarized fields will shift in the region right around $g = 1$ from π to 0, with the phases when $g = 1$ being in quadrature. The equation also shows that the electric field is proportional to the input source current I_s, and the total number of spokes N, as well as the final term $\hat{\mathbf{u}}_\mathbf{x} + j\hat{\mathbf{u}}_\mathbf{y}$, which yields circular polarization for any values of g or N. The summation of all of the ports produces a traveling current wave around the ring, while producing standing waves at various phases in each ground spoke as expected.

3.2.7 Poynting Vector and Effective Isotropic Radiated Power

The radiated power captured by a recieve antenna is what will ultimately be measured, and converted to the effective isotropic radiated power (EIRP). To calculate EIRP, the Poynting vector ($\mathbf{S}$) in the far field is first calculated and is expressed as [1]

$$\mathbf{S} = \frac{|\mathbf{E_x}|^2 + |\mathbf{E_y}|^2}{2\eta} \hat{\mathbf{u}}_\mathbf{r}, \tag{3.34}$$

where $\eta = \mu c$ is the impedance of free space. The far-field broadside Poynting vector for the MPD antenna, using (3.33) is then given by

$$\mathbf{S_{MPD}} = I_s^2 \frac{\mu cN^2}{16\pi^2 R_0^2} \left[\frac{g}{g^2-1} + \tan(g) \right]^2 \hat{\mathbf{u}}_\mathbf{r}. \tag{3.35}$$

EIRP can be calculated [1] as

$$EIRP = 4\pi r^2 |\mathbf{S}|, \tag{3.36}$$

and thus

$$EIRP_{MPD} = I_s^2 \frac{\mu cN^2}{4\pi} \left[\frac{g}{g^2-1} + \tan(g) \right]^2. \tag{3.37}$$

As expected the EIRP is related to the square of the electric field, and maximizing the electric field will also maximize the power flux in the far field.

3.3 Design and Simulation of a Multi-Port Driven Antenna

Electromagnetic simulation is used first to verify the design insights of the MPD antenna in free space, and then to model non-idealities that cannot be fully analyzed by hand. Figure 3.6 shows the instantaneous current at 0 phase on the ring and on the spokes with a circumference of 0.7λ, or $g = 0.7$ with just ports 0 and N being driven. At this instant, port 0 is driving maximum current into the ring, and port N is pulling maximum current from the ring. As expected when $g < 1$, the currents on each half of the ring all point in the same direction along the ring (counterclockwise for $\phi = 0$ to π, and clockwise for $\phi = -\pi$ to 0). This produces an overall integrated current vector pointed in the $-\hat{\mathbf{u}}_\mathbf{x}$ direction. Similarly the current along the entire spoke also points in the same direction, in this case in the $\hat{\mathbf{u}}_\mathbf{x}$ direction. This ends up looking similar to the currents observed on a transmission line, and as expected, it does not radiate effectively.

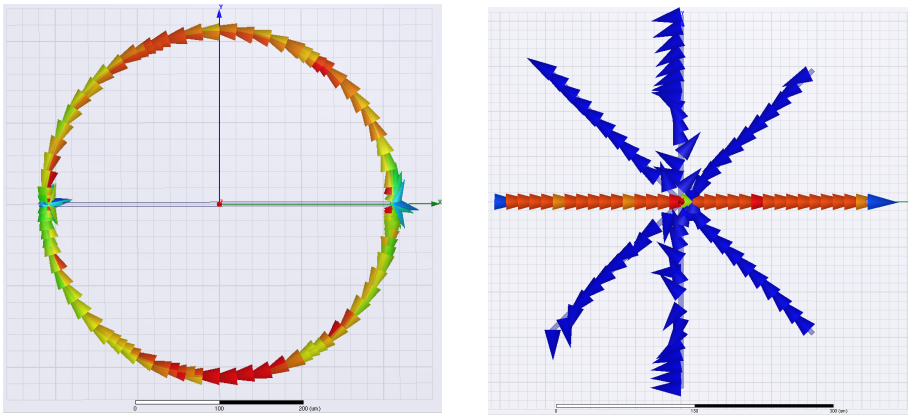

(a) **(b)**

Figure 3.6: Simulated instantaneous current at phase 0 on a ring (a) and spokes (b) driven by ports 0 and N, when $g = 0.7$.

In contrast, Figure 3.7 shows the same instantaneous current at 0 phase on the ring and spokes but for the case of $g = 1.3$. The spokes look similar to Figure 3.6b, but because g is now greater than 1, a node appears in the current pattern for the ring, and the current changes direction along the ring between the port on the x axis and the maximum of the current on the y axis. This means that the integrations of both the spokes and the ring produce vectors pointed in the $\hat{\mathbf{u}}_\mathbf{x}$ direction and thus will radiate effectively. It also is noted that current on the spokes that are not being driven for this part of the superposition is much less than that of the spokes being driven, by more than two orders of magnitude, so the assumption to neglect coupling between adjacent spokes and assume the non-driven spoke current was 0 is shown as reasonable.

The biggest non-ideality that the previous analysis ignores was the effect of the substrate on the radiation. Due to the desire to not require any external cooling elements such as fans, the thermal paths for heat generated on-chip need to be considered. This path was created by designing the chip to be mounted on a ground

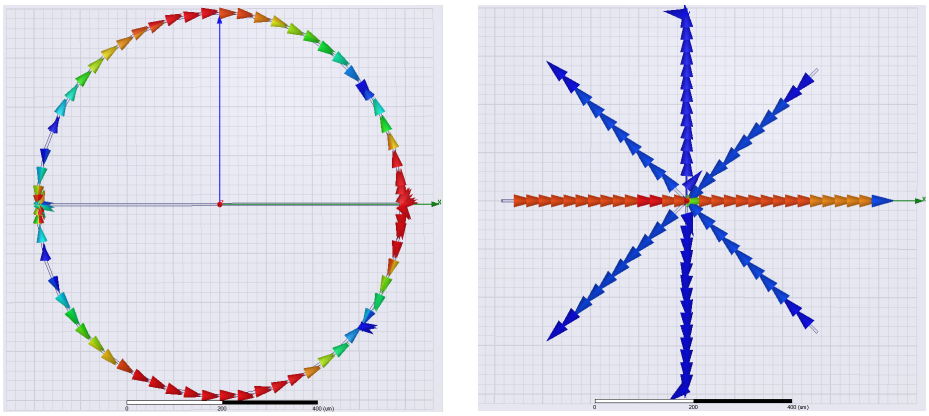

Figure 3.7: Simulated instantaneous current at phase 0 on a ring (a) and spokes (b) driven by ports 0 and N, when $g = 1.3$.

plane on a printed circuit board (PCB). Because of the high thermal conductivity of the silicon substrate and the ground plane that it is mounted on, this strategy enables continuous operation of the radiator without performance degradation due to temperature increases. However, because the dielectric constant of silicon is 11.68 times higher than that of air, most of the power initially radiated by the antenna located on the boundary of the two materials is directed down into the substrate initially [12, 74]. The electric fields reflect off of the bottom ground plane and radiate upwards. In order to have this reflected wave add coherently with the fields radiated directly upwards, the substrate is designed to be about $1/4$ wavelength in silicon thick. This will provide a $\pi/4$ phase shift as the wave travels down, a $\pi/2$ phase shift off of the conductive reflection, and a $\pi/4$ phase shift traveling back up, for a total phase shift of 2π.

Another design concern had to do with feeding the driver circuity DC power. To increase symmetry, the driver circuitry is placed at the center of the antenna, as is explained in more detail in Section 3.4. The DC feed lines need to be run to that core without harming the radiation, so the ground spokes are extended out from the ring to a ground plane that is pulled back from the ring by around $\lambda/4$, as seen in Figure 3.8. This allows DC power lines to run from the pads located around the outside of the chip, underneath the ground spokes to the driver core. Running these lines directly under the spokes shields them from the radiation in a similar manner to a transmission line. The ground plane is pulled back around $\lambda/4$, which produces the highest impedance seen on the spokes at each drive point looking outward, which means that almost all of the ground spoke RF current goes toward the center of the antenna, as assumed in the analysis, and shown in the blue arrows of Figure 3.8.

The two sources of loss for this antenna are loss due to substrate modes, and conductive loss in the antenna metals themselves. To separate out these two effects, electromagnetic simulations are done with both lossless and lossy metal conductors and a lossy substrate. The resulting antenna pattern with lossless metals is

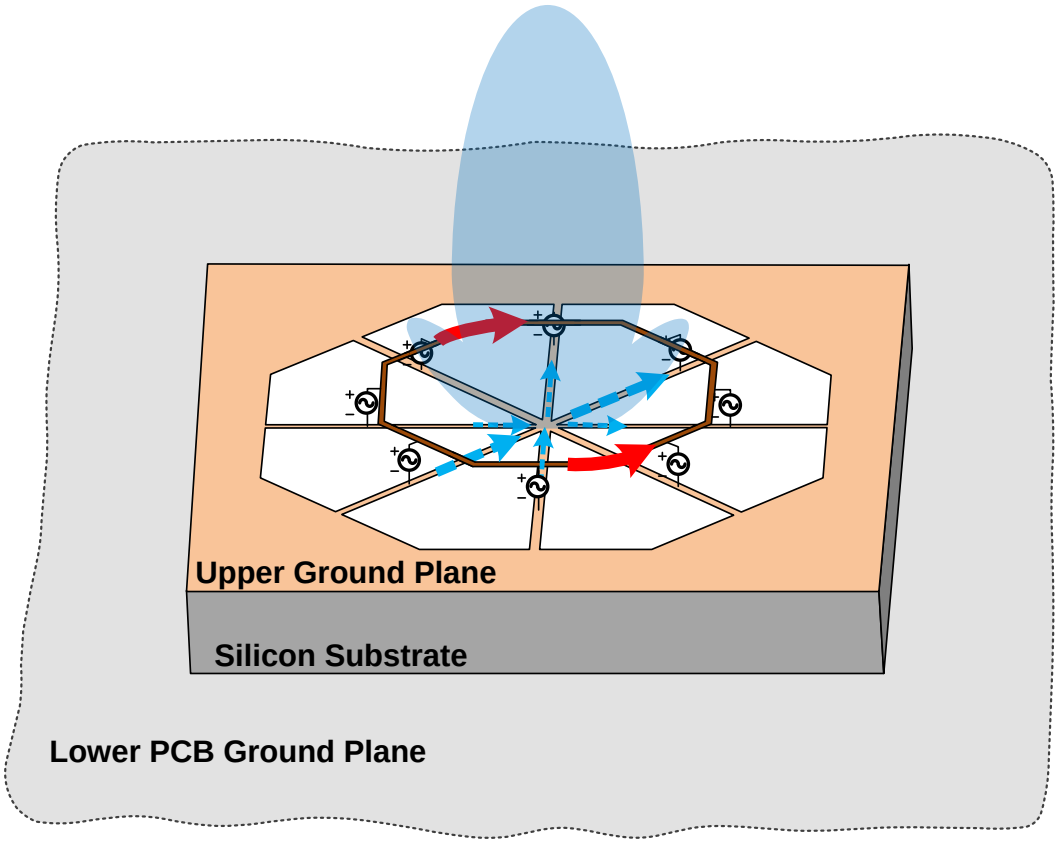

Figure 3.8: MPD antenna with substrate mounted on a PCB ground plane, with traveling wave currents on the ring marked as solid red arrows, and standing wave ground spoke currents marked with dashed blue arrows, and an overlay of the radiated beam.

shown in Figure 3.9, with 30% antenna efficiency, 9.7 dBi maximum directivity, and 4.4 maximum dBi gain. Including the metal losses those numbers dropped slightly to a 28% antenna efficiency, 9.1 dBi maximum directivity, and 3.6 dBi maximum gain. This shows that the dominating loss mechanism for this design is the loss due to the substrate.

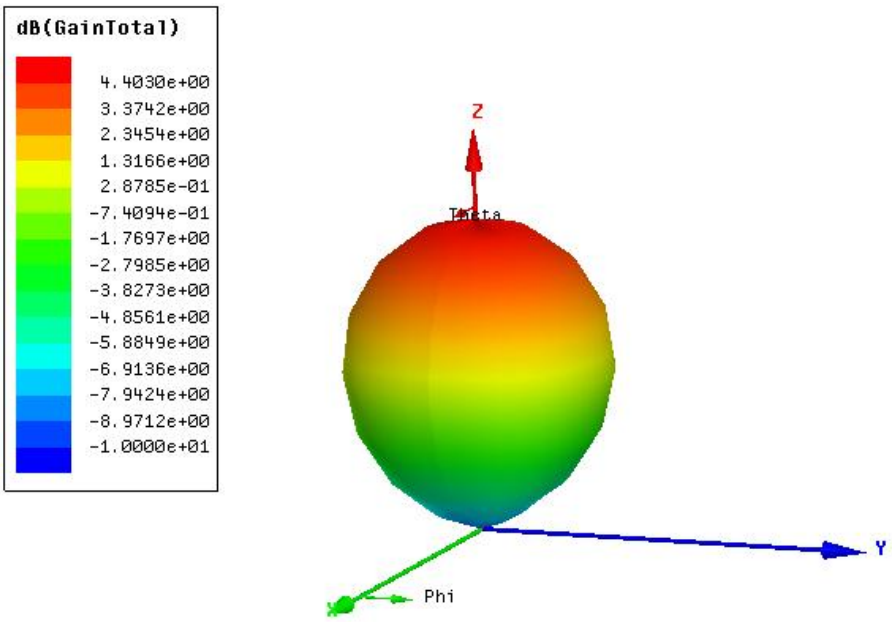

Figure 3.9: Simulated radiation pattern of the gain in dBi of the MPD antenna of Figure 3.8 with lossy 10 Ωcm substrate and lossless conductors shows a maximum 4.4 dBi gain.

3.4 Driver Circuitry of a Multi-Port Driven Radiator

The driver circuitry of this MPD radiator consists of an 8-phase ring oscillator, with each phase feeding a power amplifier stage that drives a port of the MPD antenna. This unilateral driven architecture, with large isolation between the reactive elements of the radiating structure and the oscillator tank, is important to prevent blocker signals that are nearby in the environment from coupling into the oscillator tank and either pull or, in the worst case, lock the oscillator to an undesired frequency, which is a concern with self-oscillating radiating structures where the same reactive elements are used for both oscillation and radiation. Maintaining symmetry to keep all of the phases correct was also critical, so the circuitry was placed at the very center of the antenna, in the driver core, as shown in the block diagram of Figure 3.10.

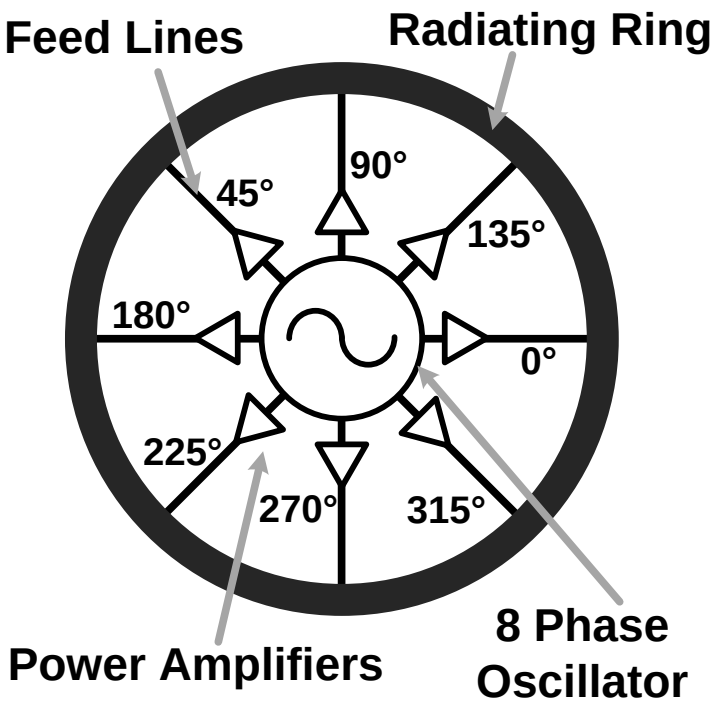

Figure 3.10: Block diagram of the MPD radiator showing driver circuitry blocks as well as ring of MPD antenna (ground spokes omitted for clarity).

3.4.1 Oscillator Design

The oscillator is an 8-phase tuned ring oscillator at 160 GHz with bipolar transistors, because an MPD antenna with $N = 8$ requires 8 phases, each 45^o apart. Due to the extremely limited layout space at the core, there is only room for a single stage of amplification after the oscillator, and thus the oscillator is designed to have each phase provide -13 dBm to the amplifier stage. In order to accomplish that, and to increase the gain at these frequencies, a differential cascode topology was used (Figure 3.11).

A small signal simulation of the cascode stage shows that the optimum power gain (G_m) of a single stage occurred with a voltage gain (A_{opt}) of 5.5 and a phase shift ϕ_{opt} of 28^o [75]. Normally in a ring oscillator, the voltage gain of each stage at steady state is required to be 1, but using series capacitance, a captive divider is created that allows for voltage gains for the transistor stages to be greater than one. Once at steady state, the phase shift is 45^o, which is not far off from the optimal, and the oscillator is designed with a voltage gain per stage of $A = 2.0$. This is because above that amount some non-linear effects that the small signal simulations did not take in to account degraded the performance. Also, as can be seen from Figure 3.12, most of the improvement in G_m at 160 GHz from $A = 1$ to $A = 5.5$ is realized by only increasing A to 2.0.

The tuning range of an 8-phase ring oscillator is relatively narrow, and due to the low quality factor of varactors at these frequencies, varactor frequency tuning is not used in order to increase loop gain and ensure startup. Thus the only frequency tuning is achieved by adjusting the biasing points on the bases of

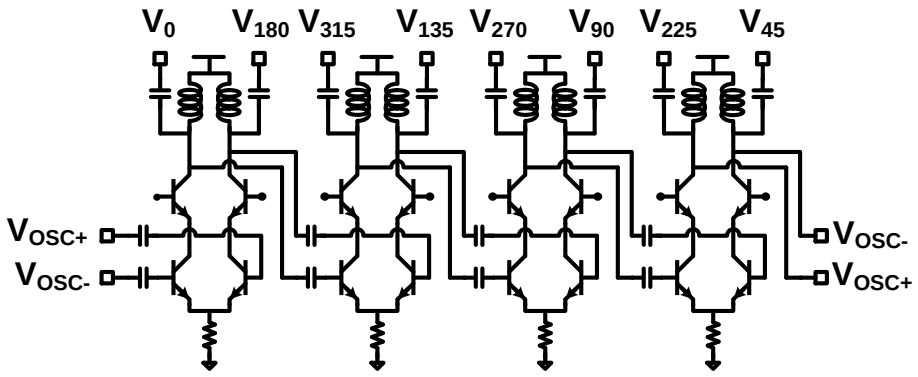

Figure 3.11: Schematic of the 8-phase ring oscillator that uses AC coupling capacitors to allow a cascode architecture as well as to create capacitive dividers that enable single stage voltage gains greater than 1 even at steady state. Bias circuitry omitted for clarity.

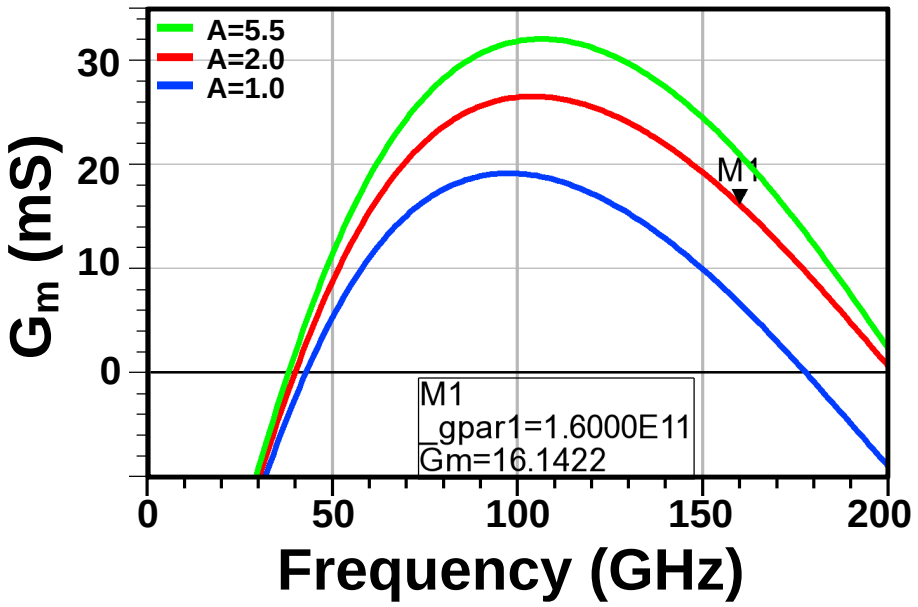

Figure 3.12: Plot of the power gain G_m of a single cascode stage with voltage amplitude gains (A) of 1.0, 2.0 and 5.5, with voltage phase shifts of 45^o. At 160 GHz, with $A = 2.0$, $G_m = 16.1$ mS.

the common source and cascode stages. In order to demonstrate the wide-band nature of the antenna, laser trimmable differential transmission line inductors are used for the tanks of the oscillator. By trimming off the parts of the inductor indicated in Figure 3.13, the effective length of the transmission line is increased, which increases the inductance and decreases the frequency. This enables 3 separate bands to be measured from the same design, at 145 GHz, 152.5 GHz and the default 160 GHz.

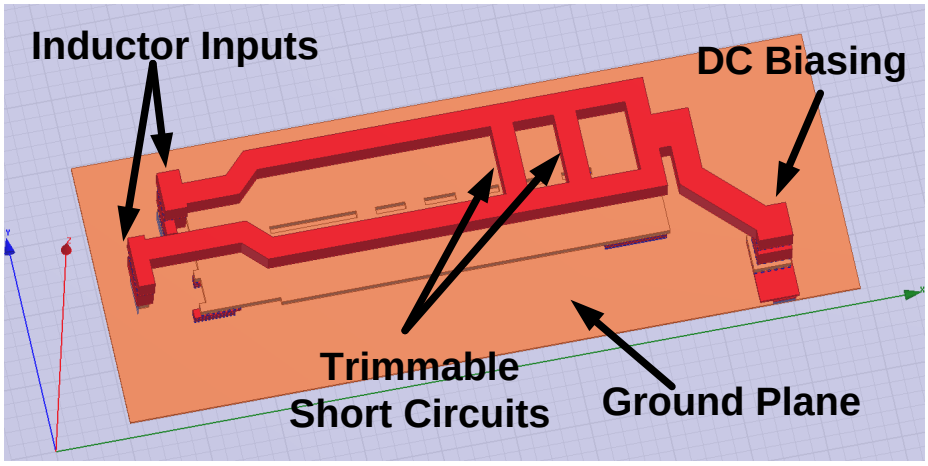

Figure 3.13: 3-dimensional rendering of the oscillator inductors implemented as differential transmission lines that have trimmable short circuits that can be cut to enable 3 different frequency bands.

3.4.2 Amplifier Design

The 4 differential power amplifier pairs are implemented as class A amplifiers to maintain maximum power gain. The schematic of one of the differential pair amplifiers is shown in Figure 3.14. Custom metal coupling capacitors are used both between oscillator stages as well as in the interstage matching networks between the oscillator and amplifier to enable cascode designs and provide greater power gain from a single stage. By controlling the capacitance value of the series metal capacitor of the interstage match, the portion of the power that is delivered to the amplifier stage compared to the next oscillator stage can be controlled. Making the capacitor too large draws too much power from the oscillator, requires lowering the value of the oscillator inductance to maintain oscillator frequency, reduces the loop gain and lowers the overall power delivered to the amplifier. Conversely, a capacitor that is too small will not draw enough power from the oscillator, so while the oscillator will have large swings, not enough of the power will be sent to the amplifier. The optimum found for this design is a capacitance of 17 fF between oscillator stages, and a capacitance of 12 fF for the interstage network between the oscillator and amplifiers.

Layout constraints within the driver core mean that the available room for bypass capacitors within the core is minimal. To help alleviate that, virtual short circuits in the design were taken advantage of wherever possible. This means that differential pairs were kept next to each other in layout, and thus little bypass

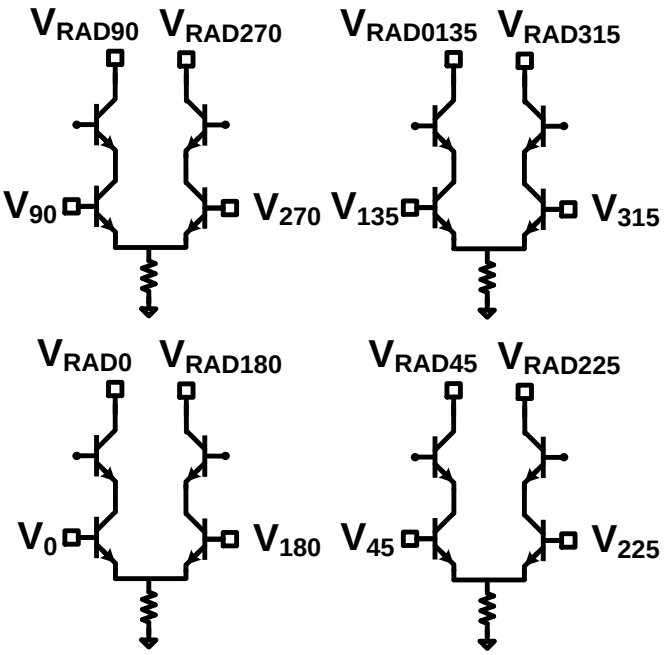

Figure 3.14: Schematic of the class A cascode power amplifier stages that each produce -7.0 dBm at 160 GHz for a total available power to the MPD antenna of 2.0 dBm in simulation. Bias circuitry omitted for clarity.

is required for the biasing of the cascode transistor's base node, or on the differential inductor V_{cc} of the oscillator.

The optimum output load of each amplifying stage is $Z_{opt} = 24 + 45j\ \Omega$ at 160 GHz. By adjusting the overall size of the antenna ring, as well as the widths of the spokes and the feed lines for the antennas, each transistor is loaded with an impedance that produces within 1 dB of the optimum output power without any explicit lossy output matching elements. In simulation, each amplifier provides -7.0 dBm at 160 GHz, which corresponds to a total of 2.0 dBm input power into the antenna from all 8 ports.

Biasing of the amplifying stages is done through the antenna. A $\lambda/4$ transmission line was extended out on one spoke from the radiating ring and is AC shorted with bypass capacitors at the other end. This allows for a DC current path to all 8 amplifiers through the ring and feed lines while providing a high impedance at the point of contact with the ring, thus minimizing the effect of the line on the radiation.

3.5 Top Level Design and Simulation of a Multi-Port Driven Radiator

As mentioned in the previous section, the differential pairs for both the oscillator stages as well as the amplifying stages are kept next to each other in layout to conserve the virtual short circuits. Differential signals that are next to each other at the amplifier output need to be routed through feed lines to opposite sides of the driver core, and in order to maintain both phase and amplitude relationships in the antenna, feed lines where

the driver stage is close to the antenna port need to be meandered, while others have to be routed across the core, as can be seen in Figure 3.15. Electromagnetic simulations of all of the metal structures in the driver core are necessary to ensure amplitude and phase matching. These simulations also include the oscillator inductors to ensure isolation between the feed lines and the inductors, as avoiding unwanted feedback paths is critical. Including all of these effects, the top level antenna design with the feed lines at 160 GHz is simulated

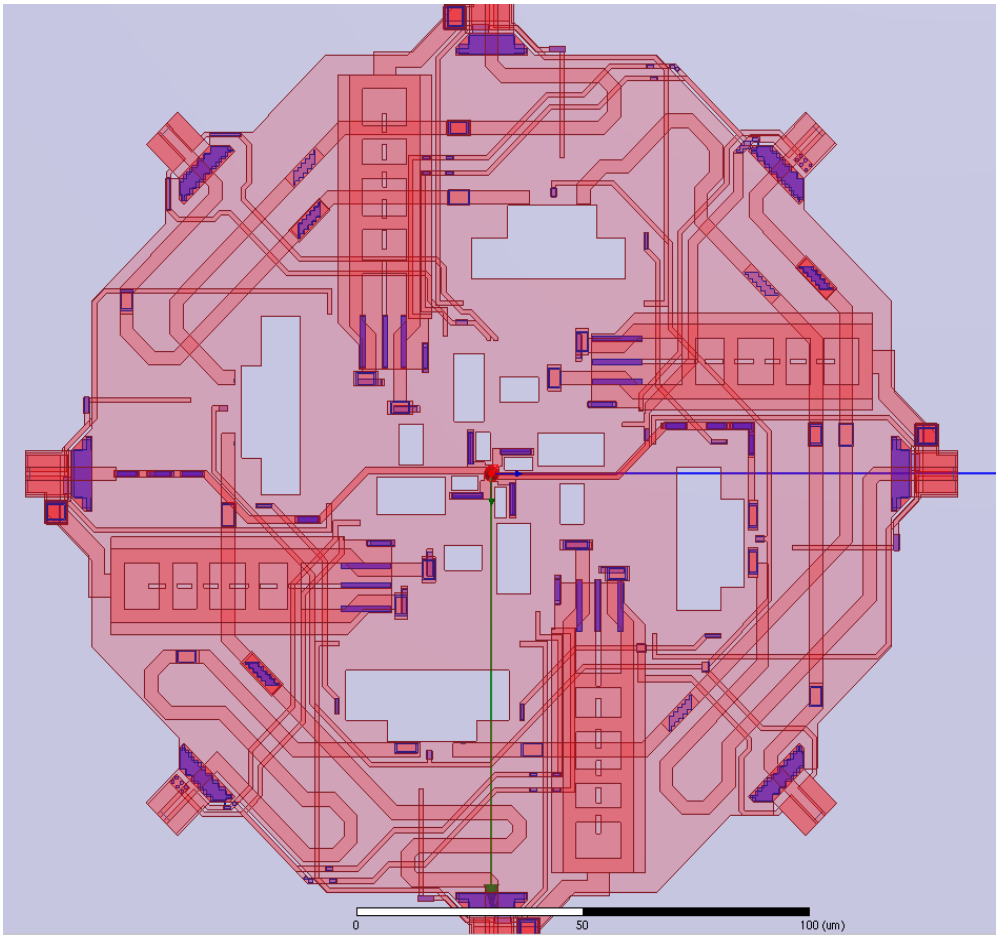

Figure 3.15: Layout for the driver core feed lines electromagnetic simulation. All 8 feed lines, as well as the 4 differential oscillator inductors and all ground lines were simulated to ensure amplitude and phase matching for the feed lines, and isolation between feed lines and oscillator inductors.

to have 8.9 dBi maximum directivity, 2.4 dBi gain, 23% radiation efficiency, and a radiation pattern shown in Figure 3.16. The gain and radiation efficiency numbers include all of the feed lines and vias down to the transistor levels of the amplifying stages, and encompass what would traditionally be the loss of the antenna, matching networks, and any loss due to power transfer between the amplifier stages and the antenna. With a simulated output power of 2.0 dBm from the 8 amplifiers, the simulated EIRP is 4.4 dBm, for a total output radiated power of -4.4 dBm.

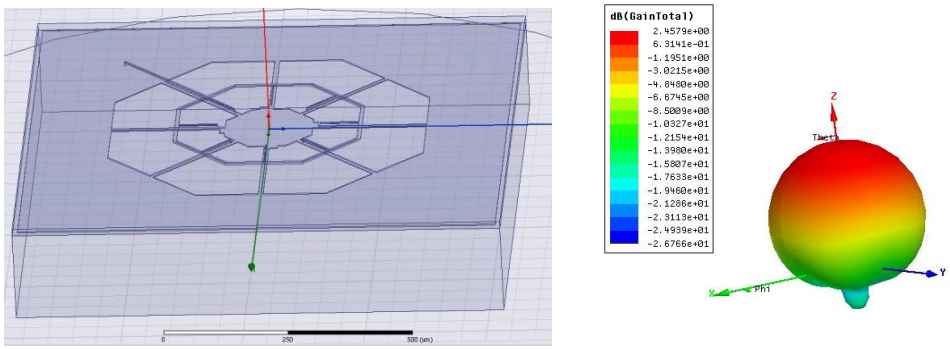

Figure 3.16: The feed line simulation was coupled with an antenna simulation that included the effects of the driver core ground and the PA biasing line (a), with the resulting radiation pattern of the gain in dBi in (b) showing a maximum gain of 2.4 dBi.

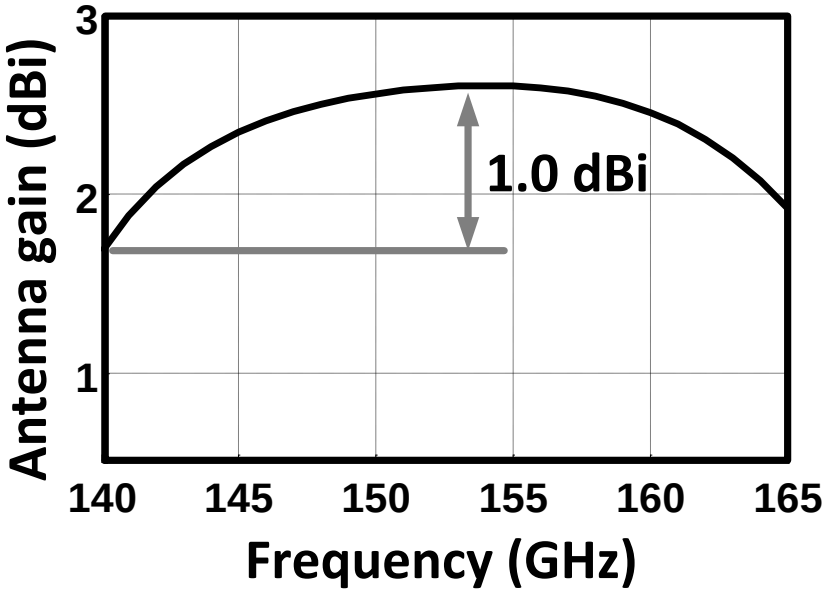

Figure 3.17: Simulated gain of the MPD antenna including feed lines from 140 to 165 GHz shows a maximum gain deviation of 1 dBi across that band and demonstrates the expected wide band performance of the MPD antenna.

3.6 Multi-Port Driven Radiator Measurements

The design was fabricated in a 0.13 μm SiGe BiCMOS process, occupies an area of 1x1 mm^2 and is thinned to 185 μm. The chip is placed on a PCB that is mounted on a 2-dimensional stepper motor (Figure 3.18a). The radiated power is detected by a 23.5 dBi gain linearly polarized horn antenna placed 75 mm (40λ) away, that feeds either a tenth harmonic mixer connected to a spectrum analyzer for frequency measurements, or to an Erikson power meter for absolute power measurements, as seen in Figure 3.18b. The chip is powered by a 3.3 V DC source that draws 117.5 mA. No external optics of any kind (e.g. silicon lens or dielectric superstrate), or any external fans for cooling are used.

The spectrum of the broadside radiation is measured using the tenth harmonic mixer and the calibrated output is shown in Figure 3.19. The chip radiates at 161.45 GHz, and the power that the receive antenna captures is -25.8 dBm. Taking into account the receive antenna's distance from the chip and the gain of the receive antenna, this corresponds to 4.6 dBm EIRP.

Next, the polarization of the broadside radiation is measured by rotating the chip around the broadside axis, and measuring the received power of the linearly polarized antenna to verify circular polarization. A good metric to report this is the polarization ratio, which is the ratio of the electric field in the clockwise polarization over electric field in the the counterclockwise polarization. This is calculated (as derived in Appendix B) from the ratio of the maximum power received by the linearly polarized antenna over the minimum power received as

$$\text{Polarization Ratio} = \frac{|\mathbf{E_{CW}}|}{|\mathbf{E_{CCW}}|} = \frac{\sqrt{\frac{P'_x}{P'_y}} + 1}{\sqrt{\frac{P'_x}{P'_y}} - 1}. \tag{3.38}$$

The measured ratio of maximum to minimum received power is 1.315, which corresponds to a polarization ratio of 14.6 (Figure 3.20). This shows that the chip is radiating almost entirely in the clockwise polarization.

The antenna pattern is measured by rotating the chip in two dimensions with the stepper motors. Measurements are taken across the radiating half-space twice with perpendicular polarizations. Measurements and simulations of two slices of the pattern are shown in Figure 3.21. The radiation is single-beam, which matches well with the simulations, and the beam-width of the radiation is measured to be slightly wider than the simulations predicted. However, the overall radiated power is larger, which kept the maximum EIRP close to the simulated value. The larger radiated power is likely due to an overestimation of the antenna's metal losses at 160 GHz.

The output power verses frequency of the three frequency bands is shown in Figure 3.22. The tuning within each band is accomplished through the biasing of the bases of the oscillator transistors, while the frequency band was tuned by laser trimming the oscillator inductors. All three frequency bands were close to their designed frequency, and can radiate more than 0 dBm EIRP. The reason for the lower power levels

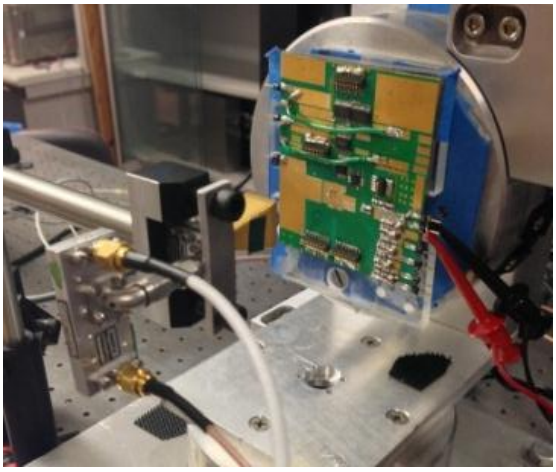

(a)

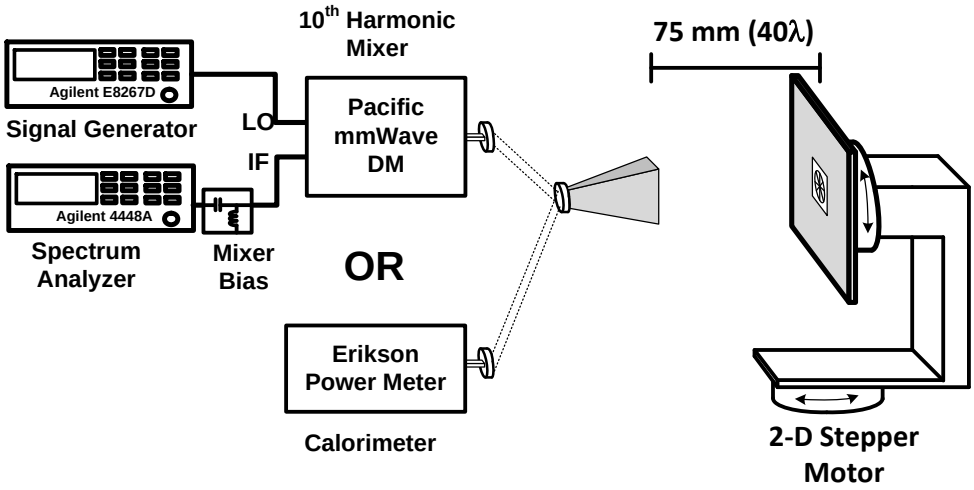

(b)

Figure 3.18: Photo of the measurement setup with EM dampers removed and the harmonic mixer attached to the receive antenna (a), and diagram of the measurement setup where received power goes either to a harmonic mixer for frequency measurements or to an Erikson power meter for absolute power measurements (b).

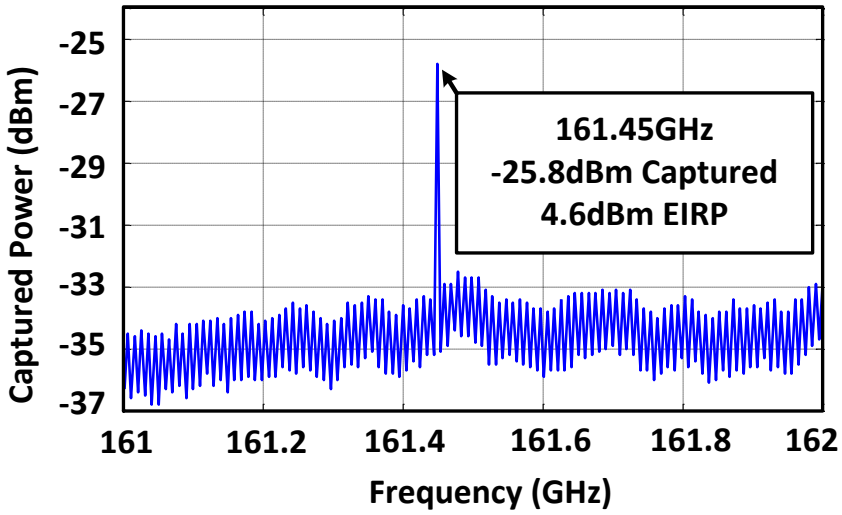

Figure 3.19: Measured spectrum of the single element MPD antenna shows -25.6 dBm captured power at 161.45 GHz which corresponds to 4.6 dBm EIRP.

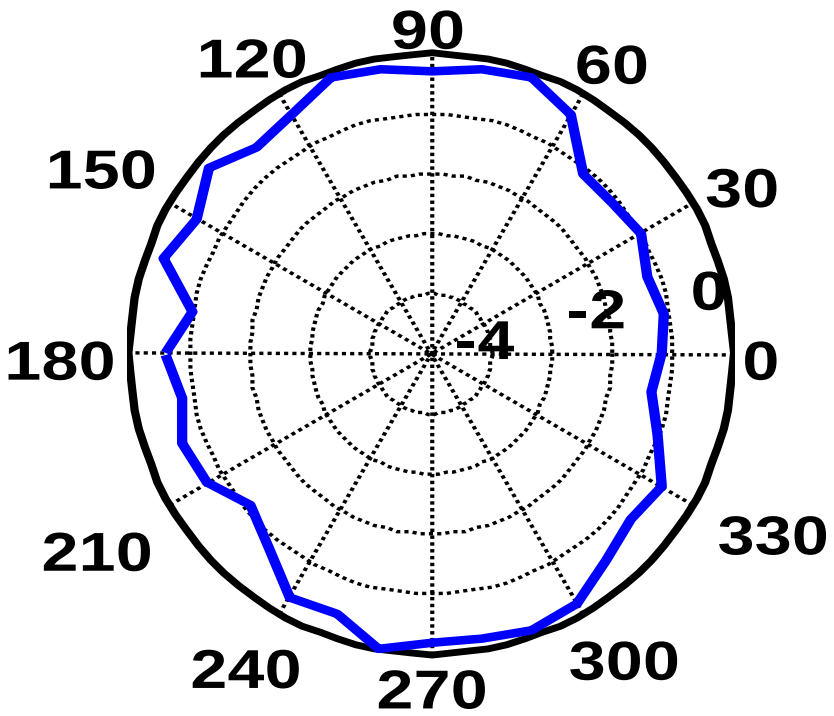

Figure 3.20: Normalized power captured in dB by a linearly polarized antenna as the chip is rotated in the X-Y plane. Ratio of maximum to minimum powers of 1.315 imply a polarization ratio of 14.6, which means that the radiation is almost entirely in the CW circular polarization.

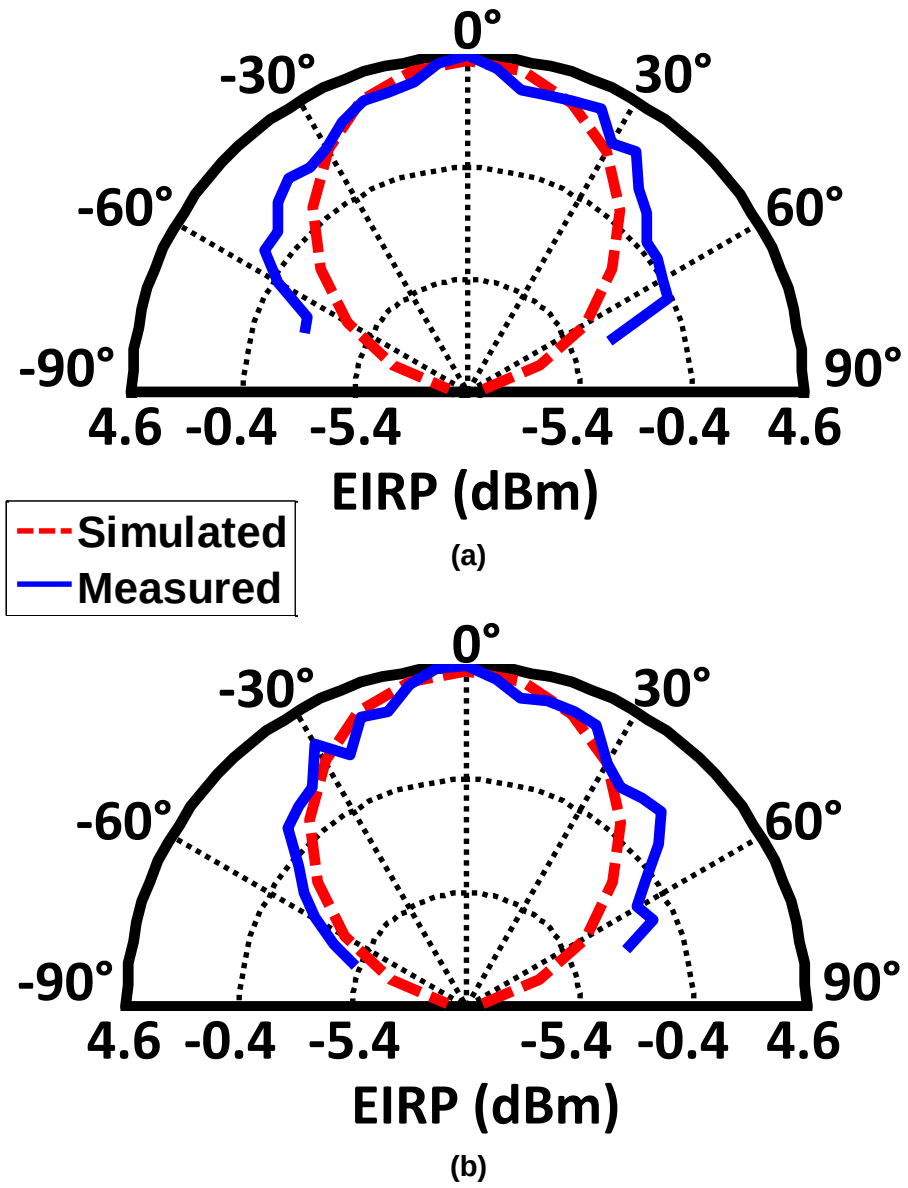

Figure 3.21: Simulated and measured radiation patterns of EIRP in dBm in the elevation (a) and azimuth (b) planes show a single beam radiating broadside to the chip that matches well with simulation.

Table 3.1: Comparison of integrated radiating sources in silicon above 110 GHz without external lenses.

Reference	[33]	[34]	[38]	[8]	[37]	**This Work**
Frequency (GHz)	191	280	165	260	210	161
Radiated Power (dBm)	-12.4	-7.2	-27	n/a	n/a	-2.0
Maximum EIRP (dBm)	-1.9	9.4	n/a	5	5.13	4.6
Number of Elements	4	16	1	4	4	1
EIRP/Element (dBm)	-7.9	-2.6	n/a	-1.0	-0.9	4.6
Process	65 nm CMOS	45 nm CMOS SOI	130 nm SiGe	65 nm CMOS	32 nm SOI CMOS	130 nm SiGe
DC Power (mW)	77	817	800*	688	240	388
Area (mm^2)	1.1	7.3	1.3**	6**	3.5	1

*Power includes transmitter and receiver. **Area includes transmitter and receiver.

in the lower two bands is that the matching between oscillator stages as well as between the oscillator stages and the power amplifiers was tuned for the 160 GHz band.

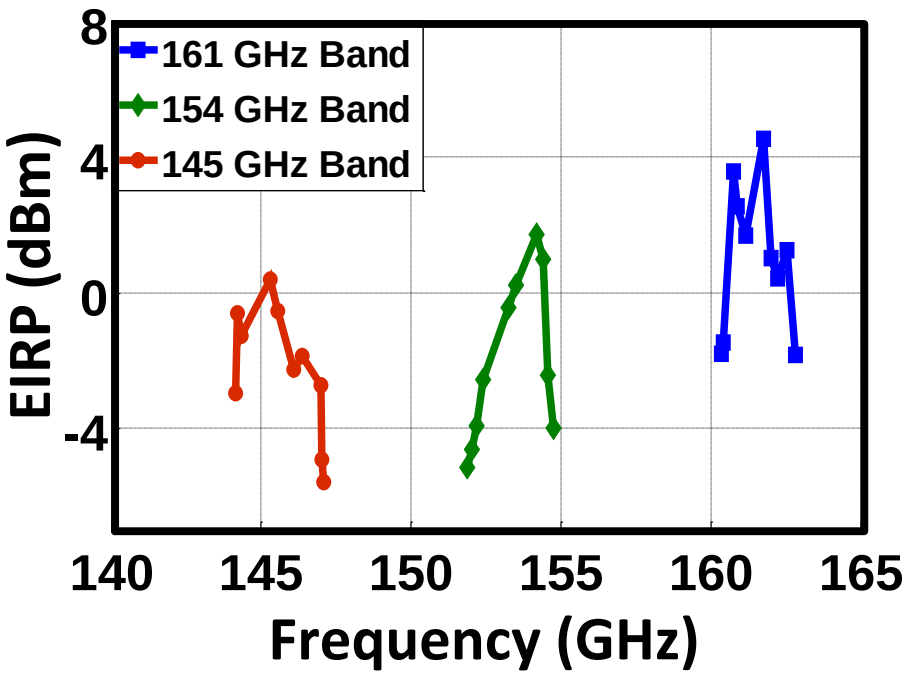

Figure 3.22: Measurements from the three frequency bands at 145, 154 and 161 GHz show radiated EIRPs above 0 dBm for all 3 bands.

A die photo of the chip is shown in Figure 3.23, and a performance summary with comparison to other works above 110 GHz without external optics is presented in Table 3.1.

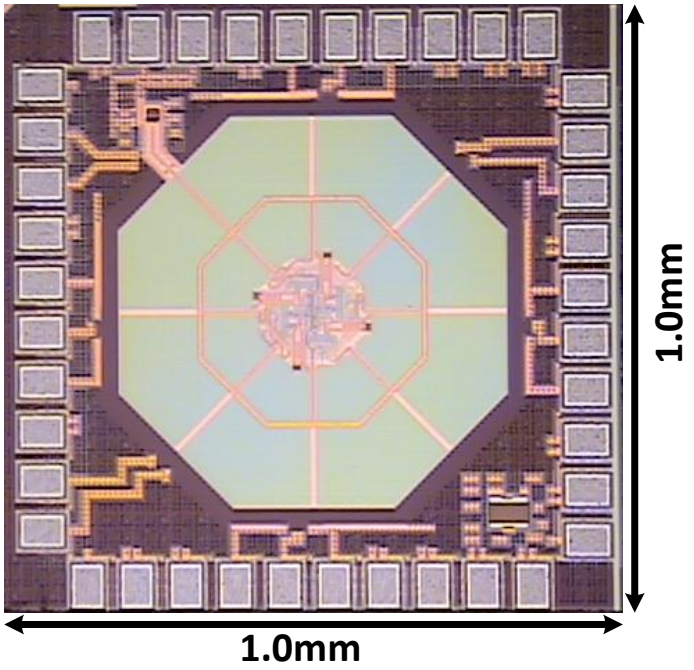

Figure 3.23: Die photo of the single element MPD radiator with area of 1 mm^2.

3.7 Slot Based Radial Multi-Port Driven Radiator[1]

Another version of the radial MPD radiator can be achieved by using electromagnetic duality theory [76], which enables a slot antenna to be designed and analyzed by introducing a fictitious magnetic current that flows through the slot, and then comparing that to a line antenna that is created by inverting the entire antenna plane's metal. The lines of the antenna now become slots, and the open space around the lines become metal. The E and H fields of the slot antenna are then replaced by the vectors H and -E for the line antenna. And through this relationship, the fields of a single linear slot cut into a conducting plane is identical to the fields of the dipole.

This same duality technique is used to create the radial slot MPD radiator. In this case, the antenna that we are going to invert is simply the ring of the radial MPD, and the spokes will be neglected. This means that when inverted, the antenna becomes a slot ring as shown in Figure 3.24, where there is an internal metal plane that is surrounded by a thin slot ring, which itself is surrounded by a large outer metal ground plane. The antenna is then excited by driving one metal plane against the other at multiple points around the ring. This configuration has advantages and disadvantages compared to the line antenna version of the radial MPD radiator. One of the main advantages is that the ground plane provides shielding for any other circuitry for

[1]This project was a joint project with my colleagues Amirreza Safaripour and Kaushik Dasgupta. I was responsible for the antenna design and electromagnetics, and the locking network distribution for the array, and Amirreza was in charge of the locking network oscillator and circuitry within the radiator core as well as the digital signal routing for the radiator array. Kaushik designed the digital to analog converters.

the vast majority of the chip, reducing the area footprint on-chip of the antenna to a fraction of the area that the line-based radiator requires due to pulling back the ground plane by $\lambda/4$.

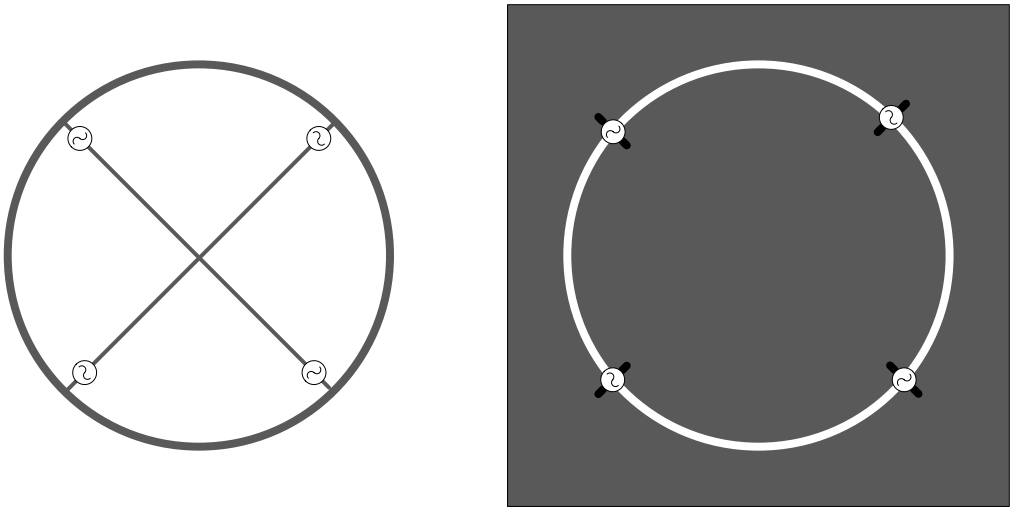

Figure 3.24: Comparison between the line version of the MPD radiator (left), and its dual, the slot MPD radiator (right). Each of the four drive points drive the outer metal plane against the inner metal plane and are phased evenly ($0°$, $90°$, $180°$ and $270°$) to create a circularly polarized radiated field.

A disadvantage of this design is that because the entire surface of the chip except for the thin slot is covered in metal, reflecting the wave off of a lower ground plane and radiating upward is not possible, so this radiator must be designed to radiate out the bottom. Techniques to effectively deal with providing a robust thermal path for heat dissipation are still an active area of research.

The feeds for this antenna operate by using micro-strip transmission lines whose ground reference is the inner metal plane, and directly connecting the signal line of the T-line to the outer metal plane. This is a simple way to solve the antenna feeds, but it brings up some DC biasing issues. Ideally both of the metal planes that make up the antenna would be ground planes at the same DC bias point. However, the signal line of the drive circuitry needs to be biased at the supply voltage, so when the signal line is shorted to the outer metal plane, a solution for the DC biasing needs to be found. Two different approaches have been used in structures that are currently out for fabrication. For the first case, the inner metal plane is grounded, and the outer plane is a VDD supply plane. The second case implements a 2x2 slot radiator array and involves using series AC coupling capacitors in the feed transmission lines, and uses a special tri-level transmission line to keep both metal planes at the same DC bias point without affecting the RF current flow. Both of these were designed at 120 GHz in a 32 nm SOI process and are currently out for fabrication. A brief description of each and simulation results follows below.

3.7.1 DC Coupled Slot Multi-Port Driven Antenna

The driver circuitry for both of these designs is identical to the circuitry designed originally for the dynamic polarization controlling radiator array described in Chapter 4. In this case, the oscillator within the single radiator is free-running and not locked to any other reference oscillator. The block diagram of this radiator is shown in Figure 3.25. The antenna is simulated to have a maximum gain of 3.0 dBi, with a radiation pattern shown in Figure 3.26 and a radiation efficiency of 39%. While this arrangement for the DC voltages works for a single radiator with all circuitry within the inner ground metal plane, it poses problems when an array is desired or other circuitry requires the outer metal plane to be at DC ground. The second design addresses these issues.

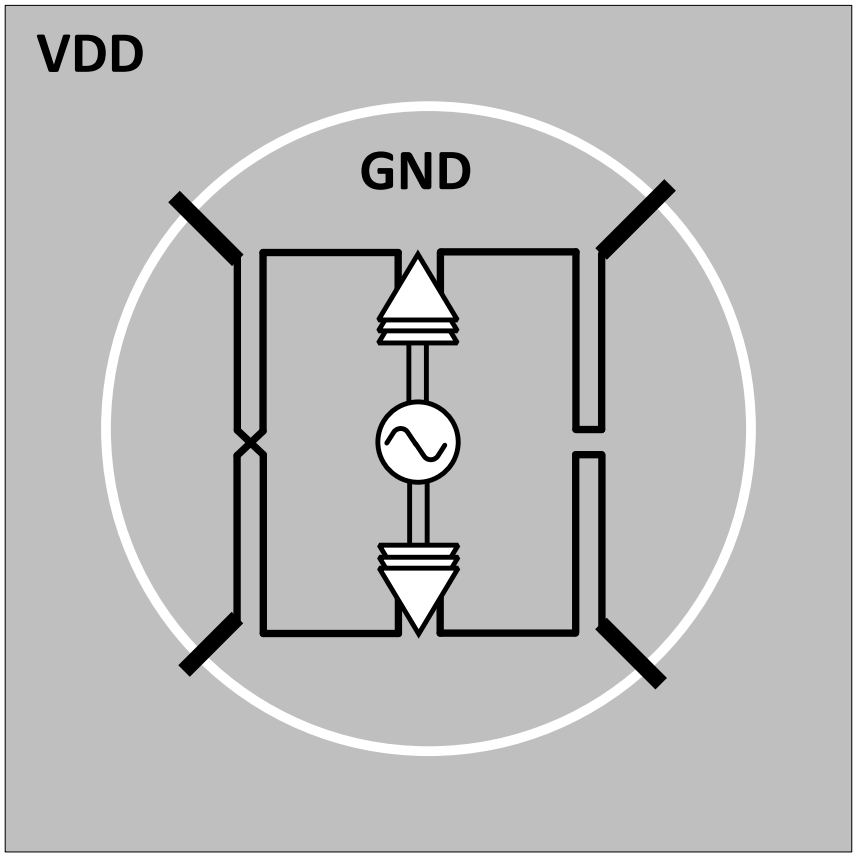

Figure 3.25: Block diagram of the DC coupled slot radiator, where the outer metal plane is biased at VDD, and the inner plane is biased at 0 V.

3.7.2 2x2 AC Coupled Slot Multi-Port Driven Radiator Array

To incorporate this design into an array, a locking signal needs to be generated outside of the radiator and ground referenced to the outer metal plane and delivered to an oscillator that is ground referenced to the inner

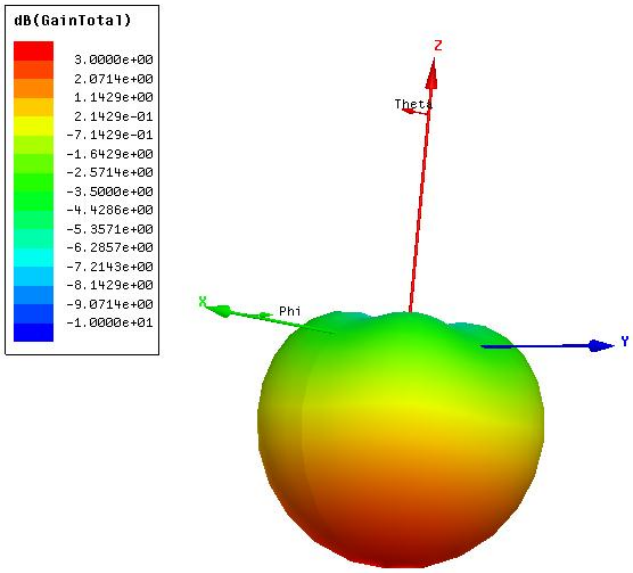

Figure 3.26: Antenna pattern of the DC coupled slot MPD radiator showing a maximum gain of 3 dBi with downward radiation.

metal plane. In order to accomplish this, a consistent reference path must be made available to the transmission line across the entire path from locking oscillator to locked oscillator. This poses a particular problem at the slot itself, where the ground plane is cut. To rectify this, a quarter wavelength tri-level transmission line scheme has been designed. The two reference planes are connected to each other using a $\lambda/4$ transmission line shorted to the outer ground plane. This means that at the connection to the inner ground plane, the impedance seen looking out is very high (an open circuit if losses in the metal are neglected). This open circuit at RF means that the RF performance of the antenna will not be affected much. At the same time, it provides the ground biasing to the inner metal through the signal line of the transmission line. It also means that there is a solid piece of metal that connects from the ground of the locking oscillator to the ground of the radiator's oscillator, and that metal can be used as the unbroken reference for the locking signals to go through. For the part of that path that is the $\lambda/4$ transmission line, there will be three levels. The lowest will be the normal outer ground plane that stops at the slot. The signal line of the $\lambda/4$ transmission line is in the middle and references to the lower metal, and finally the locking signals run on the top and are referenced to the middle line. A 3-dimensional depiction and a cross section of this tri-level transmission line is shown in Figure 3.27.

Thus signals can be passed from outside the radiator to inside the radiator, and DC ground connections are made as well. To increase the symmetry of the radiator, and to enable other DC biasing to be passed to the inside of the radiator, the tri-level transmission line structure is repeated on all four sides of each radiator,

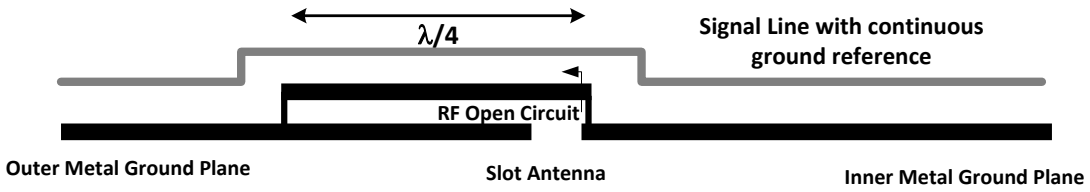

Figure 3.27: Cross section of the tri-level transmission line that maintains an RF open circuit between the two ground planes while simultaneously providing a continuous ground reference for the locking network signals, and a DC ground path for current to flow.

as can be seen in the 3D depiction of the 2x2 slot radiator array shown in Figure 3.28. The microprocessor controls all of the digital to analog converters (DACs) that feed the bias nodes to control amplitude and phase within the cell. Having both metal planes connected means that the output of the radiator amplifiers can no longer be directly connected to the outer metal plane, and thus series AC coupling capacitors are used to AC couple the drivers to the antenna. The maximum gain of this antenna array is simulated to be 5.4 dBi with the simulated radiation pattern shown in Figure 3.29. Considering an expected output power from the amplifiers combined of 20 dBm, this yields a simulated EIRP of 25.4 dBm.

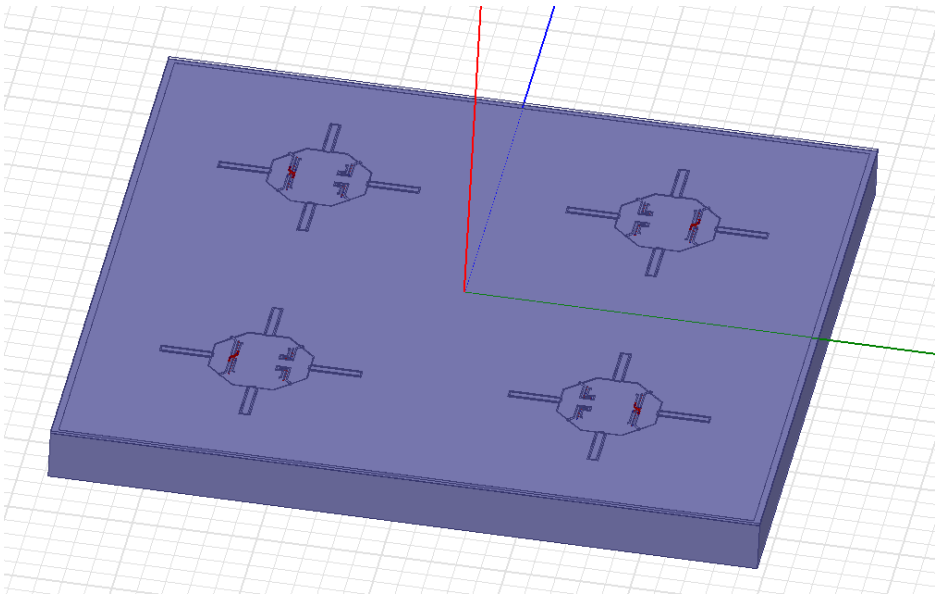

Figure 3.28: 2x2 slot MPD array with tri-level transmission line feeds on all 4 sides of each radiator.

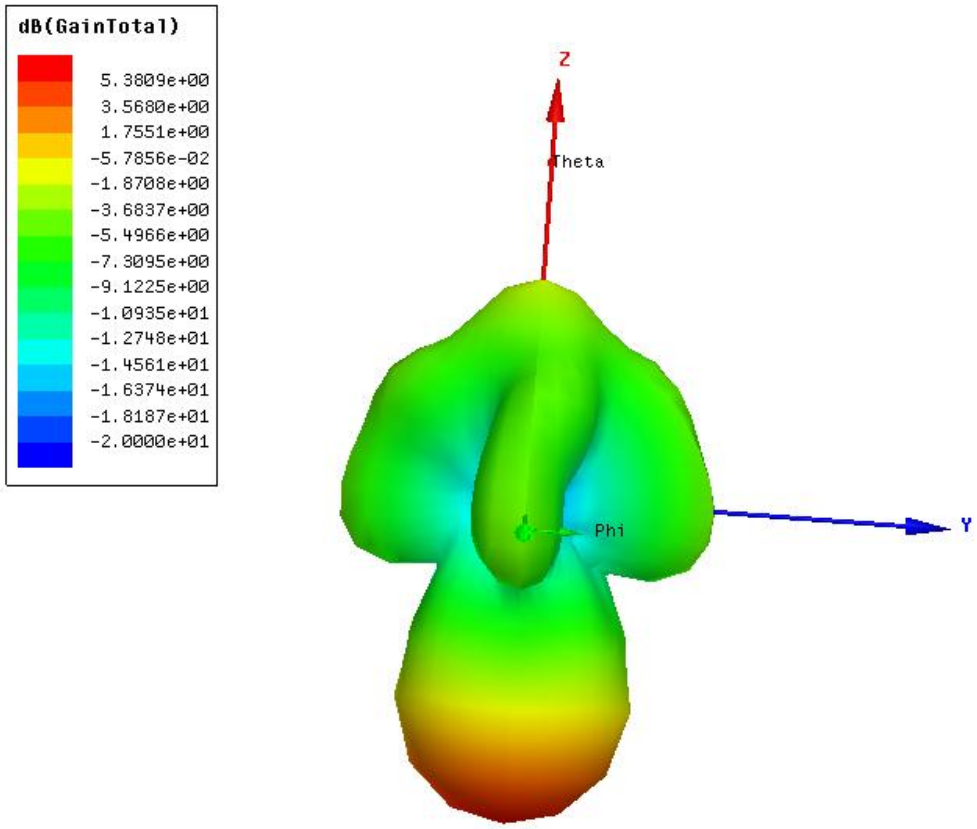

Figure 3.29: Radiation pattern of the 2x2 slot MPD array showing downward radiation with a maximum gain of 5.4 dBi

3.8 Differential Drive Multi-Port Driven Antennas

The radial MPD presented in the previous sections uses a single-ended port to drive a signal ring against a ground spoke. Another approach to designing MPD antennas is using differential drives. One of the benefits of a differential drive MPD radiator is that it can separate the feed line transmission lines from the radiator lines. If the feed line is run perpendicular to the antenna, and differential drives are used, there will be no net current on the feed line. The differential drives can also be used to lower the input impedance of the radiator, something that is important for high power designs in integrated processes. This is because the breakdown voltage is fixed for any given process, so in order to increase the power, the current must be increased, and the load impedance must thus be decreased. The reason that the differential drives can create this low impedance is because they can be used to create virtual short circuits near the drives that can effectively sink the current provided by the drives.

3.8.1 Virtual Shorts to Lower Input Impedance in Multi-Port Driven Radiators

As is shown in Appendix A, the far-field radiation is dependent on the integration of all of the currents on the chip. Currents running in opposite directions cancel, and currents that are in the same direction add constructively and radiate. Line-based single port antennas can either be open at the end, as in the case of a dipole, or can loop back to another part of the antenna, producing a virtual short, as in the case of a loop antenna.

To gain intuition about the lowest input impedance of one of these lines, we use the model of the line antenna as a lossy transmission line, and it becomes apparent that the lowest input impedance for a line that is open on one end is when the line is around one quarter of the wavelength λ, as shown in Figure 3.30. Even in this case with a lossless antenna, the input impedance is still at 73 Ω [1]. On the other hand, if the end of the line is a virtual short like a loop antenna, the low impedance point will be $\lambda/2$ from the short, or when the loop is λ in circumference, yet that is still around 100 Ω [1]. Neither of these types of line antennas will provide an inherently low impedance on the order of 10 - 30 Ω without an additional matching network that will introduce significant amounts of loss to the system.

The final way to achieve a low impedance is by using a very short line that loops back to itself, or a very small loop antenna. Though it achieves a low input impedance, this antenna is also not suitable for integrated radiators because, due to the small radius of the loop, the currents on opposite sides of the loop cancel each other out in the far field, and the antenna becomes very inefficient.

What is needed is a way to keep these virtual shorts close to the drive points, without looping the signal back around to the input port, as shown in Figure 3.30. WIth only a single port, this would not be possible, but by utilizing several ports, this configuration can be achieved.

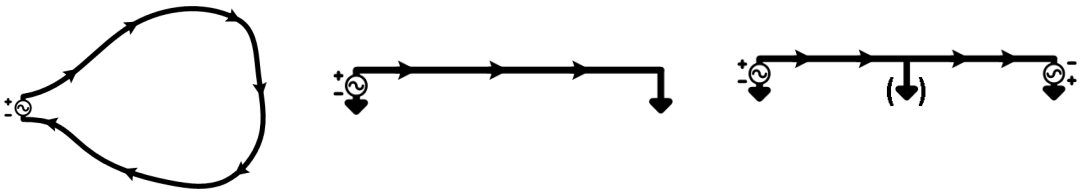

Figure 3.30: Small loops have currents that cancel in the far field (left). What is needed is a current sink that creates a short circuit without needing to loop back (center). The use of multiple drive points with virtual short circuits between them provides this current sink (right)

3.8.2 Linear Differential Multi-Port Driven Radiator

The current pattern from Figure 3.30 can be achieved if it is repeated several times, as shown in Figure 3.31, and using a differential drive from the next set of ports over to create the virtual short between the ports. Clearly at some point this has to stop, but for the moment we will assume a low impedance can be provided at each end, and will look into the operation near the middle of the radiator. At any given instance, the current will always be traveling one way, and thus the antenna will radiate effectively, but at the same time, the drive points are very near the virtual shorts, which means that the input resistance will be low. In fact, the short lines will act like short transmission lines, which means that they will appear slightly inductive, and if spaced correctly, will resonate out the capacitance of the drive transistors.

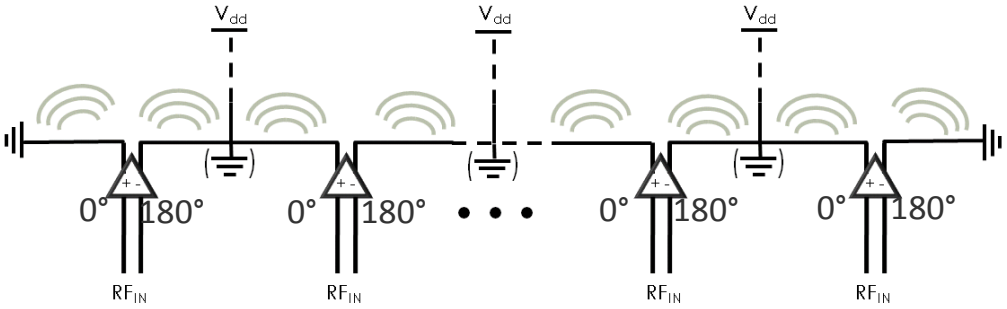

Figure 3.31: Block diagram of the differential linear MPD radiator showing differential drivers that are fed perpendicular to the antenna, and VDD lines that are connected in parallel to the virtual short circuits on the antenna.

In this way, a good impedance match to the drive transistors can be provided directly from the antenna, without any impedance matching networks in between. The overall current pattern will look similar to a dipole, but rather than the current decaying as a sinusoid to nothing at the end of the line, the current will instead remain relatively constant across the entire length of the radiator, as shown in Figure 3.32. In order to not cause any mirror currents that could harm radiation, the feed lines to the radiator are run perpendicular

to the radiator lines, and because the current is differential, there is no net current that radiates from the feed lines. The virtual short circuits that appear between each of the lines provides an ideal place to provide VDD, as the impedance looking into the VDD line will always be in parallel with the virtual short, and thus will not be seen by the radiating structure. The VDD can then be run from these points perpendicular out on the other side of the radiator, again to minimize any mirror currents that would negate radiation.

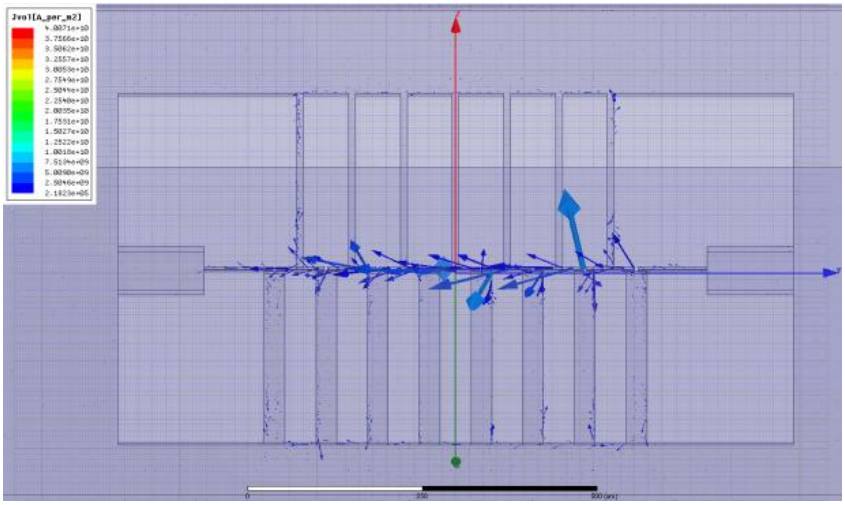

Figure 3.32: Simulated current distribution of the differential linear MPD radiator showing instantaneous current all pointed in the same direction, leading to effective radiation.

Finally the ends of the radiator need to be dealt with. The final drives must drive something without the benefit of the virtual shorts, so very wide lines are used that can distribute the current out and provide something of a match from an open circuit at the end to a low impedance at the drive point are used, as shown in Figure 3.32. The mismatch of having these higher impedance nodes at each end affects the other side of the ports on each end, but by the second set of lines, the variation in currents becomes negligible to the design.

This radiator was designed to operate at 94 GHz in a 65 nm bulk CMOS process. It is driven by a fundamental frequency voltage controlled oscillator (VCO), followed by three stages of amplification, with a 2-1 splitter between the first and second stage, and 4-1 splitters between the 2nd and 3rd stages, as shown in Figure 3.33.

Each amplifier is a cascode amplifier, and transmission line matching networks are implemented between each amplification stage. The output drivers feed the radiator directly and immediately, with no additional matching used. The saturated power at 94 GHz that is input into the antenna is simulated to be 20.4 dBm. The simulated small signal gain of the three amplifying stages is 25.5 dB, and the maximum antenna gain is simulated to be 4.5 dBi, with the simulated antenna pattern shown in Figure 3.34. This yields a simulated

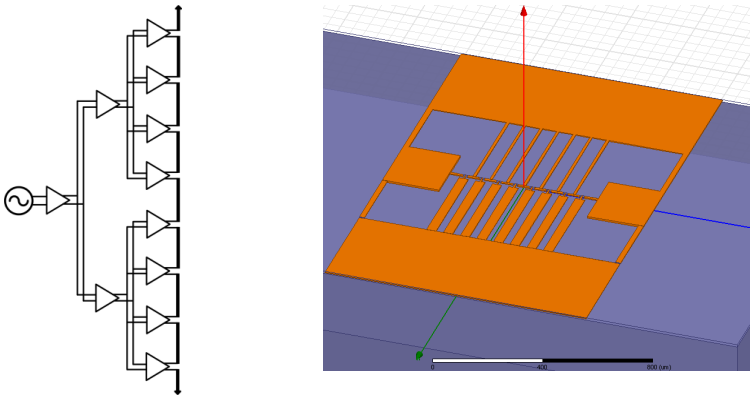

Figure 3.33: Block diagram of the differential linear MPD radiator showing on chip oscillator, and division and amplification stages is shown on the left. A 3D depiction of the metal structures of the radiator is shown on the right.

EIRP of 24.9 dBm, with a radiation efficiency of 30%, when mounted on a ground plane and with a substrate that is $\lambda/4$. While this design was never fabricated, fully finished layout of the radiator is shown in Figure 3.35.

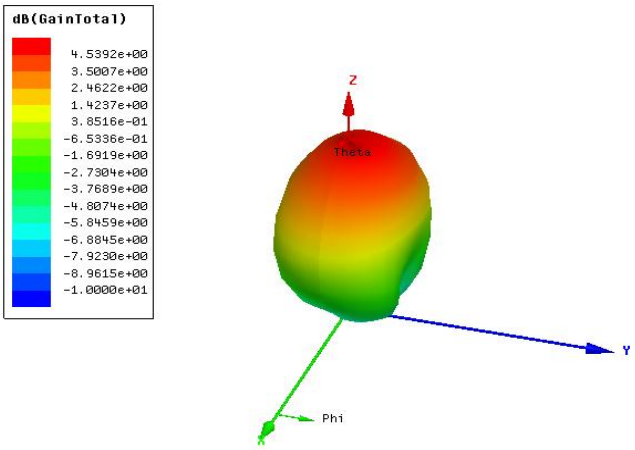

Figure 3.34: Simulated antenna pattern of the differential linear MPD radiator shows a maximum gain of 4.5 dBi.

3.8.3 Passive 2 GHz Differential Multi-Port Driven Test Structure on FR4 PCB

A passive version of the antenna was implemented by Gunnar Atli-Siggurdson, an undergraduate researcher whom I co-advised for a summer. It has the same basic 8 differential drive structure, but was scaled by

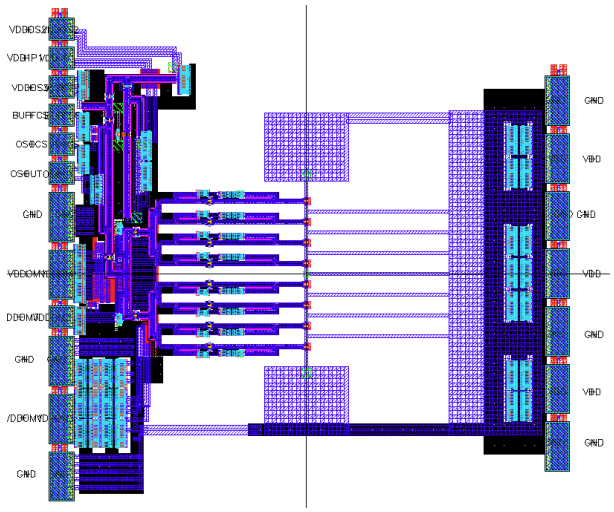

Figure 3.35: Layout of the differential linear MPD radiator in a 65nm bulk CMOS process.

around a factor of 50 to enable radiation at 2 GHz. Differences in the dielectric constants of the silicon substrate instead of the FR4 necessitated a slight adjustment to this scaling. A photo of this passive structure is shown in Figure 3.36, and it has a simulated maximum gain at 2 GHz of 5 dBi, with the antenna pattern shown in Figure 3.37.

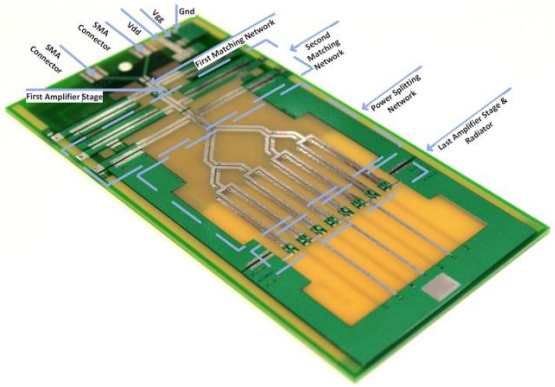

Figure 3.36: Photo of the fabricated passive differential linear MPD antenna scaled to operate at 2 GHz

After fabrication, measurements were taken of the antenna pattern, and match closely with simulations, as shown by two slices of the antenna pattern in the elevation and azimuth plane depicted in Figure 3.38. These results show the viability of this structure as a multi-port driven antenna, and it is anticipated that the 94 GHz integrated radiator version of the design would operate similar to the predicted simulations.

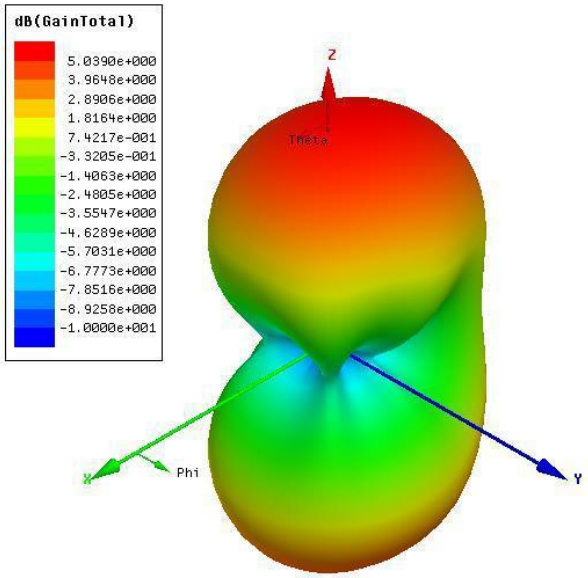

Figure 3.37: Simulated antenna pattern of the PCB differential linear MPD antenna shows a maximum gain of 5 dBi.

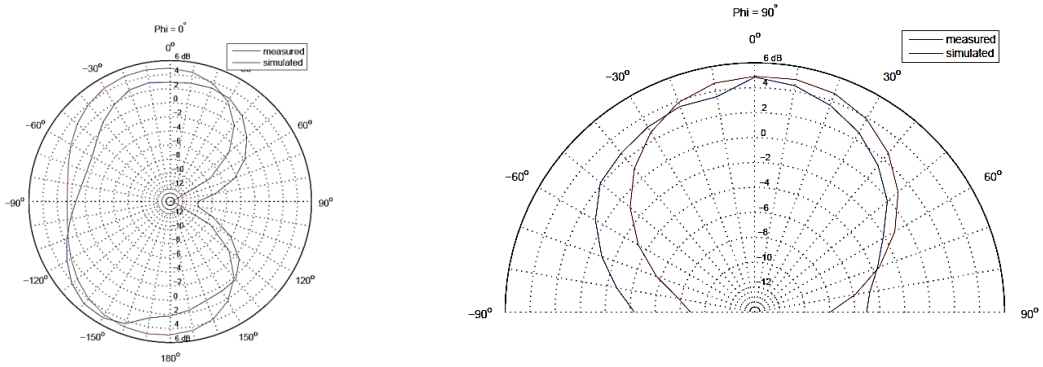

Figure 3.38: Simulated and measured antenna patterns in the elevation (left) and azimuth (right) planes match closely with electromagnetic simulations.

3.9 Conclusions

The MPD antenna design methodology of using multiple input ports on an integrated antenna to eliminate lossy power transfer mechanisms and output matching networks has been described. An analysis of a radial MPD antenna in free space with supporting electromagnetic simulations was presented that gave intuition as to the correct sizing of the antennas as well as showed the wide bandwidth that these antennas can support. The measurements of the chip match well in power as well as frequency to the simulations of the design, and a comparison to other state of the art designs show what we believe to be the highest output power as well as EIRP per element of any integrated silicon radiator above 110 GHz without external optics. Using electromagnetic dualities, a slot-based radial MPD radiator is also presented to increase the gain and efficiency of the antenna as well as decrease the required area footprint for the radial MPD antenna. Differential drive MPD radiators were also presented, along with the design and simulation of a linear differential MPD radiator at 94 GHz. A proof of concept version of the antenna at 2 GHz was fabricated and measured to closely match simulations.

Chapter 4

Dynamic Polarization Control and Modulation

Efficient wireless connectivity inherently relies on maximum coupling of electromagnetic energy from a transmitter to a receiver over the propagation medium. One of the important parameters determining the strength of this coupling is the polarization of the transmitted electromagnetic waves with respect to the receiver antenna, where a polarization mismatch can severely degrade the coupling.

Solutions have been devised to enforce polarization matching in any given transmitter/receiver pair in advance, such as using linearly polarized antennas with the same orientation (vertical or horizontal). However, a fixed linearly polarized link has to rely on accurate alignment not suitable for mobile applications, necessitating some sort of antenna diversity with two antennas having orthogonal polarizations. Alternatively, an antenna radiating with circular polarization can be utilized to couple to a receiving antenna with the same handedness to reduce the orientation dependence of the polarization matching. Unfortunately, most circularly polarized antennas have narrow polarization bandwidths and beamwidths, where they rapidly lose their circular polarization away from the center frequency and primary axis of propagation. More importantly, they can only couple to other circularly polarized antennas of the same handedness without any significant loss of energy. This leads to signal degradation if they were to couple to antennas with other polarizations or orientations. Dual polarization antennas (e.g., [77]) intended for higher data throughput have not addressed these polarization matching issues either. Generally, in most of these solutions the choice of polarization modality has to be made, a priori, limiting the reconfigurability and performance of the systems.

Dynamic polarization control (DPC) is a way to help ensure polarization matching, even as the antennas move or are rotated, and can even be used as a means to transfer additional data or to encode data in the polarization itself.

The Friis equation (4.1) can be used, along with the gains of the antennas (G_t, G_t), the wavelength (λ), and the distance between the antennas (R), to calculate the power transfer between two antennas

$$P_r = G_t G_r \left(\frac{\lambda}{4\pi R}\right)^2 P_t \tag{4.1}$$

.

One critical requirement for this equation to be valid is that the antennas must be polarization matched. Polarization mismatch can result in up to a complete loss in transferred power. The goal of DPC is to maintain maximum coupling between the transmit and receive antennas regardless of the receiver antenna's orientation or polarization. This can can be done through several antennas working in unison, or though a mult-port driven (MPD) antenna configured for DPC.

First, a brief background on electromagnetic polarization will be discussed, followed by a look at how polarization matching affects power transfer in antenna systems. Next, DPC methods with multiple antennas are examined followed by DPC through a single MPD antenna. The design and measurement of an 2x1 DPC radiator is then presented followed by a discussion of polarization modulation, and concluding remarks.

4.1 Electromagnetic Polarization

Electromagnetic waves traveling in a free-space, linear, isotropic medium eventually reach a region known as the far field, where only transverse electromagnetic (TEM) waves remain [1]. For these TEM waves, the instantaneous electric and magnetic field are perpendicular to each other as well as to the direction of propagation. For the rest of this chapter, it will be assumed that the direction of propagation is along the Z axis and the electric and magnetic fields in the far field are contained within the XY plane.

The relationship in amplitude between the electric and magnetic fields is determined by the free-space impedance of the medium (η), which can be calculated [1] from the dielectric constant (ϵ) and the magnetic permeability (μ) as

$$\eta = \sqrt{\frac{\mu}{\epsilon}}. \tag{4.2}$$

The polarization of these waves can be defined as the trajectory of the instantaneous electric field for a given XY plane over one period. This trajectory results in an elliptical pattern. A polarization diagram that traces this elliptical trajectory can be written, and the polarization can be expressed in several ways that all define that ellipse, and the direction of rotation. In a fundamental way, it can be written as the spacial angle of the major axis of the ellipse with respect to the coordinate system, and the ratio between the major and minor axes (the axial ratio), and the direction of rotation, left-hand circularly polarized (LHCP) or right-hand circularly polarized (RHCP). Large axial ratios result in polarizations close to linear, and axial ratios that approach unity are close to circular. Some examples of linear, elliptical and circular polarizations are shown in Figure 4.1.

The polarization can also be written as a sum of two predefined independent polarizations, where the polarization is given by the ratio between amplitudes of the two polarizations, and the phase difference between them. Two special cases of this include using the linear X and Y polarizations as a basis, and

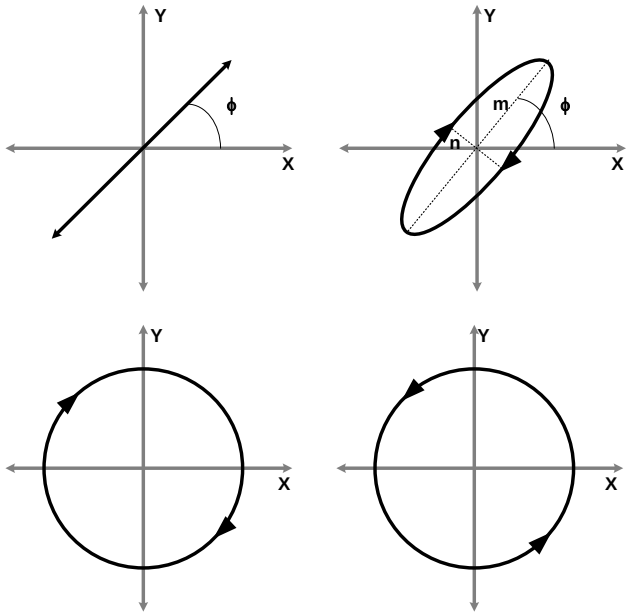

Figure 4.1: Examples of electromagentic polarization, with linear polarization having polarization angle ϕ (upper left), elliptical polarization having polarization angle ϕ and axial ratio m/n (upper right), left-hand circular polarization (lower left) and right-hand circular polarization (lower right).

using left-hand and right-hand circularly polarized fields as a basis. Using circular polarizations as a basis another polarization ratio, the circular polarization ratio, can be defined by looking at the LHCP component of the field divided by the RHCP component of the field. In this case, circular polarization ratios near zero or much greater than unity result in fields close to circular polarization, while circular polarization ratios near unity result in fields that are close to linear polarization. For the purposes of this chapter, we will report polarization as the spatial angle of the major axis of polarization from the positive X axis, and the axial ratio of the polarization.

For two antennas to be polarization matched, the polarization of the fields produced by the transmit antenna in the far field must match the polarization of the fields that can be received by the receive antenna, and vice versa when the transmit and receive antennas are switched. This way there is no loss in the transmission due to polarization mismatch.

For example, two dipoles that are oriented in the same direction, as shown in Figure 4.2a are polarization matched, as they each radiate linearly in the X dimension. On the other hand, the dipoles in Figure 4.2b are oriented in orthogonal directions, with antenna A aligned with the X axis, and antenna B aligned with the Y axis. This means that the polarizations are completely mismatched. The electric field in the far field radiated by antenna A is linearly polarized in the X direction, but the field that antenna B can receive is oriented in the Y direction, so antenna B will not pick up any of antenna A's radiation. Of course these systems are all reciprocal, and if antenna B is being driven, it produces an electric field in the far with linear Y polarization,

while antenna A can only pick up radiation in the X direction. These two examples show the extremes of polarization matching, where the power transfer is either at a maximum or at zero, but there is also a full range of partially polarization matched antennas where only some of the power predicted by the Friis equation is actually transferred.

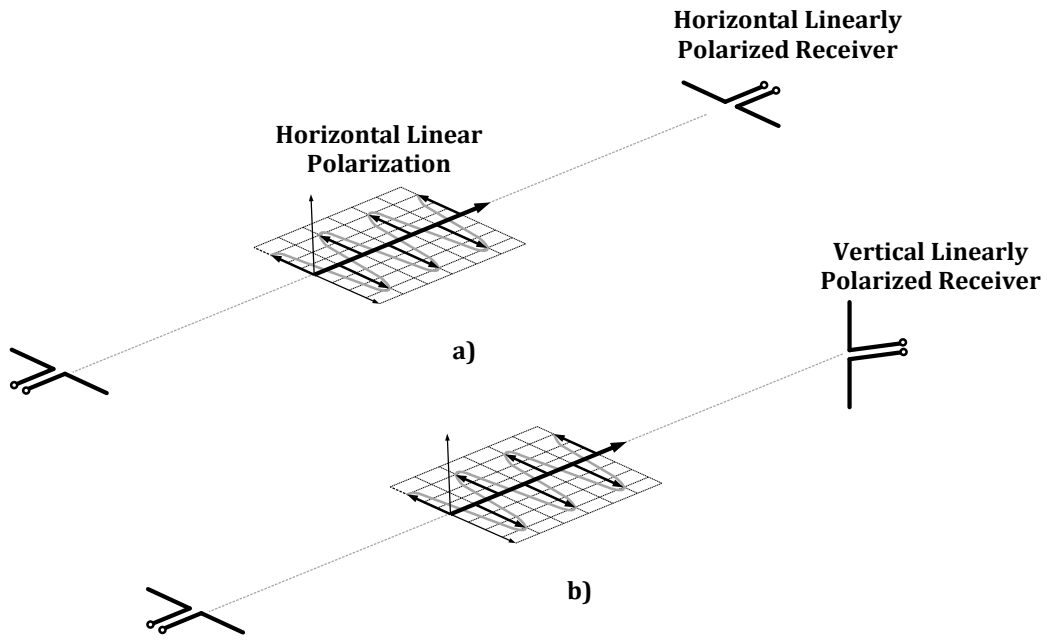

Figure 4.2: 2 dipoles in same orientation that are polarization matched and thus have maximum coupling (a), and 2 dipoles that are orthogonal and thus no signal will be picked up by the receive antenna (b).

In linear polarizations, rotation around the axis of propagation of the transmit or receive antenna can cause this polarization mismatch, while in circularly polarized systems, polarization mismatch can often occur due to rotation around one of the other two axes, causing the circular polarization to become elliptical. One benefit of circular polarization is that rotation about the axis of propagation does not cause any polarization mismatch, but rather will result in a change in phase of the received signal.

Two polarizations are said to be orthogonal if the projection of a field in one polarization onto the other polarization results in a field with zero amplitude. In this way, linear polarizations oriented 90° from each other are orthogonal, as are LHCP and RHCP fields. Because they are orthogonal, they can act independently without changing the fields in the other polarization. Also, because the overall polarization is a trajectory in two dimensions, having two orthogonal polarizations creates a basis for that space, and means that for any polarization there is always exactly one polarization that is orthogonal to it.

4.2 Dynamic Polarization Control

Dynamic polarization control is a method where the radiating structure can adjust its polarization dynamically, and entirely electronically to match the polarization of any receive antenna regardless of its polarization or orientation, as shown in Figure 4.3.

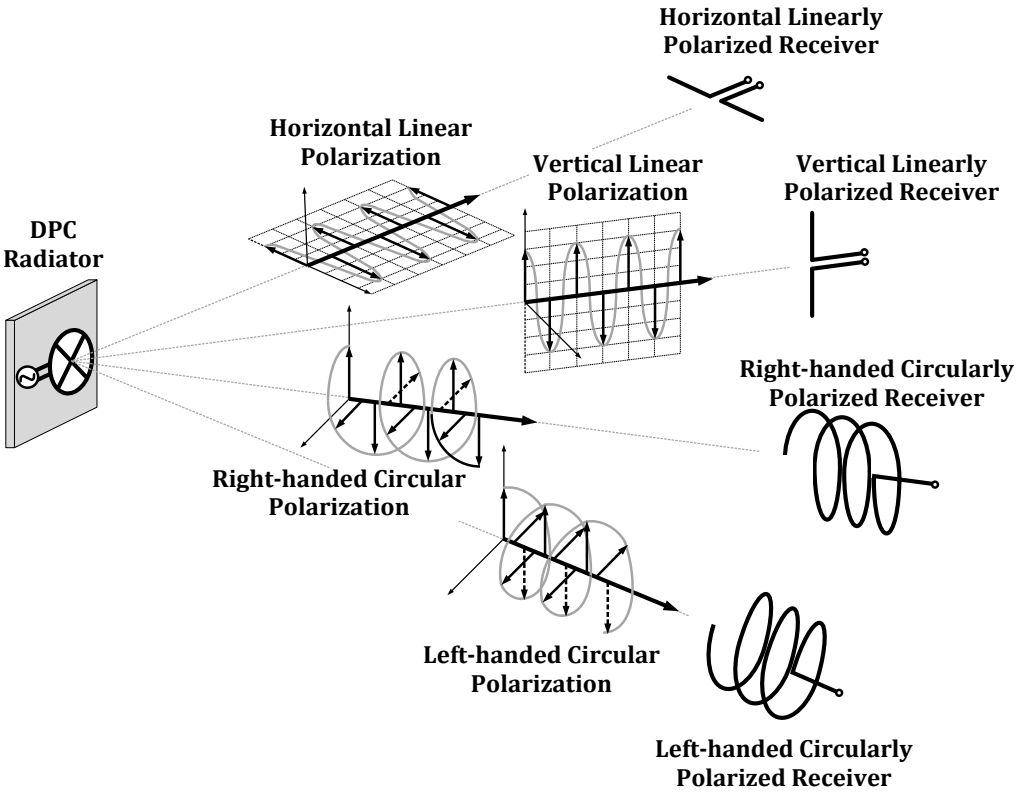

Figure 4.3: Dynamic Polarization Control (DPC) radiator can switch between various electromagnetic polarizations to match different receivers regardless of their polarization or spacial orientation.

4.2.1 Dynamic Polarization Control with Multiple Antennas

DPC using multiple antennas uses the transmission of two or more antennas with independent polarizations to create a third polarization that can change simply due to the input drives of the two transmitting antennas. One way to achieve DPC is to produce a linear polarization that can be dynamically controlled through the drives of two circularly polarized antennas, one RHCP and one LHCP. If they are equal in amplitude, then the resulting field is going to be linearly polarized, with a polarization angle related to half of the difference in phase between the two circularly polarized fields as shown below.

Let the electric fields of the radiation produced by each circularly polarized antenna with the same ampli-

tude $\mathbf{E_1}$ and $\mathbf{E_2}$ be defined as

$$\mathbf{E_1} = A[\sin(\omega t + \phi_1)\hat{\mathbf{u}}_\mathbf{x} + \sin(\omega t + \phi_1 + \pi/2)\hat{\mathbf{u}}_\mathbf{y}], \tag{4.3}$$

and

$$\mathbf{E_2} = A[\sin(\omega t + \phi_2)\hat{\mathbf{u}}_\mathbf{x} + \sin(\omega t + \phi_2 - \pi/2)\hat{\mathbf{u}}_\mathbf{y}]. \tag{4.4}$$

These fields have a similar amplitude (A), but have independent phases (ϕ_1 and ϕ_2), and are both circularly polarized in opposite directions. The combination in the far field is the sum of the two electric fields:

$$\mathbf{E_{total}} = \mathbf{E_1}+\mathbf{E_2} = A[\sin(\omega t+\phi_1)\hat{\mathbf{u}}_\mathbf{x}+\sin(\omega t+\phi_1+\pi/2)\hat{\mathbf{u}}_\mathbf{y}]+A[\sin(\omega t+\phi_2)\hat{\mathbf{u}}_\mathbf{x}+\sin(\omega t+\phi_2-\pi/2)\hat{\mathbf{u}}_\mathbf{y}]. \tag{4.5}$$

This equation can be simplified by letting $\Delta\phi = \phi_1 - \phi_2$ and $\Sigma\phi = \phi_1 + \phi_2$, resulting in

$$\mathbf{E_{total}} = 2A[\cos(-\Delta\phi/2)\sin(\omega t + \Sigma\phi/2)\hat{\mathbf{u}}_\mathbf{x} + \sin(-\Delta\phi/2)\sin(\omega t + \Sigma\phi/2)\hat{\mathbf{u}}_\mathbf{y}] \tag{4.6}$$

or

$$\mathbf{E_{total}} = 2A\sin(\omega t + \Sigma\phi/2)[\cos(-\Delta\phi/2)\hat{\mathbf{u}}_\mathbf{x} + \sin(-\Delta\phi/2)\hat{\mathbf{u}}_\mathbf{y}]. \tag{4.7}$$

Thus, $\mathbf{E_{total}}$ in the X and Y dimensions are always in phase regardless of ϕ_1 or ϕ_2, and linear polarization is always created. The magnitude of the field,

$$mag(\mathbf{E_{total}}) = 2A, \tag{4.8}$$

which is also independent of ϕ_1 and ϕ_2. The angle in space of the field,

$$angle(\mathbf{E_{total}}) = \Delta\phi/2 - \pi/2, \tag{4.9}$$

which shows that the angle of the polarization is dependent only on the difference between the two phases of the original fields. Finally, the phase of the field,

$$phase(\mathbf{E_{total}}) = \Sigma\phi/2, \tag{4.10}$$

means that the phase of the electric field can be controlled independently from the polarization angle or the amplitude of the field.

These results suggest that as long as the amplitudes of the RHCP and LHCP fields are equal, the resulting field will always be linearly polarized, as can be seen in Figure 4.4. The common mode of the phases of the fields will determine the phase of the output field. The differential mode of the fields will determine the

polarization angle. The common mode of the amplitudes will determine the amplitude of the output field, and if the amplitudes of the input fields are different from each other, the differential mode of the amplitudes will determine the axial ratio of the polarization.

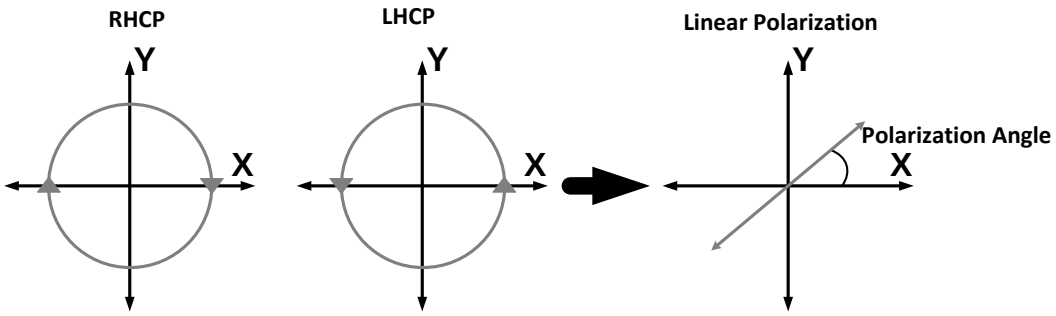

Figure 4.4: The superposition of two circularly polarized fields in the far field with opposite handedness (LHCP and RHCP) results in a linearly polarized field, whose polarization angle is determined by the phase difference of the original two circularly polarized fields.

In this way, if the desired output polarization is linear, but the polarization needs to be changed dynamically, as in the case of mobile antennas, full non-constant modulating signals can be sent through the common modes of the amplitude and phase, and polarization matching can be maintained by adjusting the differential mode of the phases.

Another way to accomplish this is to use linearly polarized transmit antennas (Figure 4.5), but in order to achieve a full 180° range of polarization angles, the amplitudes of each antenna must be adjusted even when the desired amplitude of the output field remains constant. Using two circularly polarized fields reduces the requirements for the drivers of the antennas by keeping the phase and amplitude of the output field tied to the common mode of the input phase and amplitude and only requiring a phase shift between the antennas to produce the change in polarization angle.

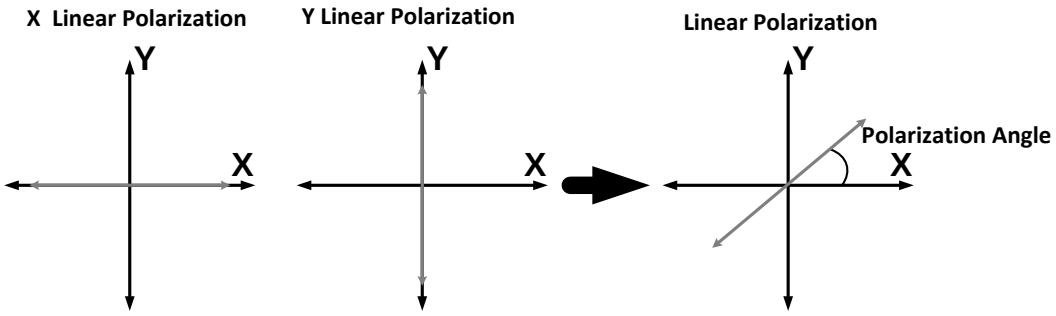

Figure 4.5: The superposition of two linearly polarized fields that are in phase will also result in a linearly polarized field, where the polarization angle is determined by the relative amplitudes of the original two linearly polarized fields.

4.2.2 Dynamic Polarization Control Through a Single Multi-Port Driven Radiator

A single port antenna in a linear medium has a fixed polarization at a given frequency as long as it is not rotated or otherwise manipulated mechanically. With only a single port at a given frequency, amplitude and phase of the incoming signal are the only controls available. The amplitude cannot change the polarization of an antenna in a linear material or it would violate superposition. A passive antenna in a linear medium is also time invariant, which means that phase changes also do not affect the operation or polarization of a single port antenna.

This fundamental limitation is lifted once more than one port is introduced. In the previous section, the second port was on a second antenna, but by using an MPD antenna, introduced in Chapter 3, additional ports can be utilized within the same antenna. The multiple ports enable the designer to set relative phases and amplitudes at different points on the antenna, which can lead to different current distributions, and thus can affect the radiated electromagnetic field and its polarization. The analysis of the MPD from Chapter 3 shows that a single spoke driven differentially produces a linear polarization aligned along the spoke. By combining these polarizations at different amplitudes and phases, the polarization of the resulting electromagnetic field can be controlled, as will be discussed in the following section.

4.3 Proof of Concept of a 2x1 Multi-Port Driven Radiator Array with Dynamic Polarization Control[1]

To verify the viability of the DPC concept, a 110 GHz 2x1 MPD radiator was designed, fabricated and tested using the IBM 32nm SOI CMOS process.

The radial DPC antenna is made up of a signal ring driven at four points against the ends of two ground spokes that are orthogonal to each other (Chapter 3). As illustrated in Figure 4.6, the operation of this antenna can be understood using superposition and analyzing two differential ports of the antenna at a time. The two opposite ports of either spoke are always driven differentially and thus virtual short circuits will appear on the ring and spokes along the axis of symmetry. This means that for this special case with two ground spokes, the driver output impedance at the other two ports is virtually shorted, and thus can be neglected. It will also provide additional isolation between the two pairs of ports. In this configuration standing current waves appear on both the spokes and the ring. The currents on these lines terminated with virtual shorts are similar to a transmission line terminated with a short circuit. If the circumference of the loop is greater than one wavelength, but less than $\pi/2$ wavelengths, a node will appear on the current wave on the loop, but not on the spoke, and the large currents at the virtual shorts on both the spokes and ring will be in phase, resulting in effective radiation (Figure 4.6).

[1]This design was a joint project with my colleague Amirreza Safaripour. I was responsible for the design of the MPD antenna, as well as the array design, and the measurements. Amirreza designed the oscillators, amplifiers, and phase shifters, as well as the locking network.

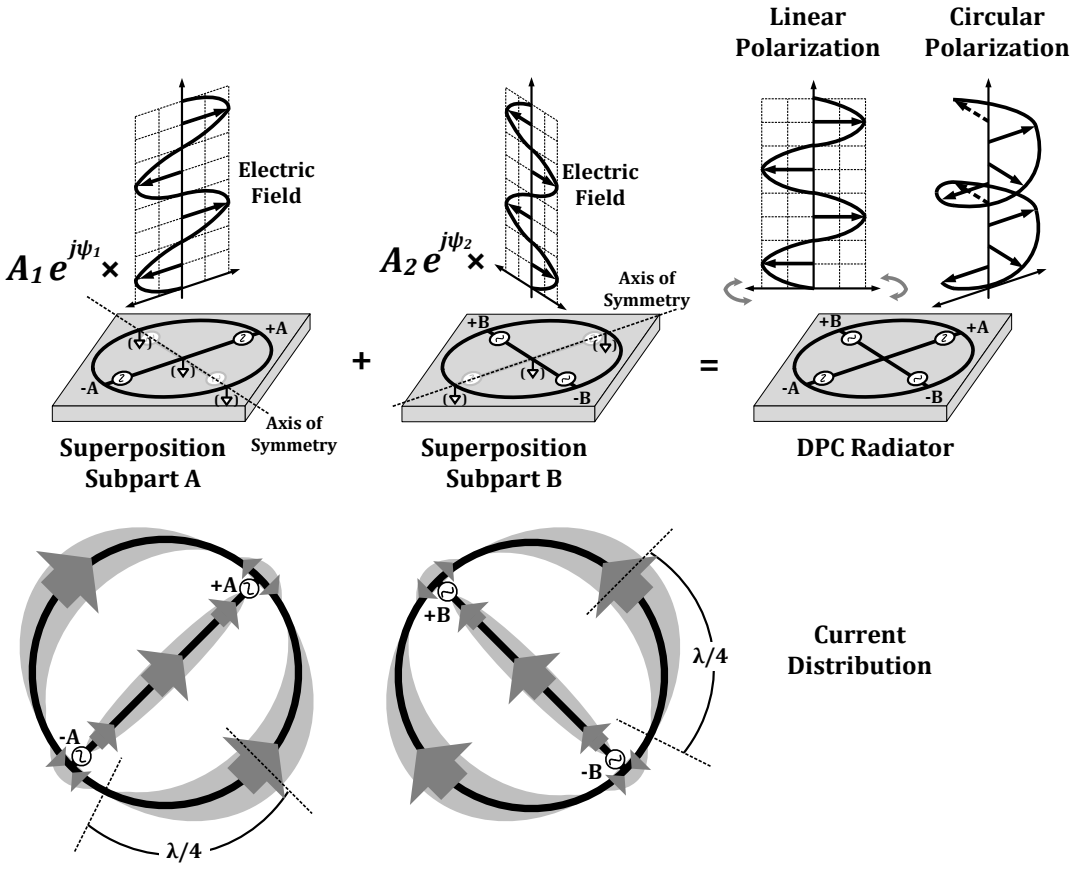

Figure 4.6: Breaking up the MPD radiator into two using superposition shows how it can be considered as two linearly polarized radiating elements that are overlapping in space and which are somewhat isolated from each other due to being driven at the axis of symmetry of the other set of drives. This enables the DPC radiator to change its polarization by changing the amplitude and phase of the two radiating subparts. Below each subpart is a depiction of the two current distributions for each subpart showing the nodes on the rings that cause the currents to line up in the same direction and create effective radiation.

These standing current waves will produce radiation that is linearly polarized along the axis of the spoke being driven. Driving the two spokes in phase or differentially produces linear polarization along the $\phi = 0°$ or $90°$ planes, respectively. If each spoke is driven individually, linear polarization is produced along the $\phi = 45°$ or $135°$ planes, instead. Linear combinations of these differential drives are dynamically set in the DPC radiator to enable linear polarization at any angle. If the radiator phase setting is chosen such that the ports are driven in quadrature, a traveling wave around the ring is created thus radiating with circular polarization. Phase differences between in phase and quadrature produce elliptical polarization, and thus the DPC radiator can transmit with any polarization .

The 2x1 DPC radiator array was designed in a 32nm CMOS SOI process. Other than being a 2-spoke design, rather than an 4-spoke design, and being larger due to the lower desired frequency of this design, the MPD antennas are designed similarly to those described in Chapter 3. It is mounted on a ground plane and radiates upward without any external lenses, dielectric superstrates or other additional post processing of any kind. Each radial radiator has an on-chip upper ground plane pulled back from the ring $\lambda/4$ to provide high impedance at the drive ports and is locked to a central locking network. The locking network is made up of a 105.5 GHz quadrature voltage controlled oscillator (QVCO) followed by three amplification stages that then feed into phase rotators within the core of each radiator (Figure 4.7). The QVCO is made up of two cross-coupled oscillators with quadrature coupling transistors that feed directly into a buffer amplifying stage, as seen in Figure 4.8. The second and third amplifiers are tuned cascode amplifiers. The radiator core has a similar set of oscillators followed by three tuned amplifiers with the addition of phase controllers that inject currents from the locking network into the tanks of the oscillators. The quadrature phase rotators are based on Gilbert cells and lock the radiator core. Each oscillator has varactor frequency tuning control, phase control though the locking network, and amplitude control by adjusting the bias voltages of the amplifying stages.

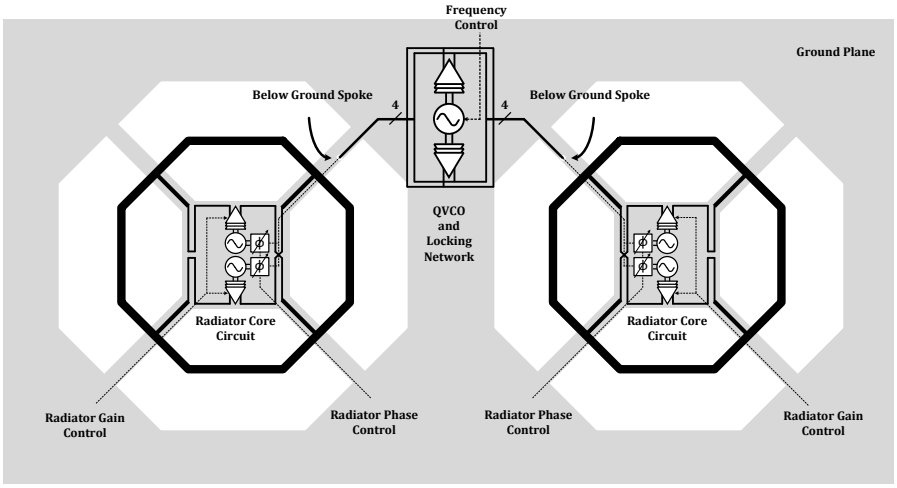

Figure 4.7: Block diagram of 2x1 DPC radiator array showing individual radiators, the locking network, and control of frequency, phase and gain to enable dynamic polarization control.

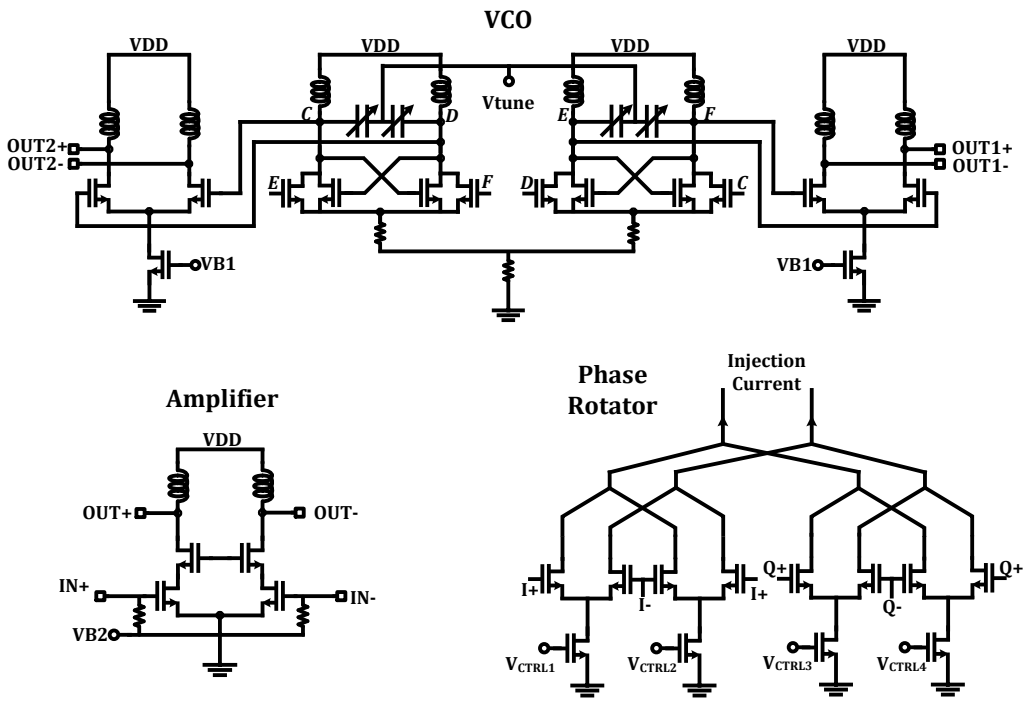

Figure 4.8: Schematics of Circuits used in the 2x1 DPC radiator array, with the VCO with first amplifier stage on top, the cascode amplifier used for the second two stages in the lower left, and the phase rotators in the lower right.

The 2-spoke design of the MPD antennas enables a more simple circuitry than the 4-spoke design discussed previously, as creating a quadrature oscillator is more straightforward than creating an 8-phase oscillator. The MPD oscillators are capable of producing -20.5 dBm of power for each phase in simulation, which after the amplification stages, produce 0 dBm going into each of the 4 ports of each MPD antenna. There is amplitude control of the amplifying stages through changes in DC current through the tail transistors that can be adjusted to adjust the power entering the ports of the MPD antenna.

Again the differential signals are kept together through the amplifier stages to enable the use of virtual shorts, and then routed after the last amplifier stage to the appropriate feeds to create a traveling wave on the ring. This means that two of the lines must be crossed to create a +A, +B, -A, -B arrangement from the +A, -A, +B, -B that comes directly from the amplifying stages. In order to match the phases at the ports of the MPD antenna, the transmission lines on the two lines that are not crossed are meandered down toward the center of the core, as shown in Figure 4.7.

Full 3-dimensional electromagnetic simulations of the radiators using HFSS show a gain of 1.37 dBi with linear polarization when the array is driven in a linear polarization mode, and the simulated pattern is shown in Figure 4.9, and 0.76 dBi gain when the array is in a circular polarization mode with both radiators radiating in the same direction, with the simulated antenna pattern shown in Figure 4.10. Linear polarization can also be achieved when the radiators are in opposite circular polarization modes (one is right-handed, one

is left-handed), and the maximum gain is simulated as 0.76 dBi, with a radiation pattern shown in Figure 4.11.

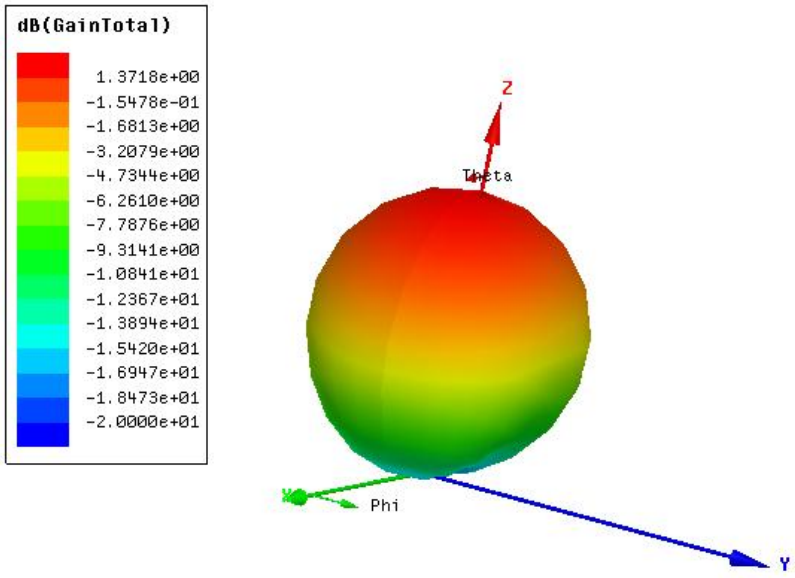

Figure 4.9: Simulated antenna pattern of the 2x1 radiator array in linear polarization mode showing a maximum 1.37 dBi gain.

4.3.1 Measurements of the 2x1 Multi-Port Driven Radiator Array

The fabricated chip is mounted onto a printed circuit board (PCB) and attached to a two-dimensional stepper motor for antenna pattern measurements. The signal is received by a 22.0 dBi gain horn antenna placed 12 cm away. The horn antenna feeds an 8th harmonic mixer that goes to an IF amplifier and to a spectrum analyzer, and is calibrated with a PM4 Erickson power meter. The entire measurement setup is shown in Figure 4.12. The radiator draws 360 mA from a 1.3 V supply and occupies 1.2 x 2.2 mm^2. Reliable polarization control is observed at all measured angles. Broadside as well as off-axis measurements of two examples of receiver locations of 20° off-axis and $\phi = 0$, as well as 30° off-axis and $\phi = 90°$ are demonstrated, with one-dimensional beam steering being utilized when $\phi = 0$.

First, a linear polarization with a high axial ratio is set, and the polarization angle is swept. The full range of polarization angles from 0° to 180° is observed in all three cases with a maximum difference in radiated power for the different polarizations of less than 2.4 dB (Figure 4.13). This full range of polarization angles means that the radiator can produce a field that is polarization matched with any linearly polarized target receiver.

Next, to show control of the axial ratio, the polarization angle was held constant and the axial ratio was

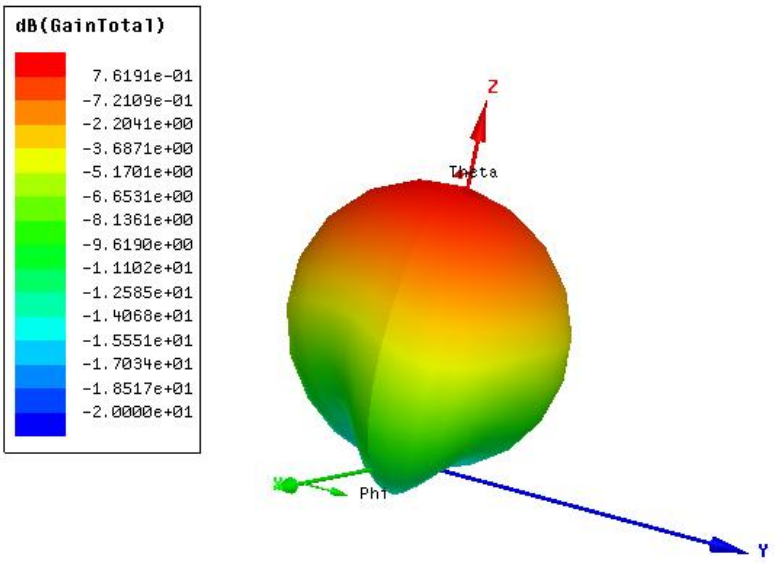

Figure 4.10: Simulated antenna pattern of the 2x1 radiator array in circular polarization mode showing a maximum 0.76 dBi gain.

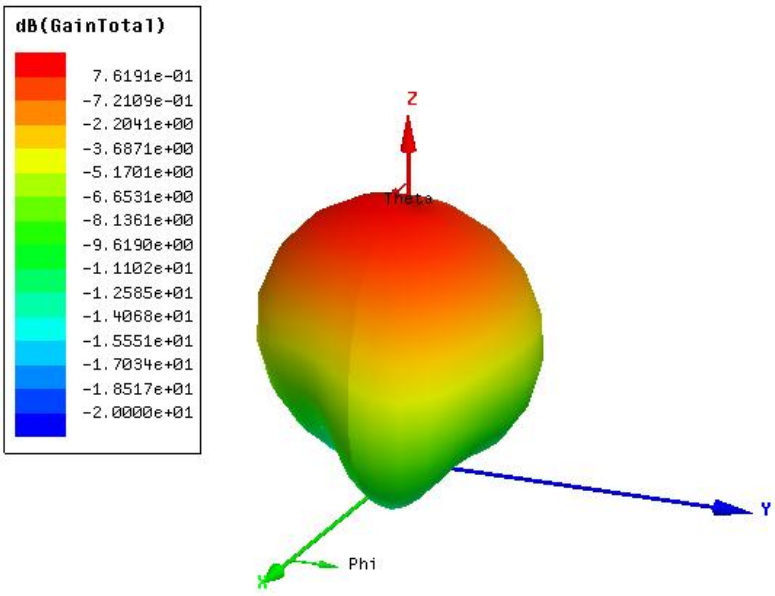

Figure 4.11: Simulated antenna pattern of the 2x1 radiator array with each radiator in an opposite circular polarization mode to produce a linear polarization, and again showing a maximum 0.76 dBi gain.

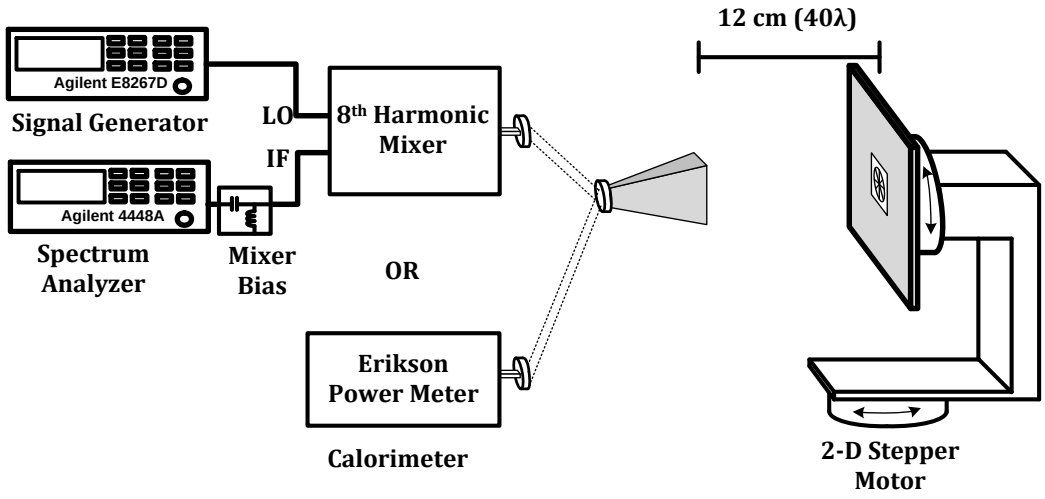

Figure 4.12: Measurement setup for the 2x1 DPC radiator array.

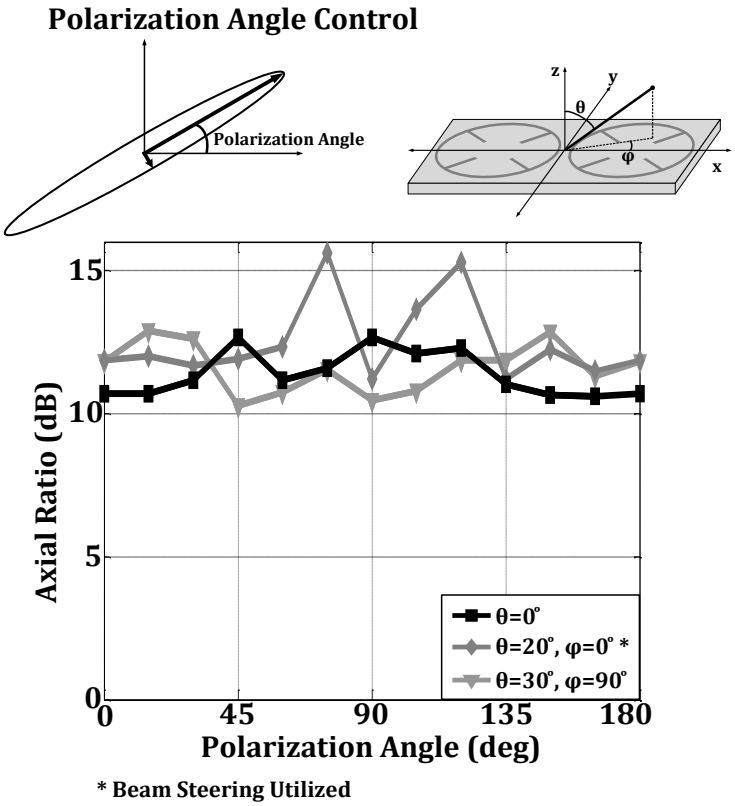

Figure 4.13: Measurements showing a full tuning range from 0° through 180° while maintaining axial ratios above 10 dB the entire time for 3 directions of radiation: broadside, with the beam aimed broadside, 20° off axis in the axis of the array, with the beam aimed toward the direction of the receiver, and 30° off axis perpendicular to the the axis of the array, with the beam aimed broadside.

swept (Figure 4.14). For the broadside case, axial ratios from 2.6 dB through 14 dB are achieved. The measurement range is limited by the noise floor of the harmonic mixer.

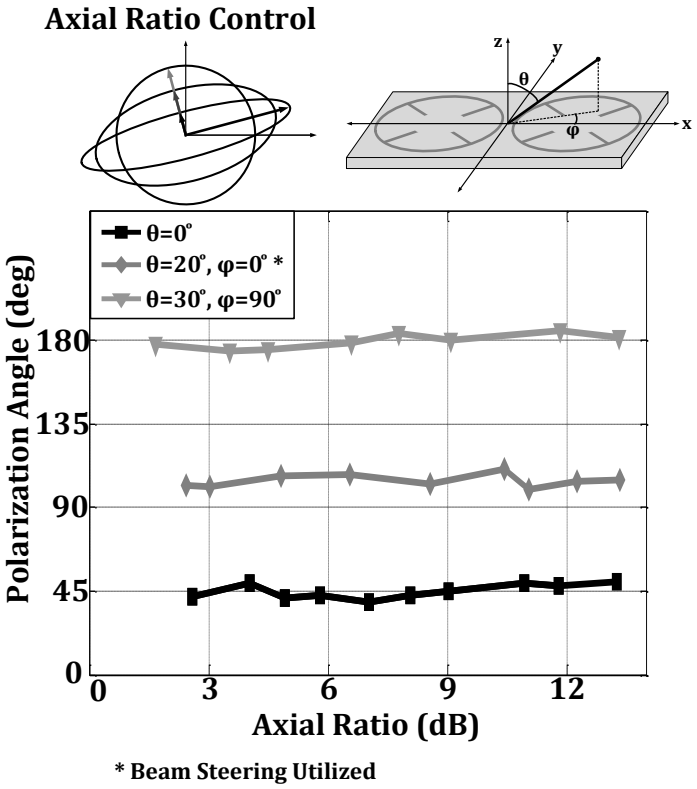

Figure 4.14: Measurements showing control over the axial ratio while maintaining a fixed polarization angle for 3 directions of radiation: broadside, with the beam aimed broadside, 20° off axis in the axis of the array, with the beam aimed toward the direction of the receiver, and 30° off axis perpendicular to the the axis of the array, with the beam aimed broadside.

The measured maximum effective isotropic radiated power (EIRP) is +7.8 dBm at 105.55 GHz producing a total radiated power of 0.9 mW. The maximum EIRP with only a single radiator on is +2.5 dBm, yielding an array gain of 2.3 dB (not counting the 3 dB due to the doubling of the source power). The antenna pattern with one-dimensional beam steering for linear polarization and polarization angle of $\phi = 0°$ is shown in Figure 4.15.

The spectrum of the power captured by the receive antenna is shown in Figure 4.16.

A performance summary and table of comparison to other fully integrated radiators with on-chip antennas in silicon without post-processing is shown in Table 4.1. To the best of the authors knowledge, this is the first fully integrated radiator with dynamic polarization control, and additionally has the highest total radiated power, as well as highest EIRP per element reported for a fully integrated radiator with on-chip antennas in silicon without external modifications.

The die photo of the chip is shown in Figure 4.17.

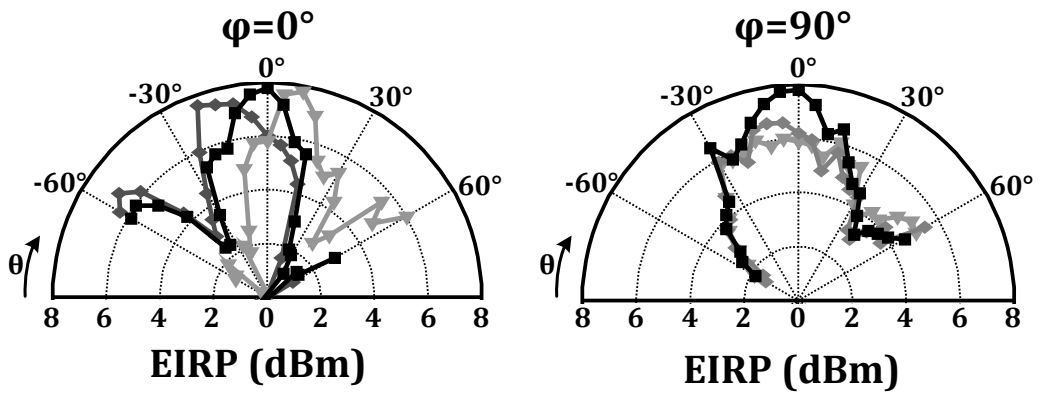

Figure 4.15: Measured antenna patterns of the 2x1 DPC radiator array along the axis of the array and perpendicular to that axis showing the ability to steer the beam in one dimension..

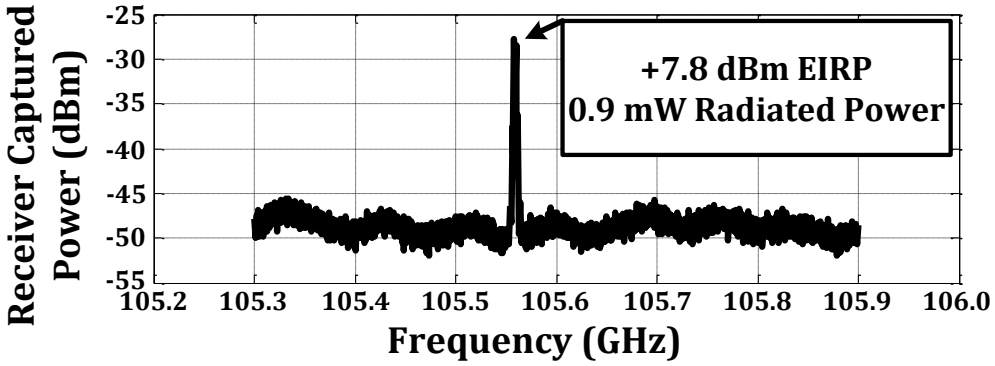

Figure 4.16: Measured spectrum of the 2x1 DPC radiator array shows a captured power of -27 which equates to an EIRP of +7.8 dBm and a total radiated power of 0.9 mW.

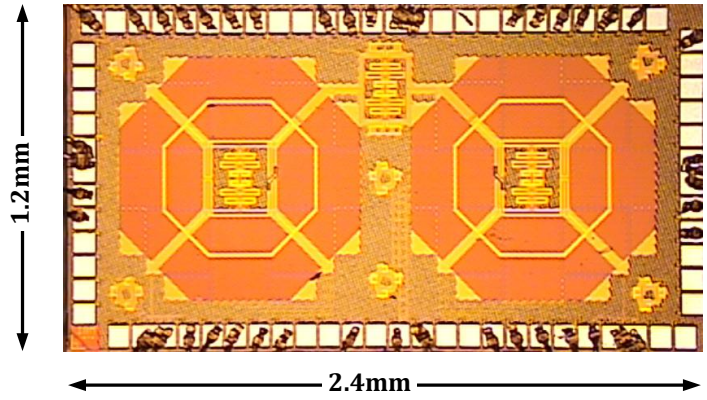

Figure 4.17: Die Photo of 2x1 Multi-Port Driven Radiator Array with Dynamic Polarization Control.

Table 4.1: Comparison of integrated radiating sources in silicon without external modification.

Reference	This Work	[78]	[33]	[37]	[38]
Polarization Control	Dynamic	N/A	N/A	N/A	N/A
Polarization	Linear, Circular, Elliptical	Circular	Circular	Linear	Linear
Linear Polarization Tuning Range	Full 0°-180°	N/A	N/A	N/A	N/A
Polarization Axial Ratio Tuning Range (dB)	2.4-13.25	N/A	N/A	N/A	N/A
Frequency (GHz)	105.5	161	191	210	165
Radiated Power (dBm)	-0.5	-2	-12.4	N/A	-27
Maximum EIRP (dBm)	7.8	4.6	-1.9	5.13	N/A
Number of Elements	2	1	4	4	1
Process	32nm CMOS SOI	130nm SiGe BiCMOS	65 nm CMOS	32nm CMOS SOI	130 nm SiGe BiCMOS
Beam Steering	1-D	N/A	2-D	N/A	N/A
DC Power (mW)	476	388	77	240	800*
Area (mm^2)	2.64	1	1.1	3.5	1.3 *

*Includes transmitter and receiver.

4.4 Polarization Modulation

In addition to being able to control the polarization and maximize coupling between two antennas, having control of two orthogonal polarizations means that four independent scaler dimensions of data can be supported by free space electromagnetic (EM) propagation in any given direction, for any given frequency. Two dimensions come from the amplitude and phase of the signal of the first polarization, and two from the amplitude and phase of the second polarization. The amplitude and phase of the signals can also be encoded through two signals that are 90° apart in phase, I and Q, which yields the same data transfer. Most antenna systems and modulation schemes currently employed only take advantage of one polarization, and are thus at most two-dimensional modulation schemes. Quadrature amplitude modulation (QAM) is one such system that takes advantage of two of the dimensions available (Figure 4.18a). By utilizing both polarizations, a 4-dimensional QAM modulation scheme (Figure 4.18b) can be envisioned, squaring the data throughput per symbol that can be transferred.

Rather than sending 4 data streams through two polarizations, the data can also be encoded in the polarization itself, as shown in Figure 4.19. One way of doing this is to encode data in the spatial angle polarization of a linearly polarized field. If the data is transferred in this way, the phase and amplitude of the signal can be set at will to mask the true data set, increasing the security of the stream. In fact, because the polarization angle requires the detection of polarizations that are linearly independent to each other, it is impossible to detect the signal with an antenna with a single output port. This can be understood by remembering that every single port antenna has a polarization, and there is always an orthogonal polarization to it. The single port antenna cannot distinguish between a signal that has zero amplitude or a signal with finite amplitude that is in the orthogonal polarization. However, a polarization agile multi-port antenna or antenna array that can detect

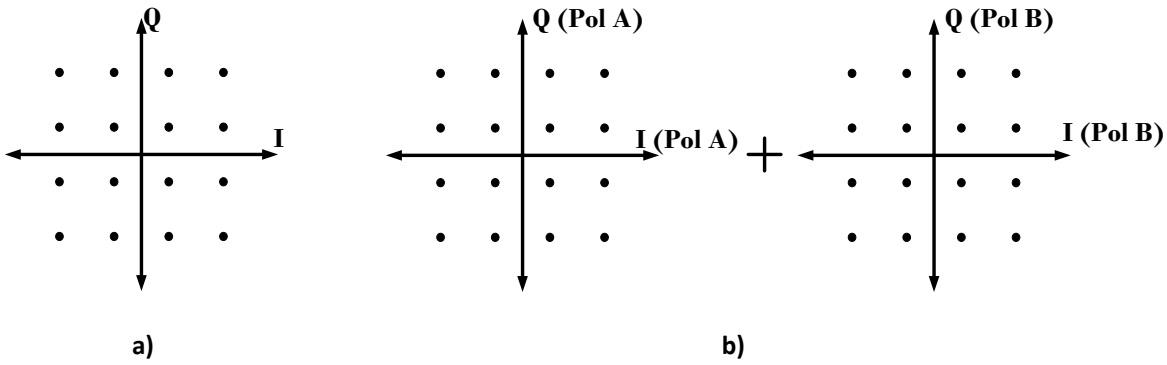

Figure 4.18: Constellation of a 16 QAM signal that uses both phase and amplitude to send 4 bits per symbol (a), and a possible four-dimensional 256 QAM that utilizes 16 QAM in two polarizations to send 8 bits per symbol (b).

both polarizations can decode the polarization angle regardless of the misinformation that may be encoded in the amplitude and phase of the signal.

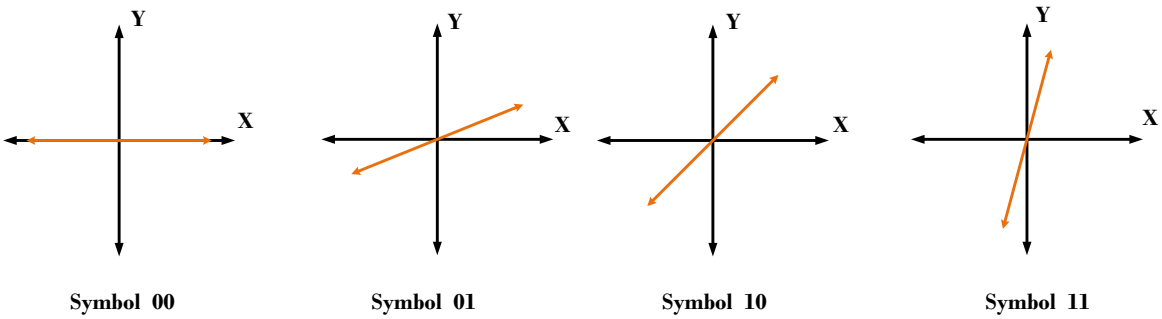

Figure 4.19: Example of a polarization modulation encoding scheme, where the two bits of data are coded into the angle of the polarization.

Implementation, fabrication, and measurements of polarization modulation on an integrated radiator are an ongoing research path.

4.5 Conclusions

Dynamic polarization control is a promising solution to combat signal power loss when the coupling between the transmit and receive antennas is decreased due to polarization mismatch. Two ways of implementing DPC have been presented. The first involves combining the output fields of several antennas in the far field to create the desired polarization. One example of this that was shown depicted using two circularly polarized antennas to create a linear polarization with a dynamically controlled polarization angle. The other implementation of DPC radiators is to use a single MPD radiator capable of radiating any polarization on its own by controlling the amplitude and phase of the various signals on its input ports. A proof of concept 2x1 MPD radiator array

with DPC was presented that is capable of tuning both the axial ratio as well as the polarization angle in a number of various modes and radiation directions. Finally, polarization modulation was proposed as a new way to increase data throughput through a channel or to increase the security of the link, as single polarization receivers would not be able to detect a signal sent through polarization modulation.

Chapter 5

Optically Generated mm-Wave Power

5.1 Integrated Silicon Processes[1]

The integration of photonics components onto a silicon substrate can enable unprecedented increases in the complexity of the design, similar to the systems level advances that occurred after the integration of circuits onto the same substrate [79–82]. Currently, there are several foundries that are beginning to offer these photonics capabilities.

The process we have been working with starts with a silicon-on-insulator wafer. This wafer has an unroped bulk silicon substrate, with a 2 μm silicon dioxide insulator on top of it. Above that is a thin 220 nm layer of silicon that will act as the active silicon for the system and also support the propagation of light around the chip. There are three available layers of silicon etching, two partial etches 60, and 130 nm deep, and one complete etch to the insulator, as seen in Figure 5.1 [83–86].

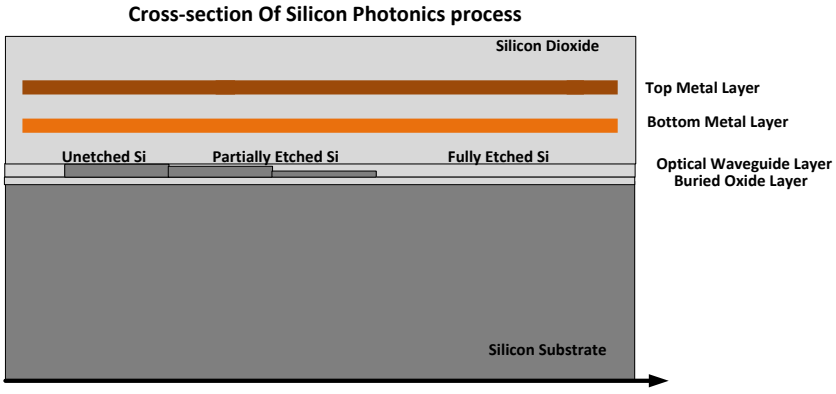

Figure 5.1: Cross section cartoon of the silicon photonics wafer showing silicon-on-insulator arrangement and the three levels of silicon etching, as well as the two metal layers (Note: not to scale)

[1]This project was a joint collaboration with my colleagues Dr. Firooz Aflatouni and Behrooz Abiri. Firooz did the device modeling of the traveling wave photodiode, and provided insightful discussions throughout the project. Behrooz was in charge of the optical distribution and optical modulators, and I handled the electromagnetics and radiation.

5.1.1 Waveguides

The primary component is the integrated silicon waveguide [87]. This can transport the light from one point on the chip to another, and is implemented in two ways. The first is a strip waveguide, which is created by etching the silicon all the way down to the insulator level, leaving a rectangular prism shaped silicon line that guides the light as a dielectric rectangular waveguide, as shown in the cross section of the device in Figure 5.2 [88]. The second way to create a waveguide is using rib waveguides [89]. In this method, the silicon on the outside of the line is not etched all the way to the insulator, but rather puts the waveguide on a pedestal of silicon, as shown in the cross section in Figure 5.3. The light is still guided by the waveguide, but the mode is distributed in a larger area and has less loss. The downside of these rib waveguides is that the radius of curvature of the rib waveguide is larger than in the case of the strip waveguide.

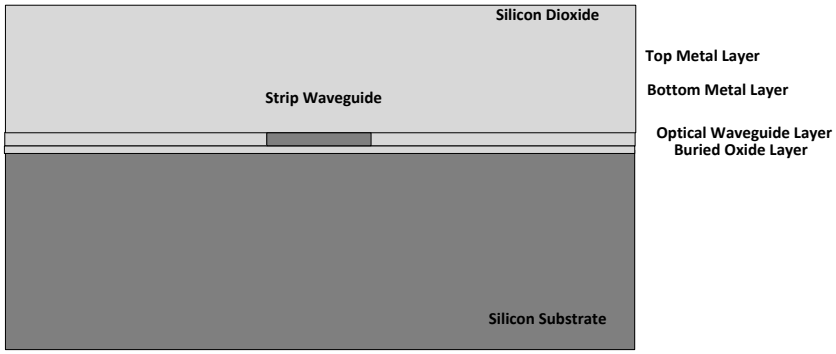

Figure 5.2: Cross section cartoon of a strip waveguide, where the light would be going into or coming out of the page, and is contained completely within the rectangular strip waveguide.

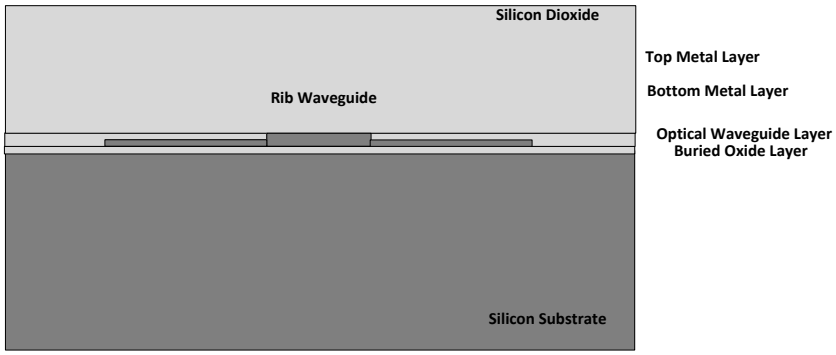

Figure 5.3: Cross section of the rib waveguide, where the light would be going into or coming out of the page, and is contained both within the waveguide as well as in the pedestal surrounding it.

5.1.2 Photodiodes

Another component of the integrated silicon photonics process is the photodiode. The photodiode is implemented as a silicon germanium PIN type photodiode. A layer of germanium is deposited onto the P-type silicon, and the top level of the germanium doped N-type. When light passes into the germanium, electron-hole pairs are created, and a current is formed that can be detected. In general, the diodes are reverse biased to increase the depletion width of the diode junction, and to ensure that as many charge carriers reach the diode terminals as possible by increasing the electric field across the junction. Device simulations put the responsivity of the diode at around 0.8 A/W. The light enters laterally into the side of the diode, and is modeled as a traveling wave diode, or a series of diodes in parallel, with inductances and capacitances modeled [90].Thus even though the phase of the signal changes as the light travels through the germanium, a similar phase shift will occur in the electronic path through the delay of the metal lines.

The diode was optimized for maximum bandwidth and responsivity, and has a 3 dB simulated bandwidth of 30 GHz. This first pole is primarily due to the junction capacitance of the diode. A second pole that is due to the transit bandwidth is expected around 120 GHz. A load-pull can be run, similar to the design of a power amplifier, and the optimal impedance for a given frequency can be shown. When operated at the 3 dB bandwidth of 30 GHz, the load pull has an optimum impedance of 300 Ω, and produces a maximum -4 dBm output power, as shown in Figure 5.4. In order to use this diode in a THz frequency multi-port driven (MPD) radiator, the load pull is run at 350 GHz, and shows an optimum load of 200 Ω that produces -25 dBm power, shown in Figure 5.5.

5.1.3 Optical Phase Modulators

A third component used in this process is optical phase modulators [91]. Phase modulation is accomplished in two ways: thermally or through carrier depletion. Phase modulators are used that can work in either mode. For a unit section of phase modulation, a P-type region is run along one side of a waveguide, and an N-type region is run along the other. Contacts are placed on both sides of both doping regions, as shown in Figure 5.6. In order to use the phase modulator in thermal mode, the two contacts on the input are tied together, and the two contacts at the output are tied together, and a voltage is passed across them. The doped regions of silicon act as resistors, and as the voltage increases, the power dissipated in the resistors increases and the waveguide's temperature increases. As the temperature increases, the index of refraction decreases, causing a phase shift in the optical signal over the same length. Carrier injection mode is achieved by tying the two P-well contacts together, and the two N-well contacts together, and controlling the voltage between them. A faster modulator, on the order of tens of GHz, can be achieved by using the diodes in reverse bias, and controlling the width of the depletion region in the waveguide, which changes the index of refraction and thus the phase, but the gain of the phase shift is very low, requiring a large length of modulator that results in higher loss. The diodes can also operate in forward bias, and the gain of the phase shift increases, but the

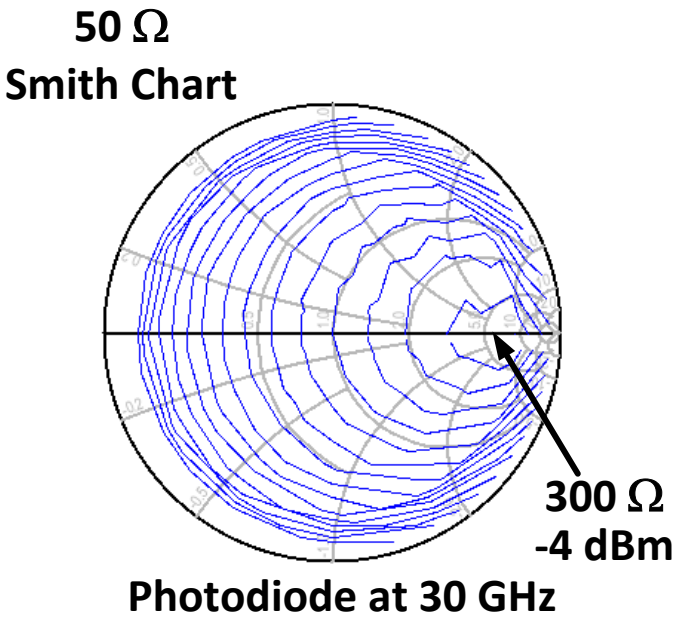

Figure 5.4: Load pull of the saturated photodiode at 30 GHz shows with an optimal load of 300 Ω, the photodiode can produce a maximum of -4 dBm.

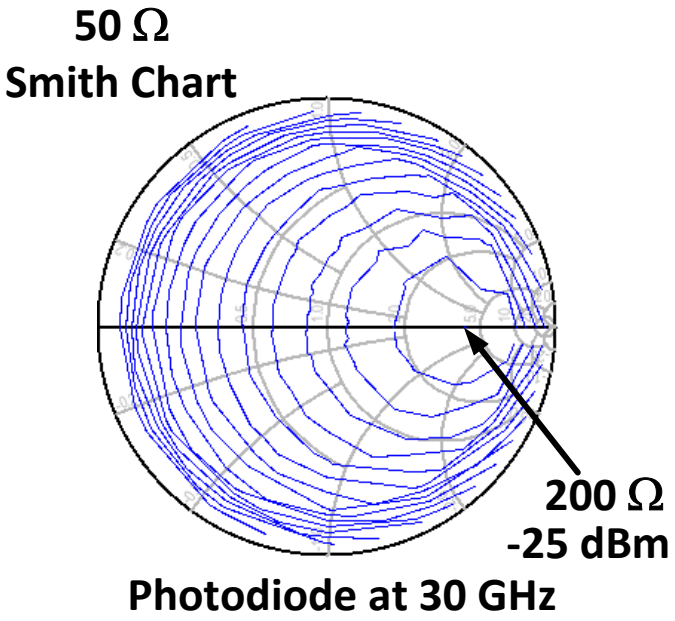

Figure 5.5: Load pull of the saturated photodiode at 350 GHz shows with an optimal load of 200 Ω, the photodiode can produce a maximum of -25 dBm.

diode turning on and off has a lower bandwidth of less than a GHz. For this design we use modulators with forward bias, as minimizing overall loss in the line was more important than the speed of the phase shift.

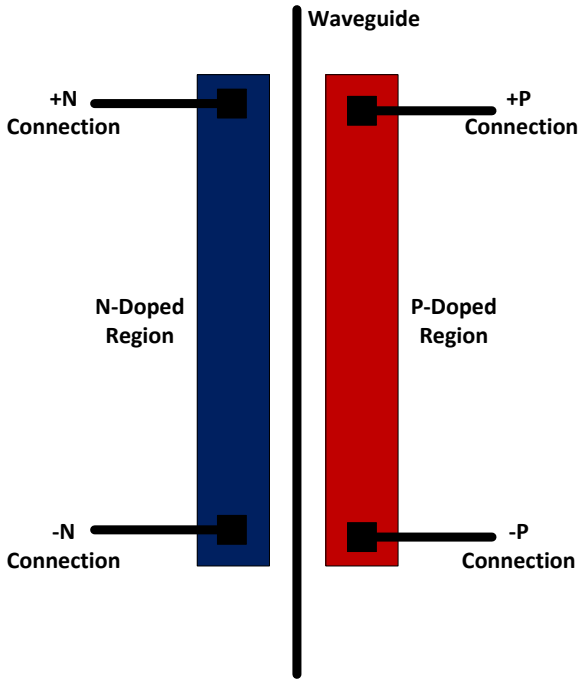

Figure 5.6: Layout of the optical modulator that has two contacts on each doping region to enable phase modulation in either a thermal or carrier injection mode.

5.1.4 Metal Stack

Finally, the available metal stack is two metal layers as shown in Figure 5.1. Having only two metal layers severely limits the electromagnetic structure possibilities compared to more mature integrated circuit processes, but because there are two metals, transmission lines, inductors, and antennas can still be implemented. The contact resistance for the original designs was expected to be less than 500 Ωcm$^\circ$, which is higher than in modern integrated circuit (IC) processes, but still allows for good connections to the diodes. The sheet resistance of the metals is similar, though slightly lower than the metals used in the ICs used for the fully electronic versions of the MPD radiator. The large benefit to the radiating structures is that the bulk of the silicon is unroped, reducing a very large source of loss that occurs in the other designs. It also makes the design very tricky, as the substrate itself becomes a resonator with a high quality factor, and all of the tolerances to substrate dimension inaccuracy become much less.

5.2 Producing THz Signals from Two Lasers

A radio frequency (RF) current can be generated from the photodiode if power from two optical frequencies are input into the diode. The diode acts as a rectifier, that is, it takes the input optical signals and outputs a current response that is proportional to the power of the optical drive. If two optical signals are added together, the rectification produces a beat frequency, and the difference of the two frequencies is produced in the electrical domain as shown in Figure 5.7.

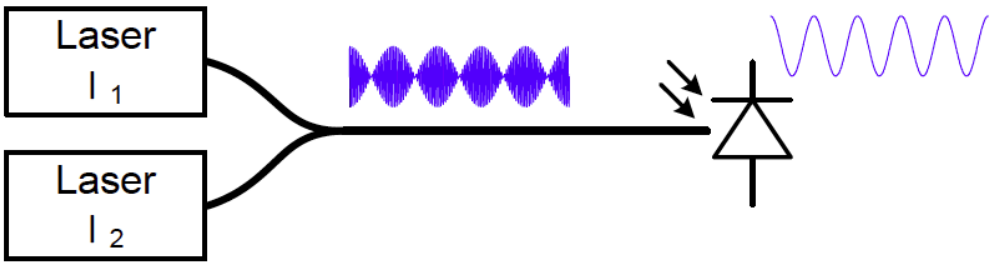

Figure 5.7: THz signal production scheme using two lasers with slightly different frequencies in the optical domain to produce a THz beat frequency at the output of the photodiode in the electrical domain.

The two optical signals, E_1 and E_2 can be written as

$$E_1 = \sqrt{P_1}\cos(\omega_1 t + \phi_1), \tag{5.1}$$

and

$$E_2 = \sqrt{P_2}\cos(\omega_2 t + \phi_2). \tag{5.2}$$

To simplify the analysis, we assume the power levels are the same, or $P_1 = P_2 = P$. When added together,

$$E_1+E_2 = \sqrt{P}[\cos(\omega_1 t+\phi_1)+\cos(\omega_2 t+\phi_2)] = 2\sqrt{P}\cos\left[\frac{(\omega_1+\omega_2)t+\phi_1+\phi_2}{2}\right]\cos\left[\frac{(\omega_1-\omega_2)t+\phi_1-\phi_2}{2}\right]. \tag{5.3}$$

The power of this field is proportional to the square of the electric field, and assuming the diode completely rectifies out the optical carrier frequency, and outputs a current, I_d, proportional to the magnitude of the beat envelope power with responsivity R, the current produced by the diode

$$I_d = 4PR\cos^2\left[\frac{(\omega_1-\omega_2)t+\phi_1-\phi_2}{2}\right] \tag{5.4}$$

$$I_d = PR + PR\cos\left[(\omega_1-\omega_2)t+\phi_1-\phi_2\right]. \tag{5.5}$$

Considering only the AC current, and removing the DC offset yields

$$I_{dAC} = PR\cos\left[(\omega_1 - \omega_2)t + \phi_1 - \phi_2\right]. \tag{5.6}$$

This shows that RF produced by the diode will be at the difference in optical frequencies, and that the phase of that signal will be the difference between the two optical phases. Like any real rectifier, there are poles associated with the photodiode that limit its performance at high frequencies, leading to the 30 GHz 3 dB bandwidth reported in the previous section.

5.2.1 Radio Frequency Phase Shifting Using Optical Phase Modulators

To control the phase of the output RF signal, the phases of the input optical signals can be delayed, causing a phase shift in the optical signal. It is important to note the difference that a time delay before or after the two optical signals have been combined affects the phase shift of the electrical signal. The phase that matters in the end is the electrical phase, which is the phase of the beat signal. The beat phase is determined by the difference between the two phases of the signal frequencies.

We can also realize that a time delay Δt causes different phase shifts $\Delta\phi$ for the different optical signals, where

$$\Delta\phi_1 = \omega_1\Delta t_1, \tag{5.7}$$

and

$$\Delta\phi_2 = \omega_2\Delta t_2. \tag{5.8}$$

Because the phase of the beat signal at the diode is determined by the difference between these two phases,

$$\Delta\phi_d = \Delta\phi_1 - \Delta\phi_2 = \omega_1\Delta t_1 - \omega_2\Delta t_2. \tag{5.9}$$

This means that if only one of the two optical signals is delayed ($\Delta t_2 = 0$), then

$$\Delta\phi_d = \omega_1\Delta t_1, \tag{5.10}$$

which produces a phase shift in the beat signal equal to the phase shift seen in the optical signal caused by that delay. However, if both signals are delayed by the same mount ($\Delta t_1 = \Delta t_2 = \Delta t$) then

$$\Delta\phi_d = (\omega_1 - \omega_2)\Delta t, \tag{5.11}$$

which is a phase shift that corresponds to the phase shift of that delay at the beat frequency. In this way, if a very high gain phase shift is desired, delaying only one of the two optical signals produces phase shifts at the beat frequency at optical phase-shifting delays. However, if a precise beat phase is desired, adding delay

after the two signals have been combined results in a phase shift on the beat frequency that corresponds to the phase shift for that delay at the electronic THz frequency. Both of these types of phase shifting are used in various parts of the design, as will be described below.

5.3 Optical Multi-Port Driven Radiators on Silicon Photonics Process

The optical MPD radiator is designed and implemented at 350 GHz. It is a 2-spoke design, that replaces the output transistors on the electrical MPD radiator with photodiodes. To increase the total power of the radiator, two photodiodes were placed at the end of each spoke. One great benefit of the optical MPD radiator is that the input lines are optical silicon waveguides, and do not interact in any way with the metal lines that carry the DC and THz signals. This means that the optical lines do not need to be routed under the spokes, but can be routed completely independently of the radiator's metals. The radiator is designed to have a circular polarization, with a 90° phase shift between each pair of driving photodiodes. The radius of the radiator ring is 100 μm, and bias to the photodiodes is provided through a single connection from the ring to the outer VDD plane over one of the ground spokes, as seen in Figure 5.8. The length of the line is $\lambda/2$, providing a high impedance on the VDD spoke, and causing a minimal effect on the RF current distribution.

The optical light signal is coupled into the waveguides by either a grating coupler from the top [92], or edge coupled from the side. The benefit of the grating coupler is that it is mechanically compatible with RF probing, and thus alignment using the current equipment available in the lab is easier. However, once aligned, the fiber can be glued using the edge coupler and is more robust, especially once the chip begins to move, for example to rotate to take antenna pattern measurements. Because of this, both couplers were used on the chip, with a Y-junction combining structure used to send signal into the main line.

The unit radiator takes an optical waveguide with both laser frequencies already combined. This signal is split into 4 using two layers of Y-junction dividers. Each of the four signals are sent to the photodiodes with the appropriate delay, and are split once more directly before feeding the diode pair, as seen in Figure 5.9. Using two diodes per MPD antenna input enables a higher total output power before the diodes saturate, and provides a better impedance match to the antenna. The phase shifting is accomplished by meandering the lines by different amounts once they split within the unit radiator to provide different time delays at each diode. Figure 5.10 shows the entire radiator with both the optical as well as the THz halves shown.

The 90° phase shifting for each diode occurs after the lasers have been combined into the same waveguide, which means that the delay must be equal to the delay for 90° at 350 GHz, rather than at optical frequencies. This is important because it means that very small differences in delay due to process variation have much less of an impact than if the phase was caused by delays corresponding to the optical frequency.

This basic unit radiator is used in 3 separate 2x2 radiator array designs. The first combines the lasers off chip, and routes the combined signal from the input coupler to the 4 radiators after two stages of Y-junction dividers, without any dynamic phase modulators (Figure 5.11). The second has phase modulators for each

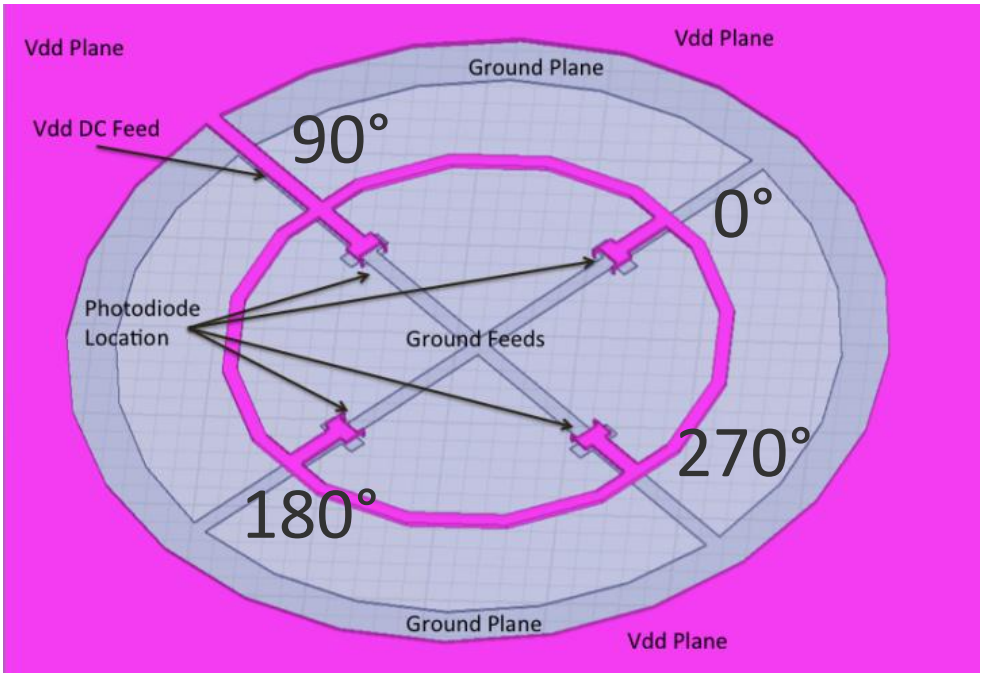

Figure 5.8: Overhead view of the metal structures of the optical MPD antenna, with the VDD lines highlighted in pink, and the ground lines in blue. The four photodiode locations that drive the signal ring against the ground spokes are marked as well.

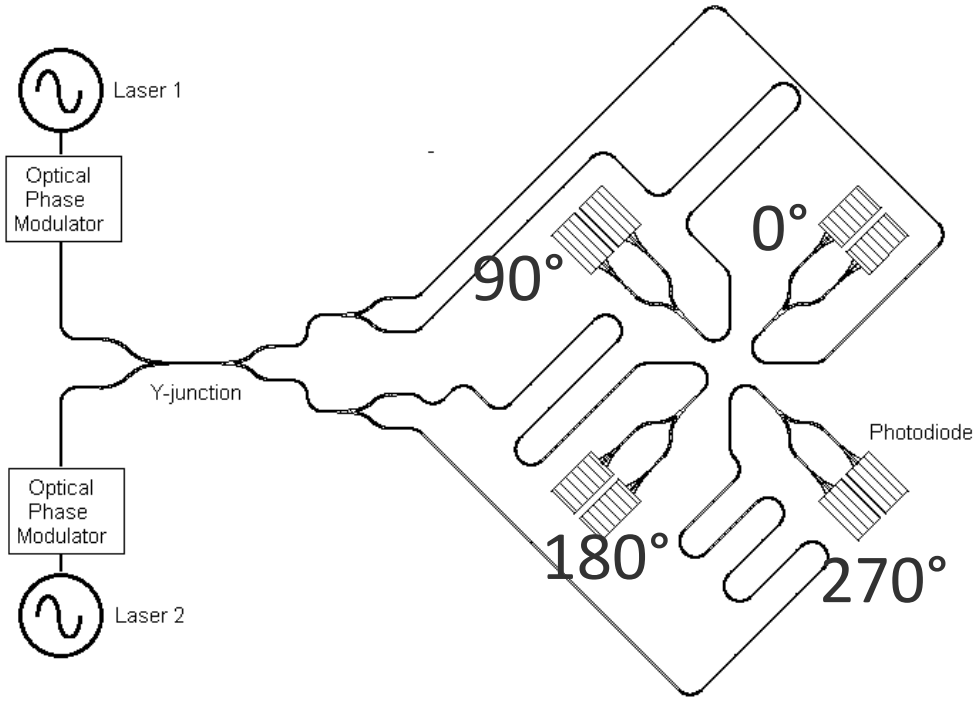

Figure 5.9: Distribution of the optical signal for a single MPD radiator shows phase shifting by using meandered delay lines.

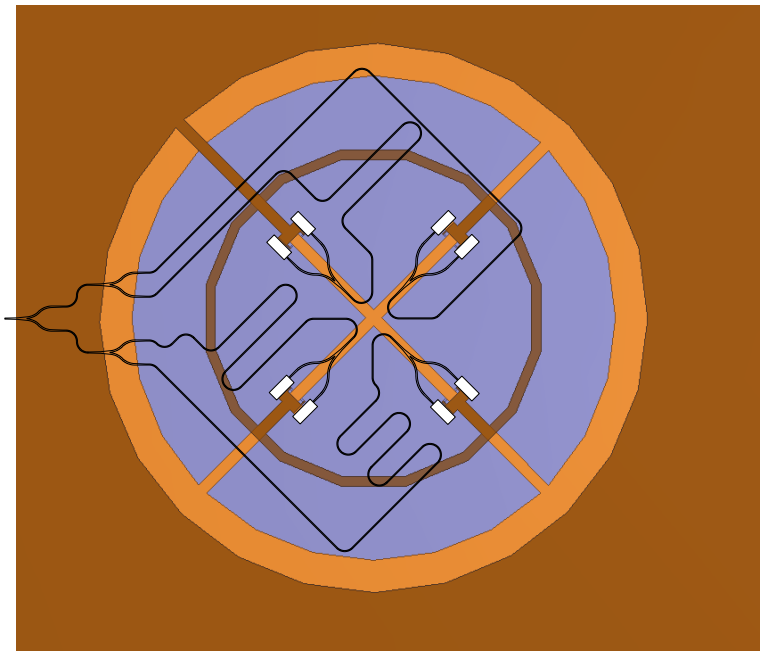

Figure 5.10: The layout of an optically driven MPD radiator, with the optical and THz sections both shown on top of each other. The optical waveguides are in black, the photodiodes are in white, the ground lines are in light orange, and the VDD lines are in brown.

of the four radiators to enable beam steering of the array (Figure 5.12), and the third uses a mirror image of the optical distribution network for two of the radiators to enable multi-antenna dynamic polarization control (Figure 5.13).

The maximum gains for the first two radiators is 12.1 dBi, with the antenna pattern shown in Figure 5.14. The third chip with dynamic polarization control (DPC) has a simulated 10.9 dB gain and the antenna pattern is shown in Figure 5.15. With each diode providing -25 dBm, the 32 diodes provide -10 dBm output power, yielding a simulated effective isotropic radiated power (EIRP) for the first two chips of 2.1 dBm, and an EIRP of the third chip of 1.0 dBm. The DPC array has phase modulators that control the relative phase of each type of polarization separately. As discussed in Chapter 4, the summation of the two circular polarizations will result in a linear polarization, whose polarization angle is determined by the difference in phase between the two polarizations. In this way, the polarization angle of the resulting linearly polarized field can be controlled fully from 0° through 180° simply by changing the relative phases of the two radiator polarizations. Figure 5.16 shows this tuning, where the polarization angle is tuned over the full range, while the gain of the antenna has a maximum deviation of 0.35 dB from the mean, and the polarization ratio, defined as the ratio of the two circularly polarized field remains within 0.045 dB of 0 dB, which implies very good linear polarization.

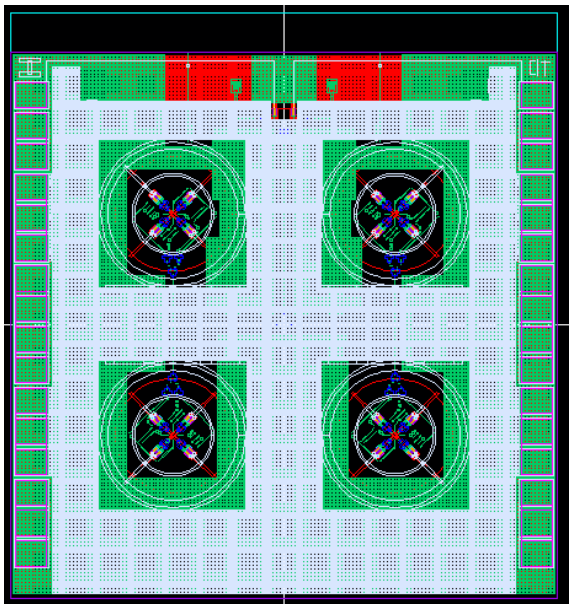

Figure 5.11: Layout of the first chip, a 2x2 optically driven MPD radiator array that does beam forming, but that does not have any optical phase modulators.

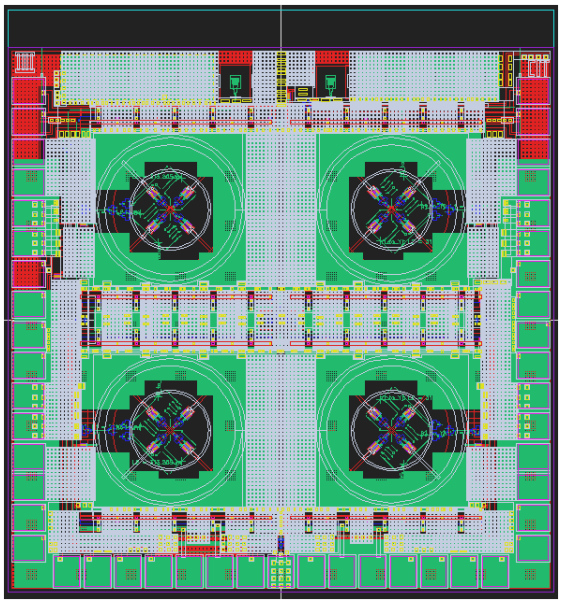

Figure 5.12: Layout of the second chip, a 2x2 optically driven MPD radiator array that has optical phase modulators for each of the radiators that enables beam steering.

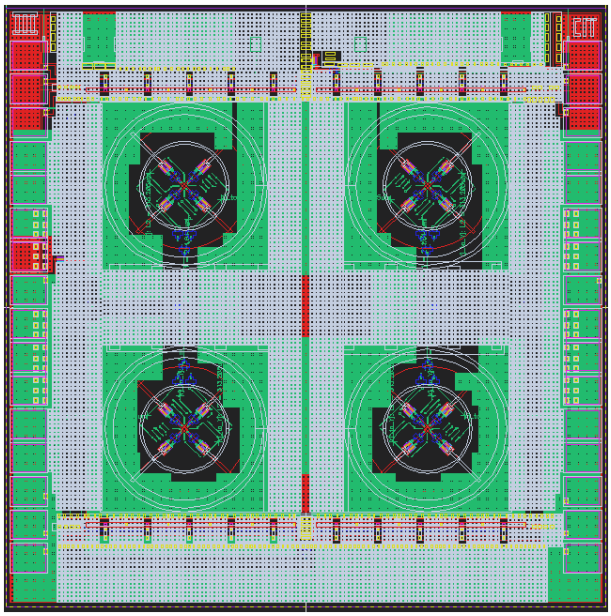

Figure 5.13: Layout of the third chip, a 2x2 optically driven MPD radiator array where two of the radiators are right hand circularly polarized, and two are left hand circularly polarized, with optical phase modulators that enable dynamic polarization control.

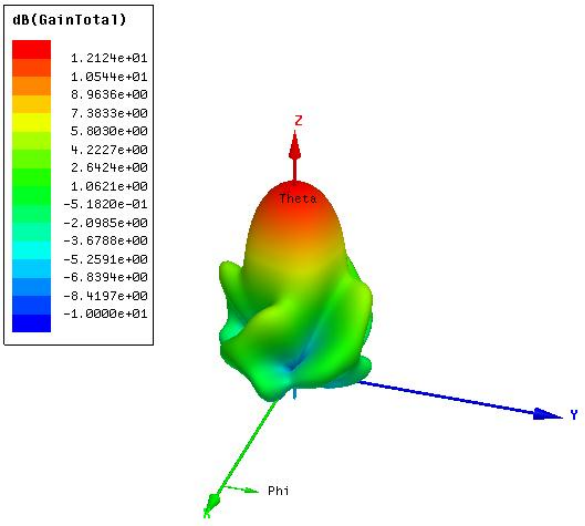

Figure 5.14: Antenna pattern of the beam forming and beem steering (Chips 1 and 2) 2x2 optically driven MPD radiator arrays showing a maximum gain of 12.1 dBi radiated upward.

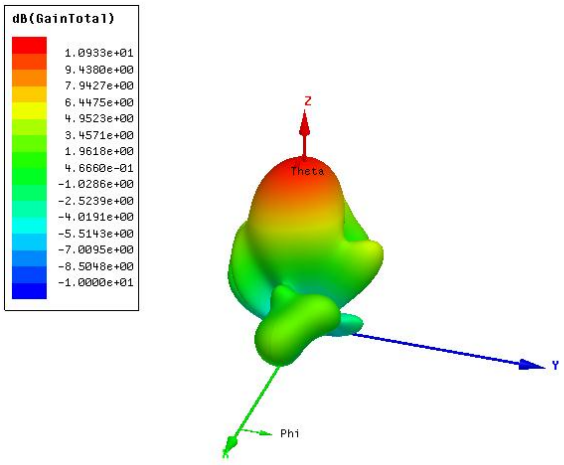

Figure 5.15: Antenna pattern of the 2x2 optically driven MPD radiator array with dynamic polarization control (Chip 3) showing a maximum gain of 10.9 dBi radiated upward.

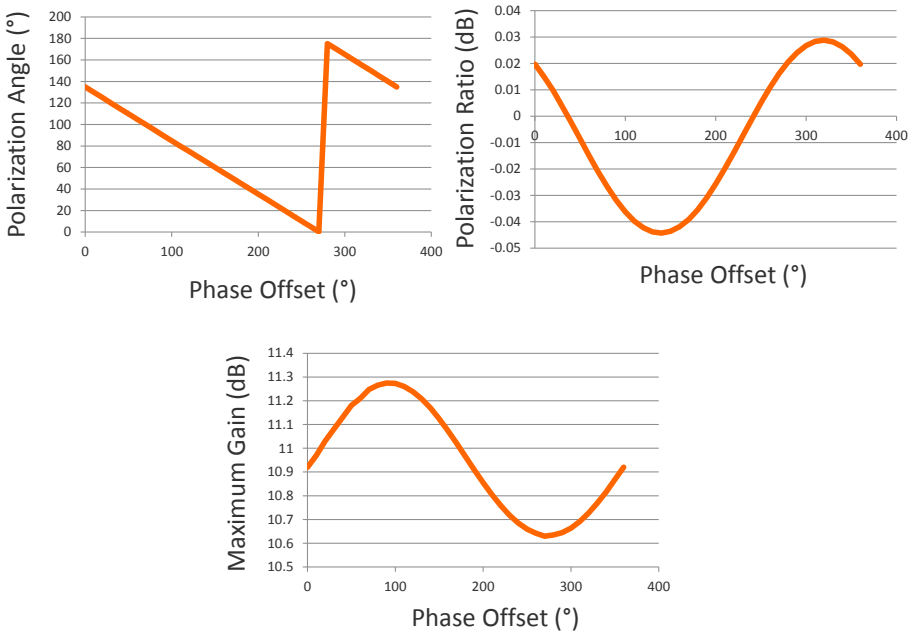

Figure 5.16: Simulation showing the full 0° through 180° control of the polarization angle of the optically driven MPD radiator array with DPC (upper left). The maximum fluctuation in the polarization ratio is -0.045 dB (upper right), with a maximum deviation of the antenna gain is 0.35 dB from its mean at 10.9 dBi (bottom).

5.3.1 Measurements from Optical Multi-Port Driven Antennas

DC measurements from the first fabrication run of the optical MPD antennas have been performed. Simulated estimates, as well as those provided by the fabrication service of the contact resistances were off by an order of magnitude, and contacts to the active regions of the substrate were often non-existent.

The dark current (current when light sources are off) of a single diode was measured in a test structure to be 2 uA at a 4V bias, meaning that for 32 diodes in the array, we expected 64 uA. However, after measuring multiple chips, they drew dark currents in the range from 16 uA through 35 uA, suggesting that many of the diodes were not being contacted.

Also, test structure Y-junction dividers show significantly more loss at 1.5 dB than the simulated 0.8 dB. In addition, in the initial simulations, an additional 3 dB of optical loss due to having two types of input couplers that are combined with a Y-junction combiner were not taken into account. Overall, this leads to an additional 11 dB less power in the optical path going to the diodes.

Taking into account the additional optical loss and the loss due to the larger contact resistances, the expected EIRP of the radiated signal falls to -20 dBm EIRP, which is below the noise floor of our current measurement setup.

5.3.2 Improvements for Second Fabrication Run

Taking these issues into account, a second fabrication run is underway. While numerous small improvements have been made to overcome the issues from the first run, the overall design has remained the same. The first change was in the layout of the diode. The contact areas were increased by a factor of 3 to lower the overall contact resistance of the diode. The Y-junction combiners were also redesigned to lower their loss to a simulated 0.3 dB. The final splitting of the power into 2 diodes per antenna port has been removed, and a single diode now drives each port. This was done to ensure that the diodes can be saturated by the lasers. Finally, the dual input couplers have been removed. For chip 1 and 3, grating couplers were kept, and chip 2 now uses edge couplers only. This removes the first Y-junction combiner that was adding 3 dB of extra optical loss to the system. With these improvements, the simulated EIRP of the first two chips improves to -2 dBm, and -3 dBm EIRP for the third chip, well within the noise limits of our measurement setup.

5.4 Conclusion

There are several interesting opportunities that are arising for designers as the integration of photonics components onto silicon increases. While these processes are currently very early in their design cycles, and performance cannot always be accurately predicted by models, it is expected that these issues will be worked out as the process matures. These photonics processes are expected to support transistors soon, and eventually integrate light sources on chip. At that point there should be many opportunities to integrate more complex

systems that currently occupy entire benchtop setups, which can enable entirely new applications that are currently not possible or cost prohibitive.

Chapter 6

Self-Healing of an Integrated Multi-Port Driven Radiator Array

The multi-port driven (MPD) radiator suffers from many of the same ailments of mm-wave power amplifiers (PAs) that prompted the investigation of a self-healing mm-wave power amplifier that is discussed in Chapter 7. The MPD radiator can also benefit greatly from a self-healing system, as phases and amplitudes matching of multiple mm-wave signals at significantly different locations of an integrated circuit can often end up mismatched. The issues of impedance mismatch due to antenna interaction with the near field are also present, as are variation due to process, mismatch, temperature, and aging, making the MPD radiator an ideal candidate for a self-healing system.

The fundamental difference between the self-healing MPD radiator and the self-healing power amplifier has to do with the detection of the performance metrics. While it is still feasible to directly detect the DC power consumed by the radiator, detecting the output power is significantly more difficult. The PA had a single output that could be tapped with a coupled transmission line to directly measure the power being delivered to the load. On an MPD antenna, there are individual drive points for each spoke, and if the radiator is an array of MPD antennas, the number of output drives increases. On top of that, the desired metric is not just the output power, but also the far-field radiated pattern, and the far-field radiation cannot be directly detected by integrated sensors on a chip. This means that indirect sensors must be employed to detect enough information about the circuit operation to make an inference about the far-field radiation.

One way of doing this it to create substrate mode pickup sensors.An integrated MPD antenna will initially radiate most of its power downward, into the substrate. It can either be radiated out through the backside of the substrate, or reflected off of a ground plane and radiated upward out through the top of the chip. Either of these use cases will create undesired substrate modes within the substrate. Accurate detection of these substrate modes can yield enough information to be able to identify the drives of the MPD antenna, and thus are enough to recreate the far field pattern.

Substrate mode sensors are ideal sensors for this application because the power coupled to the substrate is lost to the system, and is thus a waste. With direct power measurements, such as the radio frequency (RF)

power sensor in the PA, the goal is to only siphon off the minimum power required to accurately detect the total power to the load, as any additional power will degrade the performance. On the other hand, because the substrate mode power is already a waste, as much power as needed can be pulled from the substrate. It is important, however, that the substrate sensors do not affect the useful radiation of the circuit, so care must be taken to couple only power from the substrate and not from the useful radiation. An algorithm is then used to modify the state of the radiation parameters based on the collected data. This process can occur once or be used as an iterative optimization, until an optimum actuation state is found for all actuators. If variables in the system such as environment change, the healing can be run again.

6.1 Mode Pickup Sensor Simulations

Proof-of-concept simulations have been performed to validate the viability of these pickup sensors. The first preliminary designs were done with a 2x2 MPD circularly polarized radiator array designed to radiate at 94 GHz radiating downward and shown in Figure 6.1. Most of the substrate in such a radiator is covered in a ground plane, and thus slot-type antennas placed within the ground plane are a natural fit for the substrate sensors. For this simulation, slot dipole antennas were placed in alternating polarizations along the X and Y axes. The aim here is to be able to pick up information about the electric field at various points in the chip and at various polarizations.

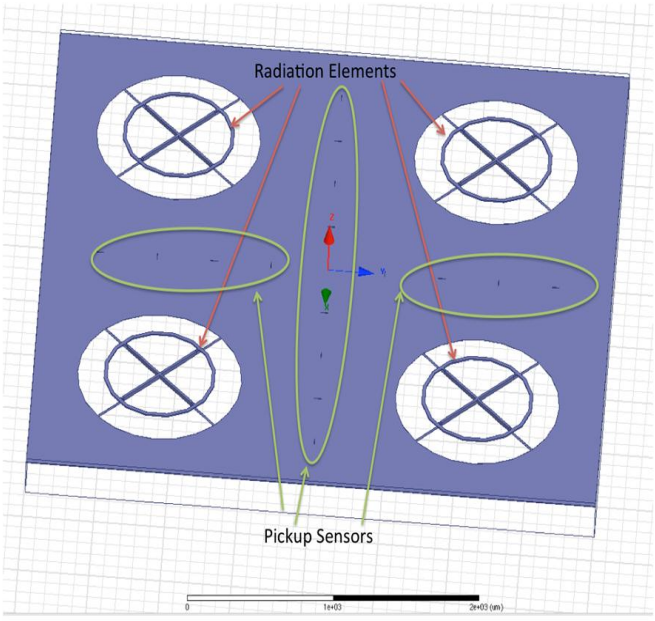

Figure 6.1: The simulation setup showing a 2x2 MPD radiator array and four sets of slot dipole mode pickup sensors, each with two sensors in each polarization.

The slots are small compared to the wavelength, meaning that they will not affect the radiation pattern much, while picking up enough power to sense effectively. The pickup sensor ports were terminated with a 50 Ω load to represent the input impedance of the power sensor they are attached to. A simplified drawing of the system is shown in Figure 6.2. The sensors are labeled A,B,C, and D, for each of the +X, -X, +Y and -Y directions. The radiators are designated as R1 through R4, with phases at each port marked.

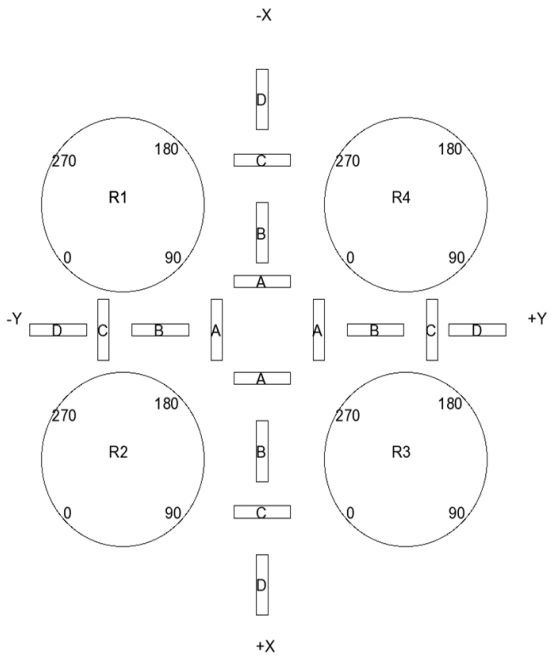

Figure 6.2: A simplified diagram of the simulation, showing each of the 16 driver input ports as well as the 16 sensor outputs.

6.1.1 Sensing Changes In Substrate Thickness

The first test was to see if there was enough data in these sensors to determine errors in the thickness of the substrate. To demonstrate the ability of these sensors to detect change in substrate height, which will affect substrate modes, substrate heights (zsub) of 175 μm, 225 μm and 275 μm were simulated, and sensor outputs compared. The entire structure was simulated in a 3-dimensional electromagnetic simulator, and the relative phases of the driver ports were varied.

The gain and directivity of the radiator array in the nominal case of 175 μm thick substrate is plotted in Figure 6.3, and show a maximum gain of 9 dB and directivity of 13 dB, which is the optimum condition for this system.

The sensor ports output the voltages seen in Figure 6.4. The Y axis of this plot shows the voltage generated at each sensor, while the X axis of the plot denotes the position on the chip where the sensor is located,

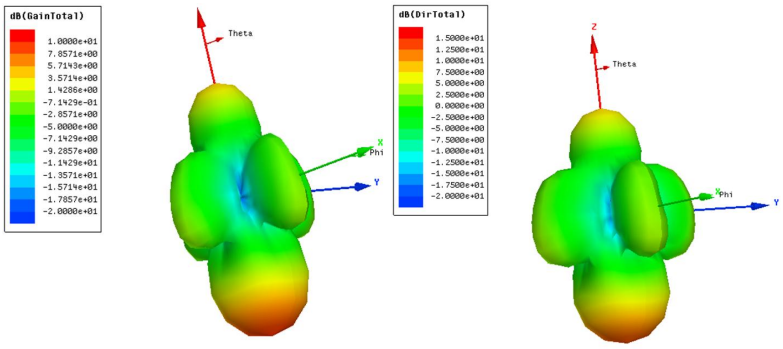

Figure 6.3: Antenna pattern for 2x2 MPD radiator array under nominal conditions showing gain (left) and directivity (right) and downward radiation

(sensors on the +X axis of the structure are plotted as 0, +Y as 1, -X as 2, and -Y as 3). The A sensors are plotted as SniffA in red, B sensors are SniffB in dark blue, C sensors are SniffC in magenta, and D sensors are SniffD in cyan. The total input power into the radiator is also shown in dBm. This format of plot will be used to consider the sensor output going forward. Note that when everything is symmetric, the sensors at the same position (A, B, C, or D) are all roughly picking up the same voltage.

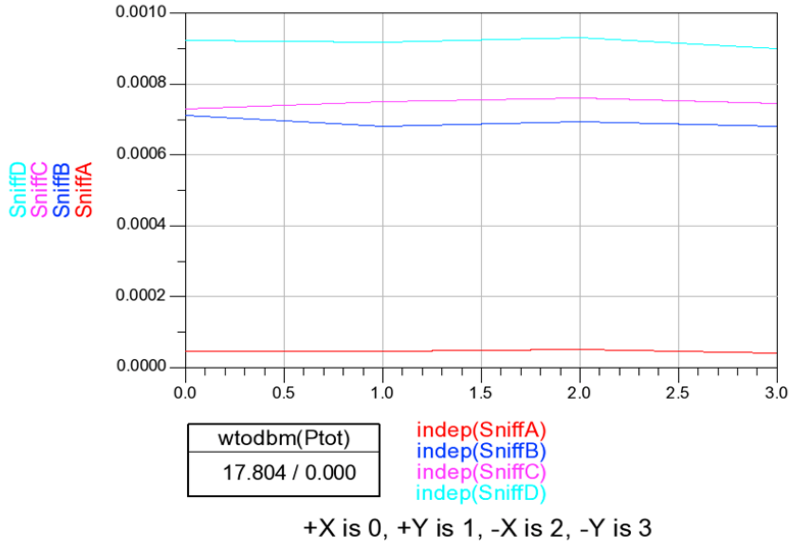

Figure 6.4: Nominal output from the 16 sensors. Each sensor location (A,B,C, or D) is plotted as a trace, with each axis being assigned a value on the independent axis, as denoted in the figure.

The next simulation was run with the substrate thickness increased to 225 μm. Figure 6.5 shows the gain, directivity, and sensor outputs at this thickness. The peak gain has reduced from 9 dBi down to 6 dBi, and

the directivity from 13 dBi down to 10 dBi. This 3 dB decrease in the performance of the antenna shows up in the sensor output voltages, with the D sensors along the outside of the chip going up by 60%, and the B sensors going down by more than a factor of 4. The entire system is still symmetric, so similar sensors on each axis are still close to the same levels.

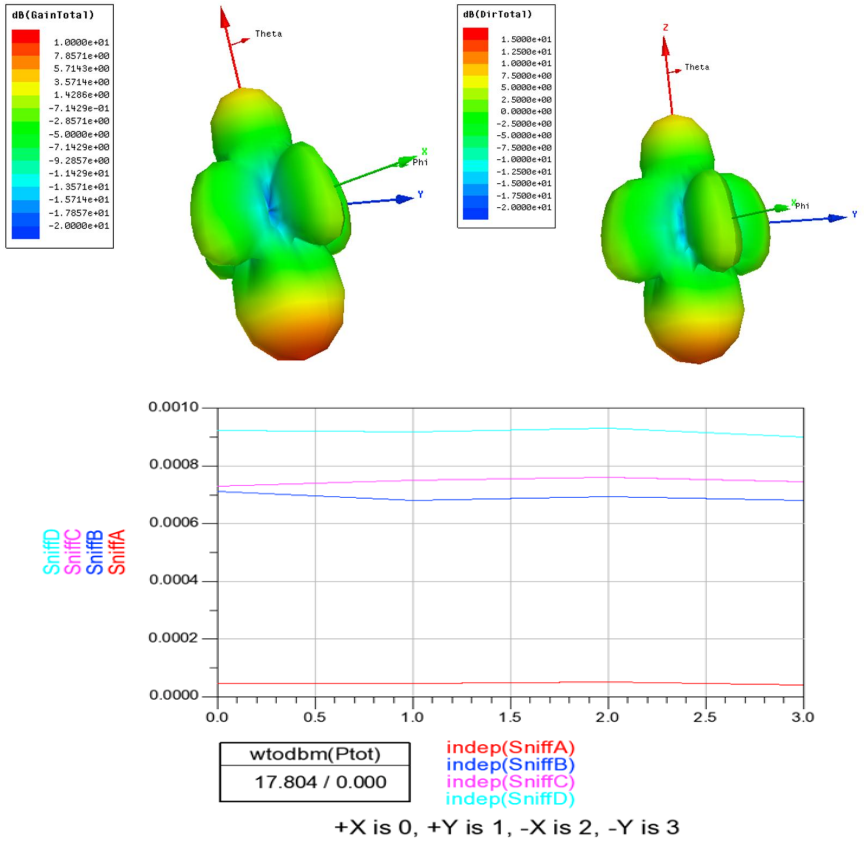

Figure 6.5: The antenna patterns for gain (upper left) and directivity (upper right) as well as the sensor output voltages when the substrate thickness is 225 μm.

Setting the thickness to 275 μm detunes the system even further. Figure 6.6 shows the gain, directivity, and sensor output under these conditions. The peak gain has now dropped to below 0 dBi, with peak directivity down to 6 dBi. The fact that the gain is falling faster than the directivity shows that the efficiency of the system is decreasing. The design is now about an order of magnitude off of its peak condition. At this point, the voltages on D have gone back down to similar levels to the nominal, but the sensors of A have continued to rise from 0.2 mV to 0.4 mV to 0.6 mV as substrate height increased. The sensors on B have also increased compared to the case of a 225 μm substrate. These simulations indicate that there is indeed enough information in the sensors to detect variations in substrate thickness, though the mapping from the sensors is complex. Different sensors are better at detecting different non-idealities. In this case, the sensors nearest the center were picking up a substrate mode that increased in power as the substrate thickness increased, which

led both to a decrease in radiated gain, and an increase in the voltage picked up by sensor A.

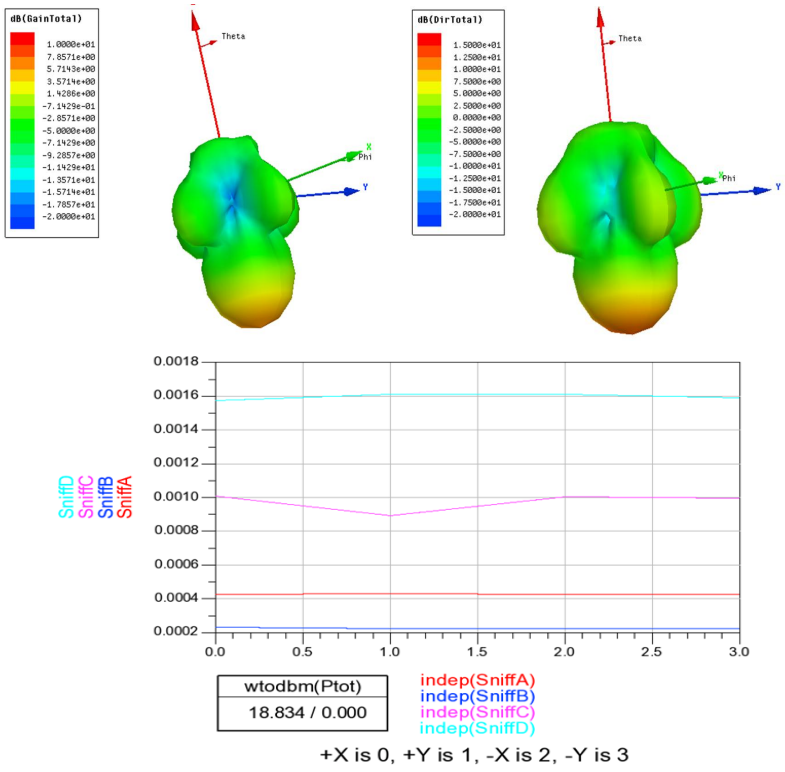

Figure 6.6: The antenna patterns for gain (upper left) and directivity (upper right) as well as the sensor output voltages when the substrate thickness is 275 μm..

6.1.2 Sensing Perturbations in Drive Phases

The second test for these mode pickup sensors is to see if they can detect changes in the phases of the drive points on the radiators. One way the phases can become mismatched is if an entire radiator's phases are off by a fixed amount. This would occur if there was mismatch in the delays of the locking networks for the radiators, where the locking signal for one radiator was different than the rest. Figure 6.7 shows the case when a single radiator has a phase shift of 30°. One should notice that the sensors on B and C are acting inverted to each other. For radiator 1 and 3, the C sensors on the Y axis increase, while the ones on the X axis decrease. The B sensors on the Y axis meanwhile decrease and on the X axis increase. This is the opposite for radiators 2 and 4, and gives a good clue as to which of the radiators is off. Sensor A can yield additional data, as the A sensor with the maximum voltage is always the one just counterclockwise to the radiator that is leading. Similarly the shape of the D sensors is unique to each radiator being out of phase, and can be used in the mapping from the sensor data back to the radiator phases.

A similar mapping can be obtained for a single radiator with a negative phase shift of 30°, as shown in

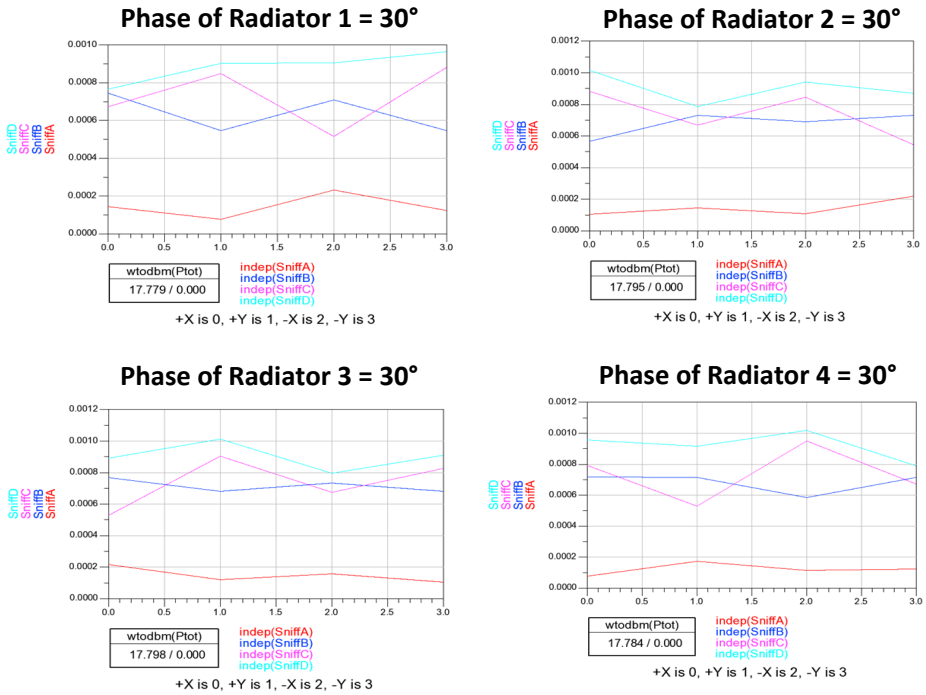

Figure 6.7: Sensor outputs when a single radiator is 30° ahead in phase.

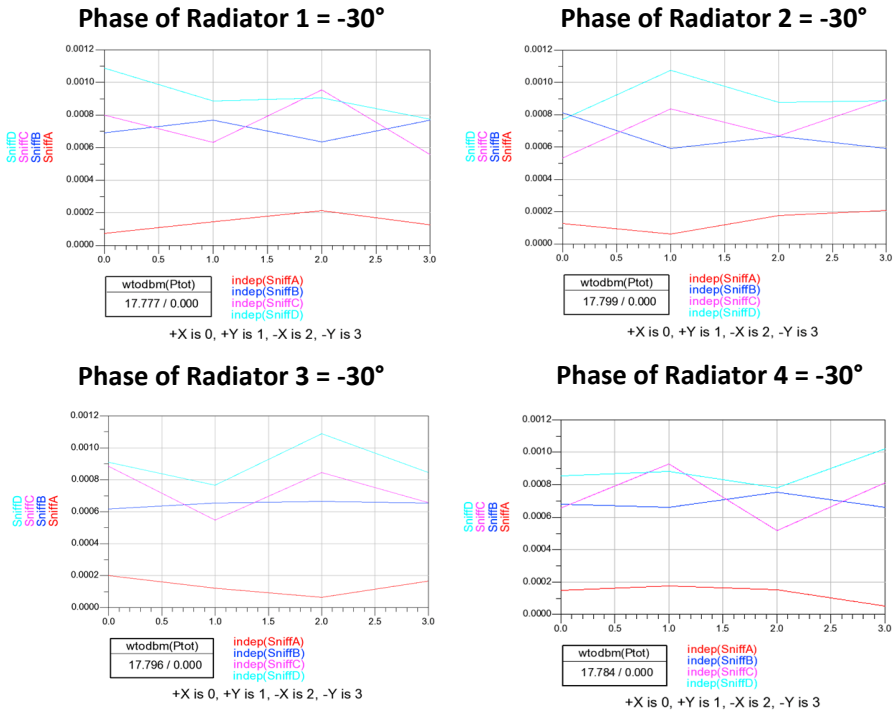

Figure 6.8: Sensor outputs when a single radiator is 30° behind in phase.

Figure 6.8. A similar, but inverted pattern for sensors B and C is observed. The key difference here is the relationship between the C and A sensors. While the A sensor always had a maximum when the C sensor had a minimum when the phase shift was positive, the maximum at the A sensor now occurs when the C sensor is also at peak. In this way, a single radiator whose phase is incorrect can be identified, and it can be determined if the radiator is leading or lagging.

Another variation that can potentially cause degradation in the system is if an individual port within the ring is out of phase. This will degrade the system as the traveling wave along the radiator ring will encounter a drive point that is not adding perfectly coherently, which will lower the efficiency and also increase power lost to substrate modes. As stated above, changing port X will change the relative phase the $X°$ ports in each of the rings, so setting port 0 = 30° will perturb the 0° ports to 30° in each of the rings. Each of the 4 sets of ports were set forward 30 °, and then set back 30°, and the sensor results are shown in Figure 6.9

These variations also produce a unique signature similar to the ones when an entire ring is perturbed. Sensors B and C now are moving in phase with each other, while sensor D is inverted from that pattern. A pattern where each plot is shifted 1 unit to the right between plots is also observed, as the ports being perturbed rotate 90° counterclockwise between each plot. This shows that individual perturbations in driver phase can be sensed by these pickup sensors.

The ability of the sensors to sense variation is not limited to variation in driver phase, but is applicable to any variation that affects the radiation of the system. Thus with the proper healing algorithm block and actuator space, knowing the cause of the variation is no longer necessary; the system only needs to be able to sense that a variation occurred, needs to be able to find an actuation that counters the perturbation, and then needs to be able to sense that the system is once again operating at an optimal state.

6.2 Mode Pickup Test Structures in 2x1 Multi-Port Driven Radiators[1]

Substrate mode slot ring sensors were added as test structures to test the feasibility of this concept on the 120 GHz 2x1 MPD radiator presented in Chapter 4. The locations of the mode pickup sensors are marked on the chip topcell layout in Figure 6.10

The slots ring sensors enable multiple polarizations of the substrate mode to be sensed at the same location because sense voltages are tapped off on each side, as seen in the layout in Figure 6.11. The sensor voltage is put to the gate of a transistor in cutoff that acts as a rectifier that converts the 120 GHz signal to DC. A stage of current amplification is followed by a conversion to voltage and two stages of differential amplifiers, where the opposite side of the amplifier is biased with a dummy rectifier circuit with no input for matching, as shown in Figure 6.12.

Simulations show that similar mapping should be possible with these slot ring sensors. An issue arose in these sensors with offset voltage mismatch causing the sensors to rail high or low. Measurements show

[1] Amirreza Safaripour also collaborated on this aspect of the project and designed the schematics of the sensor circuitry while I designed the electromagnetic slot ring pickups of the substrate modes, did the physical layout, and the measurements.

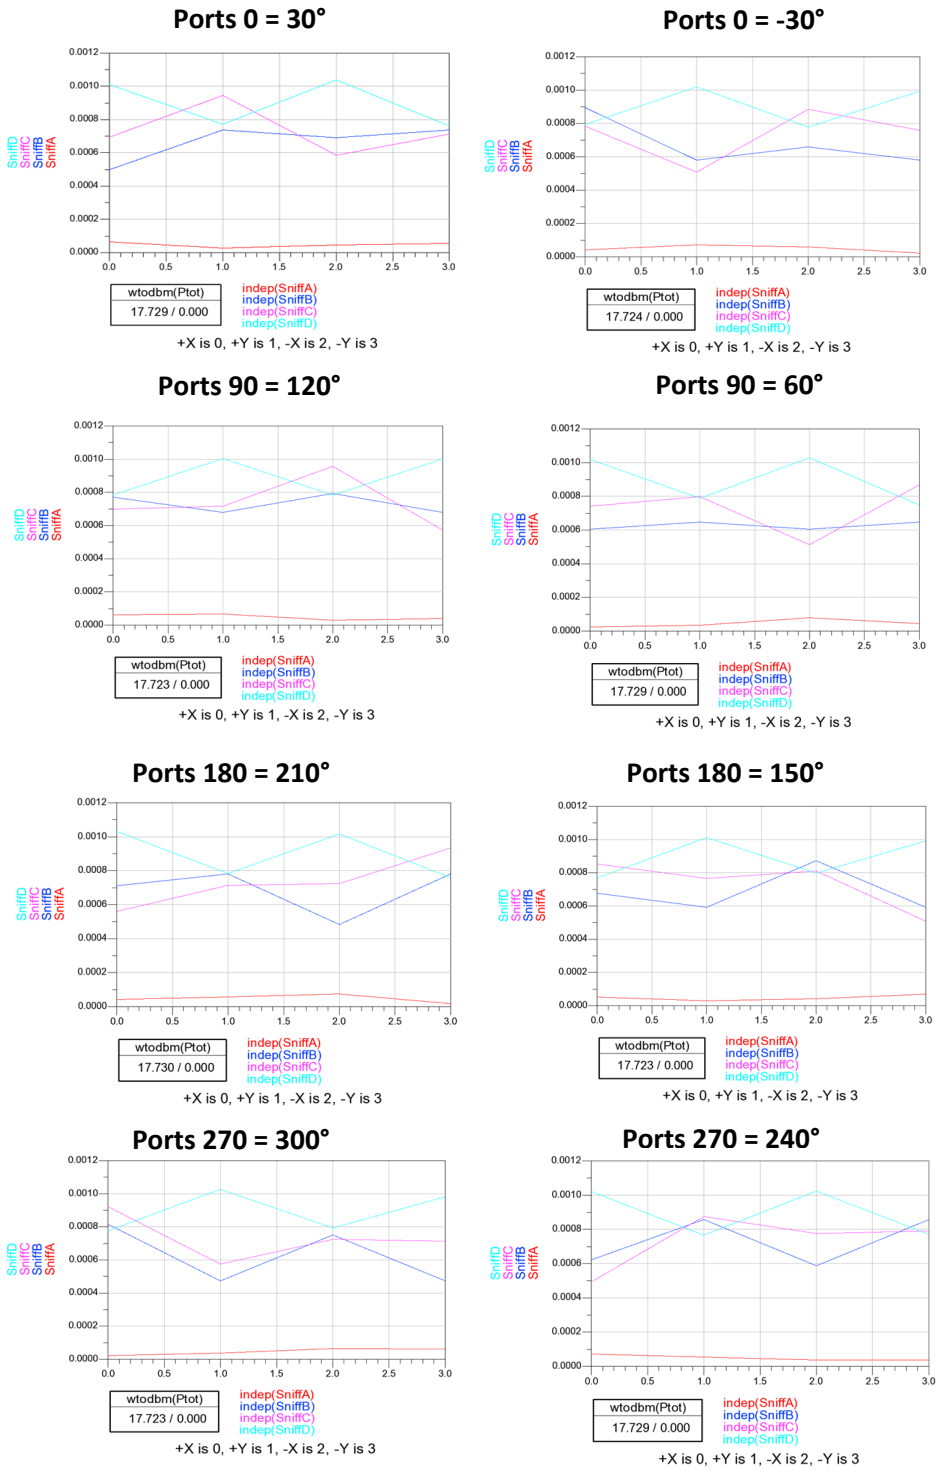

Figure 6.9: Sensor outputs when a single port of each radiator is 30° ahead or behind in phase.

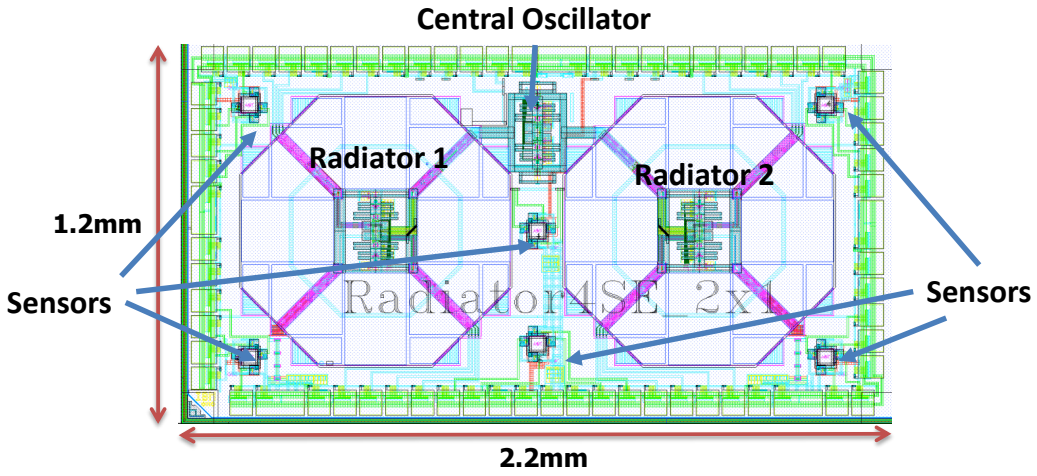

Figure 6.10: Topcell layout of the 2x1 MPD radiator array with self-healing substrate mode pickup sensors.

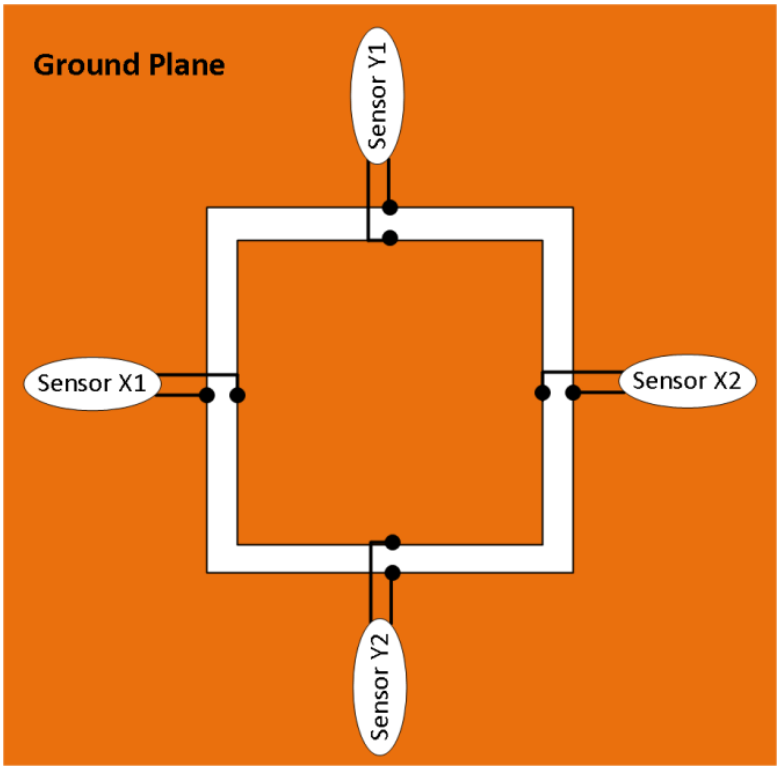

Figure 6.11: Diagram of an individual slot ring substrate mode pickup sensor, with four sense ports, two for each polarization.

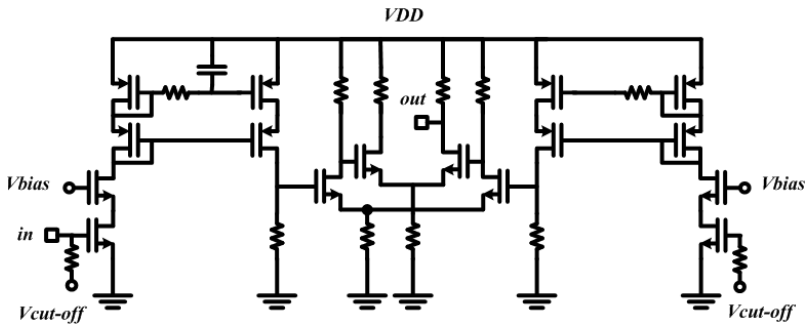

Figure 6.12: Schematic of the sensor rectifying circuitry.

the sensors are picking up substrate modes, as they will still vary up to 15 mV. However, due to the large differences in the absolute magnitude of the sensor voltage from the expected nominal amount, a mapping was not able to be constructed from the sensors to the output radiation for this version of the chip. Rectifying this issue in future designs and constructing the mapping of the mode pickup sensors to the radiation pattern is an ongoing research goal.

6.3 Self-Healing a 2x2 Radiator Array[2]

A 2x2 array of radial MPD radiators at 120 GHz with integrated self-healing has been designed and is currently being fabricated. A brief summary of the progress on this project is presented below.

A block diagram for this self-healing system is shown in Figure 6.13. Each individual radiator is the same as in the 2x1 radiator described in the previous section. The schematics of the radiator core are also the same, but the sizes of all of the transistors have around doubled to produce higher output power, and the matching networks have been retuned. The phase interpolators and amplitude biasing are controlled locally within the radiator core by on-chip DACs that are controlled by an integrated micro-controller. The sensors employed for this system are DC current sensors to detect the DC power consumption, and mode pickup sensors, that sense the substrate modes and are then mapped to the far field radiation patterns. The updated mode pickup sensors are capable of detecting both phase and amplitude, enabling more accurate mapping of the substrate modes to the radiated field than the amplitude sensors of the first design were able to do.

In terms of radiator performance, the gain of the 2x2 array has a maximum gain of 1.8 dBi which when combined with the output power of the amplifiers of 20 dBm, leads to a simulated effective isotropic radiated power (EIRP) of 21.8 dBm. This version of the chip still has full phase control of the radiator oscillators, so

[2]Contributions to this project were made by Amirreza Safaripour, Kaushik Dasgupta, Alex Pai, and myself. I was in charge of the antennas and radiation, as well as the locking network distribution. Amirreza programed the digital controller, did the digital routing, the locking oscillator, the circuitry within each radiator, the mode pickup sensors and the healing algorithm. Kaushik designed the digital to analog converters (DACs) and analog to digital converters (ADCs), and Alex designed the DC sensor.

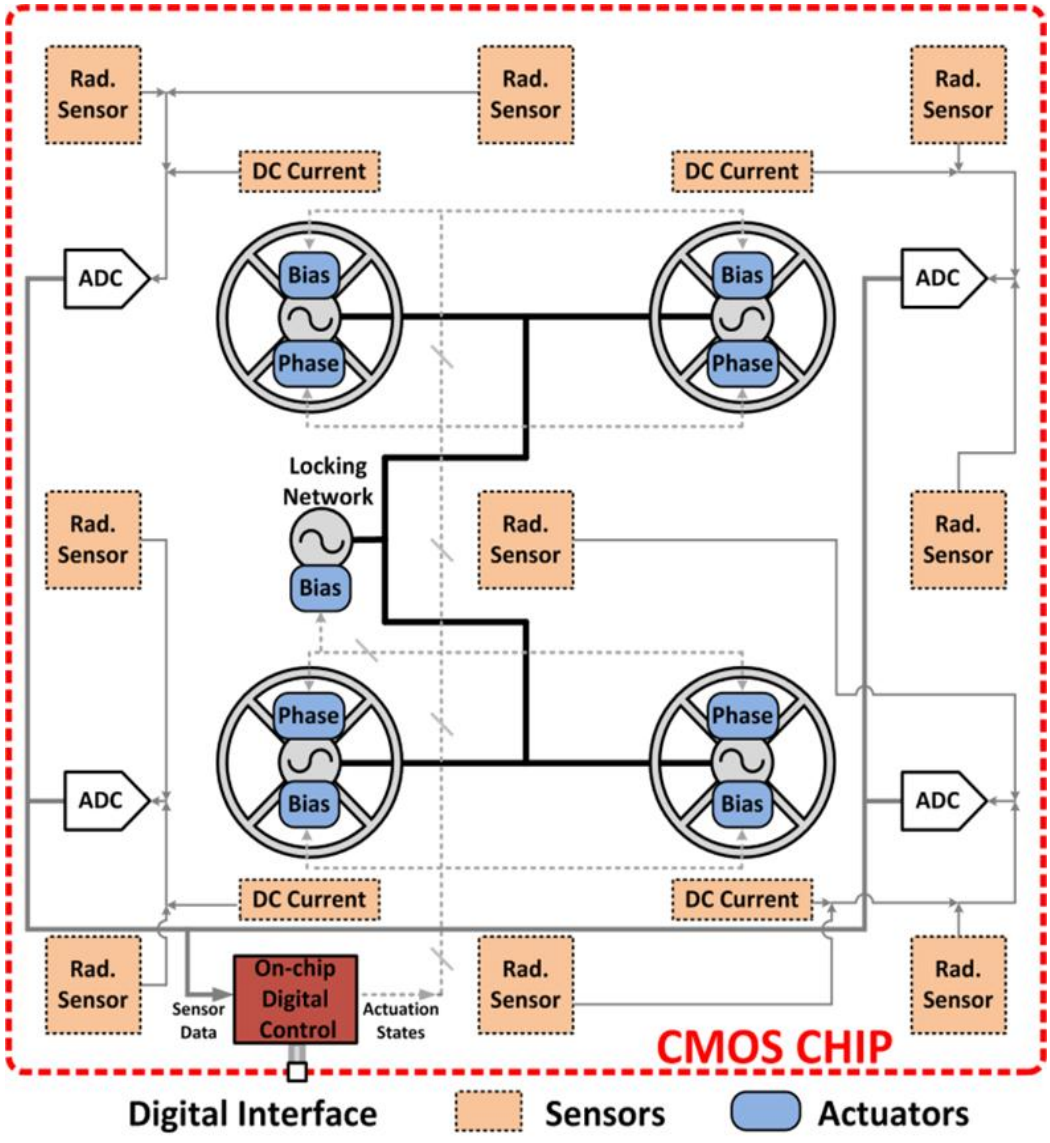

Figure 6.13: Block diagram of the fully integrated self-healing 2x2 MPD radiator array, with substrate mode pickup sensors and DC current sensors, as well as phase actuation for all four radiators and DC operating point bias actuators on the four radiators and the locking network, all controlled by an integrated micro-controller.

it is also able to do dynamic polarization control as well as beam steering.

Chapter 7

Self-Healing Power Amplifier

7.1 Motivation for Self-Healing

Continual advances in integrated circuit (IC) fabrication have opened up numerous new applications and design possibilities for millimeter-wave (mm-wave) systems that previously were not possible and/or not economically feasible [4, 5, 25, 26]. In addition, improvements in power generation in silicon processes have made silicon power amplifiers (PA's) viable [27–32, 93]. Making the PA's in a silicon process allows for greater integration with the rest of the transceiver and reduces the cost. These advances also come with new challenges because with every reduction in minimum feature size, as the industry moves to smaller and smaller process nodes, variation between ICs as well as between transistors within a single IC continues to increase [39–45]. This is compounded by the fact that the digital processing market is the primary driving force behind this scaling, leading to foundries optimizing their device models mainly for digital use. Thus models that are reliable at mm-wave frequencies are often not available early in the node's development stage. These transistor variations and model inaccuracies, as well as other sources of performance degradation, such as environmental variations, critically impact high power mm-wave designs in nanometer scale CMOS technologies and reduce the yields of such designs. Environmental variations can be caused by many things, but one of the important issues for PA's is antenna load impedance mismatch [47]. This occurs when the environment interacts in the near field of the antenna and changes the load impedance looking into the antenna. The ability of PA's to function properly under load impedance mismatch conditions becomes an especially critical issue when they are being used in a phased array [4, 5, 49–52]. In a phased array, the load impedance that needs to be driven by a PA changes when signal from other nearby elements of the array are coupled back through the antenna. Due to the fact that a mm-wave PA is tuned to provide the optimal impedance to the driving stage, any change in load impedance will be away from that optimal and will degrade the PA's performance [48]. Other sources of degradation include degradation due to aging [46, 94] and temperature variation.

This chapter will present self-healing as a method to reduce the adverse affects of process and environmental variation for mm-wave power amplifier (PA) design. Along the way, the design and measurement of

a proof of concept 28 GHz self-healing power amplifier from [53] will be used as an example to explain the various self-healing concepts[1].

Section 7.2 gives an introduction to self-healing and other reconfigurable circuit techniques and presents the design goals and architecture of the example PA, followed by an examination of some of the ways these circuits can be actuated in Section 7.3. A brief overview of the sensors, data converters, and algorithm is then presented in Section 7.4. Finally, a case study of system level measurements of the example PA are then presented in Section 7.5, with concluding remarks in Section 7.6.

7.2 Introduction to Self-Healing

There are two different approaches to solve the issue of performance degradation due to variation. The first approach is to design more and more variation tolerant systems by adopting architectures and circuit topologies which are less sensitive to process and mismatch variations. These techniques have been widely adopted in CMOS digital as well as analog designs over the years. Some of the commonly used strategies include supply voltage optimization [95], optimum device sizing [96], process, voltage and temperature (PVT) insensitive biasing methods [97], etc. The main limitation of the first approach is that with increasing variability in nanometer CMOS, these techniques become harder to implement. In addition, for high frequency designs, a lot of them cannot be implemented due to severe degradation in circuit performance.

The second, more scalable approach is to sense the performance degradation once it occurs and then adjust the system performance by using various knobs. The ability to dynamically sense and actuate critical blocks of the system eliminates the additional design complexity of variation insensitive circuits. Self-healing is a design methodology that takes advantage of the vast digital processing power that is available on modern CMOS processes to reduce the effects of process and environmental variation by using a feedback loop through a digital processing core to heal the chip back to its optimum performance levels when facing performance degradation caused by these variations shown in Figure 7.1. The self-healing loop starts with integrated sensors that detect the performance of the mm-wave circuit. These sensors need careful consideration as they are being implemented on the same chip as the mm-wave circuit, and thus are subject to the same variations. They must be designed to be robust to these variations, so that their outputs are a true measure of the mm-wave circuit's performance, and not dependent on the variation within the sensor. The sensor's outputs are converted to digital bits with an analog to digital converter (ADC) that sends that data to an integrated digital core. This core takes that data, runs an optimization on it, and controls digital to analog converters that set actuation points within the mm-wave circuit. The actuation space of the actuators has to be sufficiently large to cover all of the expected variation the chip may experience, while inducing a minimal

[1] This 28 GHz self-healing power amplifier was a joint project with Kaushik Sengupta and Kaushik Dasgupta. I was responsible for the PA design, actuator design, system integration and system measurements. Kaushik Sengupta was responsible for the sensor designs and the digital VHDL coding. Kaushik Dasgupta was responsible for the data converters and assisted in the system measurements. Aydin Babakhani and Arthur Chang also provided valuable insight during initial discussions on architecture.

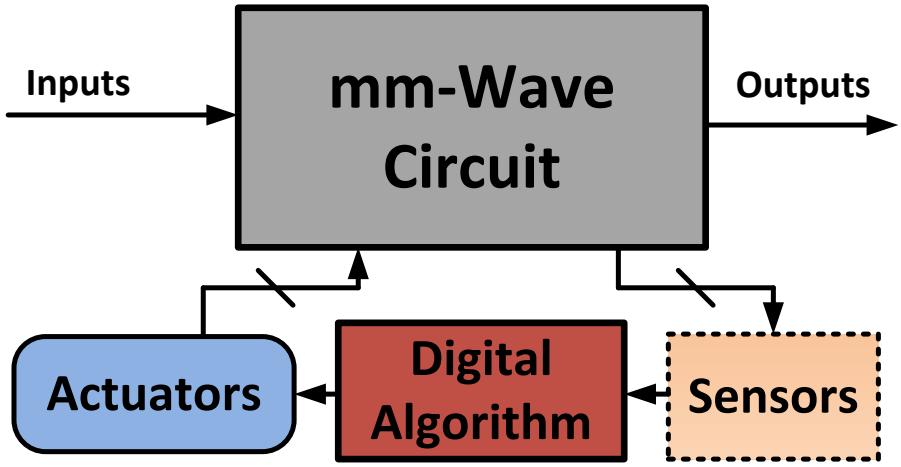

Figure 7.1: Block diagram of generic self-healing system for mm-wave circuits.

---

amount of performance loss due to their presence. This loop can then be iterated until an optimum actuation state is found.

This methodology is more amenable to digital design since for high performance analog design, any sensor or actuator directly affects system performance. [98] implements a system which utilizes several adaptive processing units in addition to regular processing units and periodically turns off the system and performs a functional check and replaces the regular units with their adaptive counterparts as necessary. A reconfigurable test scheme is reported in [99], where performance of a low-noise amplifier (LNA) is detected by reconfiguring it in a feedback configuration and subsequently healed by bias control as well as passive network switching. A healable mm-wave power amplifier is reported in [55–57] wherein based on output power levels, the bias of the core amplifier is dynamically adjusted to provide constant gain. An indirect method of sensing phase-noise is reported in [100] which utilizes correlation between phase noise and integrated sensor outputs to optimize the system for best phase noise. On-chip PVT compensation has been demonstrated in a 2.4 GHz LNA [101] where the input of the LNA switches between off-chip and an on-chip voltage controlled oscillator (VCO) during measurement and calibration phases respectively. A closed-loop method for healing phase-locked loop (PLL) reference spurs is implemented in [54] where both spur estimation and correction are performed at the VCO control voltage. The work presented in [102] demonstrates integrated phase error detection and self-healing for use in phased array based receivers and transmitters.

7.2.1 Self-Healing Blocks

- **Sensors:** The most vital aspect of sensor design for self-healing circuits is robustness. The mm-wave circuit is designed to be cutting edge, to push the envelope of possible performance, and that means that variations can significantly degrade performance. The sensors, on the other hand, can be designed more conservatively, with robust design topologies and other techniques such as using non-minimum length transistors [42]. Also, because the self-healing loop can be duty cycled, the DC power requirements are more relaxed than in the mm-wave circuit, robust designs that are more power hungry can still be acceptable in some circumstances. In order to design the sensors, the performance metrics of importance must be determined. It is also important to determine if the absolute value of the metric is important, or if it is a metric that must be maximized or minimized. In the case of mm-wave power amplifiers, the output radio frequency (RF) power is a metric that often is bounded, but, on the other hand, efficiency is one that should always be maximized. For metrics that are not bounded, the most important aspect of the sensor is that it is monotonic, so that the optimizer does not get stuck on an artificial local maximum or minimum. The exact value reported in this case does not matter as much as the optimization will always try to maximize or minimize the metric. For bounded metrics, the sensors need to be much more robust, as the optimizer will attempt to set the chip to a specific operating point that is based upon the absolute value that the sensor is reporting.

- **Actuators:** To enable the mm-wave circuit to adapt to process and environmental variations, actuators need to be built into the design. Ideally, these actuators will not affect the performance of the mm-wave circuit, but in practice there will be some cost associated with the actuator, in the form of performance degradation and/or power and area consumption. This means that the actuators must be placed sparingly and be designed with the expected variation in mind. In the case of power amplifiers, the expected variation from process variation comes in the form of threshold voltage variation and parasitic capacitance variation, while the environmental variation can include temperature, aging and load mismatch. Thus, the aspects of the design that should be controlled by the actuator are the operating points of the amplifying transistors and the matching networks.

  Possible types of actuators can be separated into four broad categories, that cover most of the parameters that a designer has control over when initially designing the amplifier. The first is control of gate bias voltages. By controlling the DC voltage on the gates of the amplifying transistors, the operating point of the transistor can be controlled. This will affect DC current, transconductance, f_{max} of the transistor, gain, RF saturation power, linearity, parasitic capacitances and stability. In the case of cascode or stacked amplifiers, a combination of these parameters can be controlled by independently controlling the various gate bias points. The second type of actuator is tuning of the passive power combining and matching elements. Tuning of the matching elements can ensure that the input and output impedances seen by the amplifier are optimal, and in the case of power combining PAs, can adjust

the power combiner in the case of mismatch between the output stages.

The third type of actuator is to actuate the supply voltage. This can enable higher efficiency in low power back off, while still enabling a high power mode with a higher supply voltage, but requires an efficient tunable DC-DC converter.

The final type of actuator involves changing the sizes of the transistors themselves. This can be done by placing many transistors in parallel and switching them in and out of the circuit. This again can help in efficiency in low power back off by reducing the size of the amplifying transistors, but the impedance mismatch and the loss associated with the switches can make this type of actuator cost prohibitive. In the example design of the self-healing power amplifier, gate bias control and matching network tuning are employed as actuators.

- **Data Converters:** In order to bridge the analog and digital domains, data converters must be implemented on chip. The sensor data must pass through an analog to digital converter (ADC) to be converted to digital signals that can be read by the digital algorithm. These ADCs must be robust, as the digital algorithm only has access to the ADC's digital output bits, so the aggregate variation of the sensors combined with the ADC determines the accuracy of the performance metric measurement. The analog actuators also need digital to analog converters (DACs) to enable the digital core to control analog voltages. To keep overhead down, low power current mode DACs are used to control the analog actuators. Some actuators such as the tunable transmission lines, however, directly take digital inputs and thus do not require the use of DACs.

- **Digital algorithm:** In order to reach the optimum actuation state, the digital algorithm core needs to first determine the performance of the mm-wave circuit from the sensors, and then decide how to adjust the actuators to improve that performance. There are several components to the digital core that can be separated and implemented independently. There will be components for sensor reading, actuator writing, optimization, and global control.

  The global control component will first tell the sensor reading component to deliver the sensor data by querying the ADCs to return the state of all of the sensors. The global control component then sends this information to the optimization component. The optimization component calculates the performance from these sensor readings and determines the next state to set the actuators to. This is where the algorithm itself is implemented. The global control component then delivers this actuation state information to the actuator writing component, which sends out the bits to the DACs (for the analog actuators) or directly to the digitally controlled actuators. By grouping the digital core into these separate module components, more of the code can be reused between different designs, and changes can be made to update the optimization component's algorithm without recoding the other supporting components.

7.2.2 Block Level Versus Global Healing

It is important also to consider block level versus global healing. While the scope of this chapter just deals with block level healing of just the power amplifier block, global healing of the entire transmitter can be useful as well. For example, if the transmitter is operating in a low power back off mode, it may be advantageous to adjust the gains of the preamplifiers and mixers, not just the bias points of the power amplifier to lower the overall DC power consumption and improve efficiency, or to tradeoff between linearity and power consumption. The entire healing digital core for the transmitter can be contained in a single centralized core, or it can be broken up into smaller block level cores that are controlled by a transmitter level global healing core. By incorporating self-healing both at the block level and transmitter level, maximum performance can be achieved.

7.2.3 Design Considerations and Architecture for an Example Self-Healing Power Amplifier

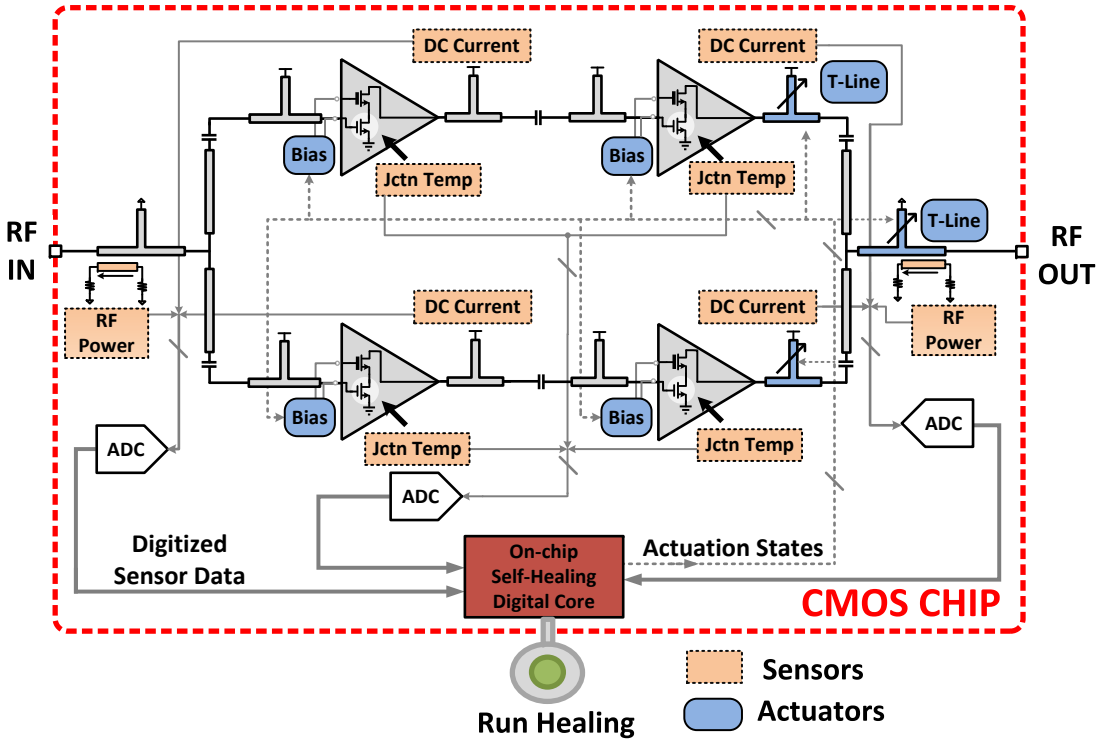

Figure 7.2: Block level architecture of the example integrated self-healing PA. Data from three types of sensors is fed through analog to digital converters (ADCs) to an integrated digital core. During self-healing, the digital core closes the self-healing loop by setting two different types of actuators to improve the performance of the power amplifier.

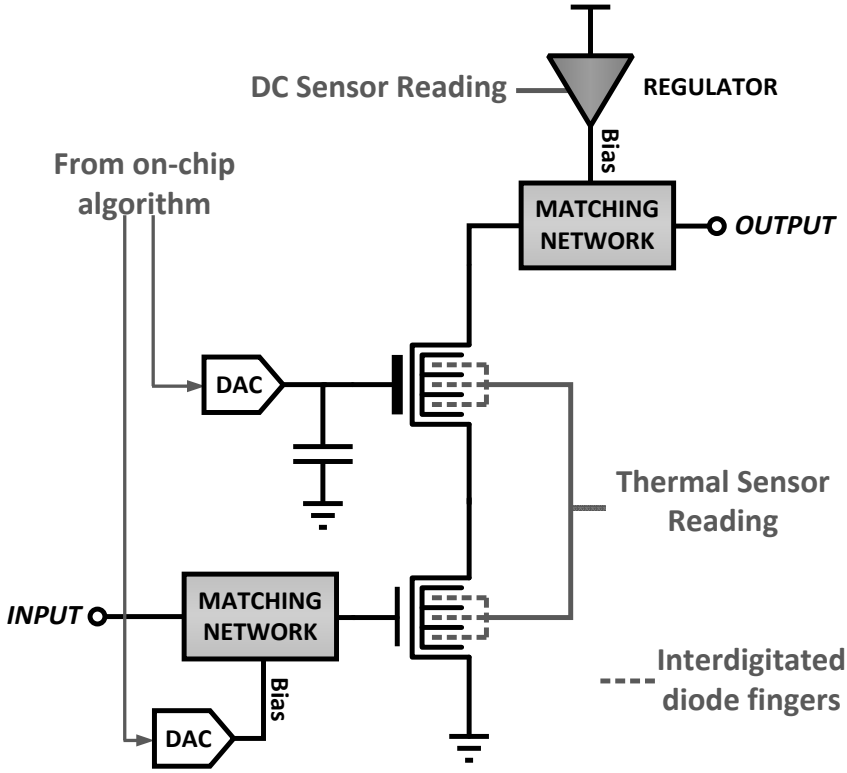

Figure 7.3: Schematic of a single cascode amplifying stage showing connections to matching networks, gate bias actuators, DC sensor and temperature sensor.

The PA will be a fully integrated self-healing PA at 28 GHz implemented in a standard 45 nm SOI CMOS process (Fig. 7.2) [53]. It is a 2-stage, 2-to-1 power combining class AB PA [103]. The input dividing network is matched to 50 Ω, and the interstage and combining network are designed to present the optimal impedance to each stage for maximum power transfer at saturation, when the interstage networks are loaded with the gates of the output stages, and the output combiner is loaded with 50 Ω. The first stage is half the transistor size of the output stage, to ensure that the output stage can be fully driven into saturation. Class AB design was chosen in order to enable linear operation and to allow for non-constant envelope modulation schemes to be implemented.

To increase the gain of each amplifying stage, each stage is a cascode amplifier, and two stages are used to further increase the gain. The common source transistors are 56 nm analog transistors, while the cascode transistor is a 112 nm thick gate oxide transistor to increase the voltage breakdown of the amplifier. A schematic of one of the output amplifying stages is shown in Figure 7.3, showing the connections through the matching networks, as well as several of the sensors and actuators that will be discussed in the following sections. The three matching networks use a 2-stub matching technique, and biasing is done through the AC

short circuits at the end of the stubs. A metal AC coupling capacitor is used to allow for independent biasing of the inputs and outputs of the amplifying stages. Full 3-dimensional electromagnetic simulations of the matching networks including the capacitors and pads were performed to ensure proper functionality.

There are three main metrics of interest that will define the performance of the example mm-wave power amplifier: output power, power gain, and DC power. This means that to be able to calculate the performance of the PA, the self-healing sensors need to be able to detect input and output RF power, and DC power. RF power sensors are thus placed at the input and output ports of the amplifier, and DC power is sensed both electronically through DC sensors as well as thermally through temperature sensors. On the actuator side, gate bias actuators are implemented on all amplifying stages, and the stubs of the output power combining matching network are tunable to enable tuning of the output network. These sensors and actuators will be discussed in more detail in subsequent sections.

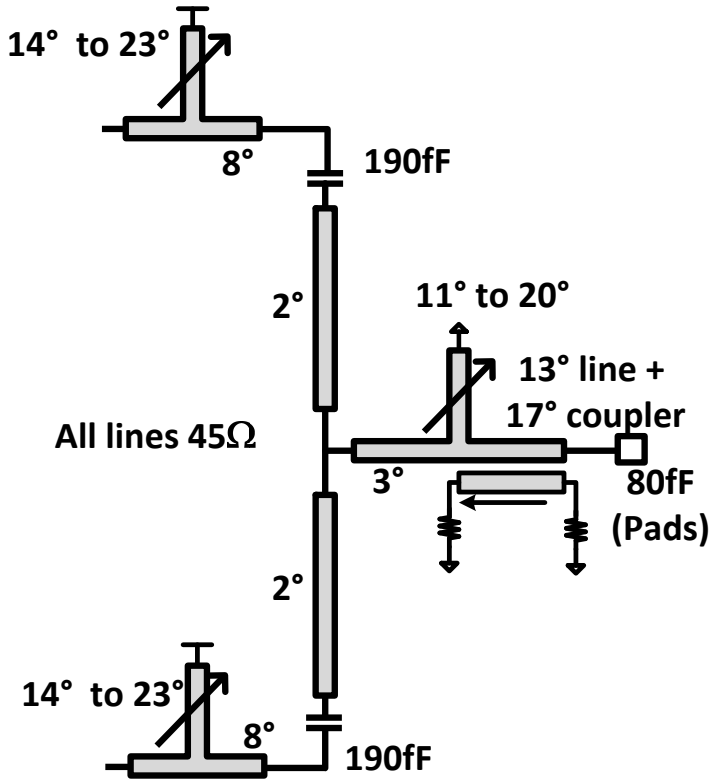

Figure 7.4: Schematic of the output power combining matching network.

An impedance mapping from the output load of 50 Ω to the optimal impedance of the output stage with typical transistors of (7+7j) Ω showing each matching element is depicted on the Smith chart in Figure 7.5. Because the output matching network is tunable, the two inputs to the power combiner were not isolated from each other, providing savings in area as well as loss for the network due to the shorter transmission lines. To design the matching network, the two output stages were assumed to be identical, and thus the impedance

seen looking into the T-junction from either stage toward the load is twice the impedance seen looking toward the load just after the T-junction [29]. If there is a mismatch between the two output stages, the self-healing algorithm will attempt to correct for that by adjusting the DC operating points of the stages or by tuning the output matching network.

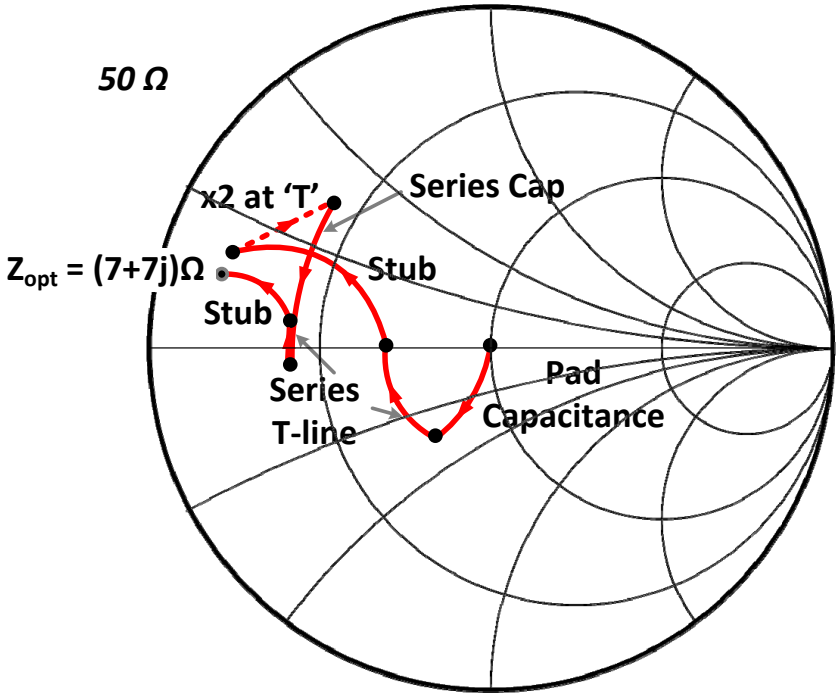

Figure 7.5: Mapping impedance transformation throughout the output matching network.

7.3 Actuation: Countering Performance Degradation

In order for the self-healing system to be able to improve the performance of the mm-wave circuit, actuators that are controlled by the digital algorithm are required. It is important for the actuation space to cover the expected space of variation, so that the optimum point for any expected variation is still within the actuation space.

The expected process variations that are considered include chip-to-chip process variation, and mismatch between devices on the same chip. The main sources of this variation are variation in dopant levels and variation in gate thicknesses that affect threshold voltage and parasitic capacitance. Within the context of power amplifier design, these types of variations will change the operating points of the transistors, can change the maximum gain and f_{max} of the transistors, and will affect the desired input and output impedances for maximum power generation. The environmental variations that will be considered include load impedance mismatch, temperature variation and transistor degradation and failure due to effects such as aging. Load

impedance mismatch will affect the output matching network, and will detune the system, lowering the output RF power supplied by the output amplifying transistors, as well as possibly increasing the loss due to the output power combiner matching network. Temperature variation will mainly show up in transistor performance, and will affect the matching, bias and saturated power of the amplifying transistors. To a lesser extent, it will affect the passives, and will change the matching networks, but this effect is much less dominant. Transistor degradation due to aging will reduce the gain and saturated output power of the transistors, and will affect the output matching, to the point where complete failure of some of the output transistors can severely degrade the matching network's ability to provide the optimum load impedance to the other functional transistors. There will thus be degradation to the output signal due both to less power being generated by the output transistors, as well as degradation due to mismatch in the output power combining matching network.

How then do we select appropriate actuators to cover the effects of all of these variations? A good actuator should cover a large actuation space while having a minimal impact on the performance of the amplifier in the nominal case. Every actuator will cause some degradation, whether it be to decreased RF power, increased DC power, or even occupying area on the chip, but keeping this degradation to a minimum will make the system have a net positive effect on the performance once the healing loop is turned on.

There are several possible actuators for mm-wave PAs that generally fall into four categories: gate bias voltage actuators, passive matching network tuning actuators, supply voltage actuators and transistor architecture actuators.

7.3.1 Gate Bias Actuators

Control of the gate voltages of the transistors is one of the most effective, lowest cost actuators available to the designer. By adjusting the gate of the common source transistors and any stacked or cascode transistors, the operating points of the transistors can be changed. These actuators can be very low cost, due to the high DC gate resistance of CMOS transistors, so adjusting the voltage of the gate does not inherently draw any extra DC power. In general, for class AB PA operation, the transistors should be biased such that the f_{max} of the transistor is maximized. This biasing point is heavily dependent on the threshold voltage of the transistors and can vary greatly over process variation. This can enable maximum performance at class AB saturation situations, but control of the gate bias also enables a lowering of the DC power consumption of the power amplifier in power back-off scenarios. There is a tradeoff on gain verses DC power in back-off.

One way to reduce the output power and achieve back-off power levels is by using a variable gain amplifier before the PA and lower the input power of the PA, but because the PA is the dominant source of DC power consumption, this leads to low efficiencies in back-off because the PA is drawing relatively the same amount of DC power as in saturation, but is producing much less RF power. The other way to achieve back-off power levels is to trade some of the gain in the power amplifier for lower DC power consumption. By lowering the gate bias of the common source transistors, the PA will saturate at lower power levels, while

simultaneously reducing the DC power consumption, significantly raising the efficiency of the amplifier in back-off conditions. These actuators require analog inputs and thus will need DACs that will be discussed in the next section to convert the digital signal sent by the algorithm to the analog voltage for the actuator. Gate bias actuators are employed on the example PA for the gates of the common source and cascode transistors of both paths of the input and output amplifying stages. Looking back to Figure 7.3, the gate bias actuators for the common source transistor are input through a stub in the matching network, while the gate bias control for the cascode stage is biased at the AC short circuit created by the AC coupling capacitor.

7.3.2 Passive Matching Network Tuning Actuators

While controlling the DC operating points of the amplifying stages can offset much of the process variation, overcoming changes to the reactive parasitics as well as load impedance mismatch requires a change in the matching networks themselves. The second type of actuators are passive matching network tuning actuators. The goal of these actuators is to be able to dynamically retune the matching networks on the chip after fabrication to account for the variations in parasitic capacitances and optimal output loads of the amplifying transistors. Actuators can be added to any type of passive element, and as an example, two types of actuators to create tunable transmission lines for transmission line matching networks will be considered in more depth.

For the output power combining matching network, keeping the loss to a minimum is critical, because the output network deals with the highest power levels, and since the loss will show up in the form of attenuation in dBs, this means that loss in the output network will result in the largest actual loss in terms of watts.

Also, for maximum performance, the amplifying transistors will be operating near their voltage breakdown and the matching network will transform the impedance to 50 Ω from a much smaller impedance, which means that the voltage swing within the output network and at the output will be much larger than the transistor breakdown, and any actuator transistors must not be placed such that they breakdown. On a 50 Ω load, a peak power delivery of 16.5 dBm creates a peak to peak voltage swing of more than 6 V, so any actuation transistors must be carefully placed in such a way that they do not experience breakdown when the PA is supplying maximum power. Finally under the above two constraints, the actuators still need to cover a large enough actuation space to cover the entire region of expected variation in the circuit as well as externally.

One way to actuate the matching networks is to change the distributed inductance (L) and or capacitance (C) of the transmission lines. This will change both the intrinsic impedance (Z_0) of the line as well as the effective relative dielectric constant (ϵ_e) due to the relationships

$$Z_0 = \sqrt{\frac{L}{C}}, \tag{7.1}$$

and

$$\epsilon_e = \frac{LC}{\epsilon_0 \mu} \tag{7.2}$$

where ϵ_0 is the free space dielectric constant of $8.85 \times 10^{-12} F/m$ and μ is the permeability of the material

and for the dielectrics within the scope of this chapter employed on standard integrated circuit (IC) processes, is equal to the permeability of free space of $4\pi \times 10^{-7} \frac{Vs}{Am}$.

The inductance can be controlled by modifying the path of the return current, while the capacitance can be controlled by switching in many small lumped capacitances into the transmission line itself without affecting the main path of the return current very much. These concepts were investigated, and the results are presented in Appendix C. In the end, the additional losses created by these types of tunable transmission lines outweighed the benefit that the tuning provided, and they were not used in the final design of the self-healing power amplifier.

The other way to tune the passive networks is to change the length of the transmission lines. While it would be significantly difficult to change the length of the series lines, it is possible to actuate the length of the transmission line stubs. However, using transistors as switches on various parts of the matching network also can lead to tradeoffs in transistor size. To have low-loss switches in the on state, the size of the transistor is increased, but this increases capacitance in the off state. This can initially lead to a compromised design that both loads the circuit with capacitance in the off state, and causes loss due to resistance in the on state.

One way to overcome both of these challenges is to tune just the transmission line stubs of the matching network that are terminated with short circuits. The effective length of the transmission line can be changed by shorting out the stub at various points near the end of the line, as shown in Figure 7.6.

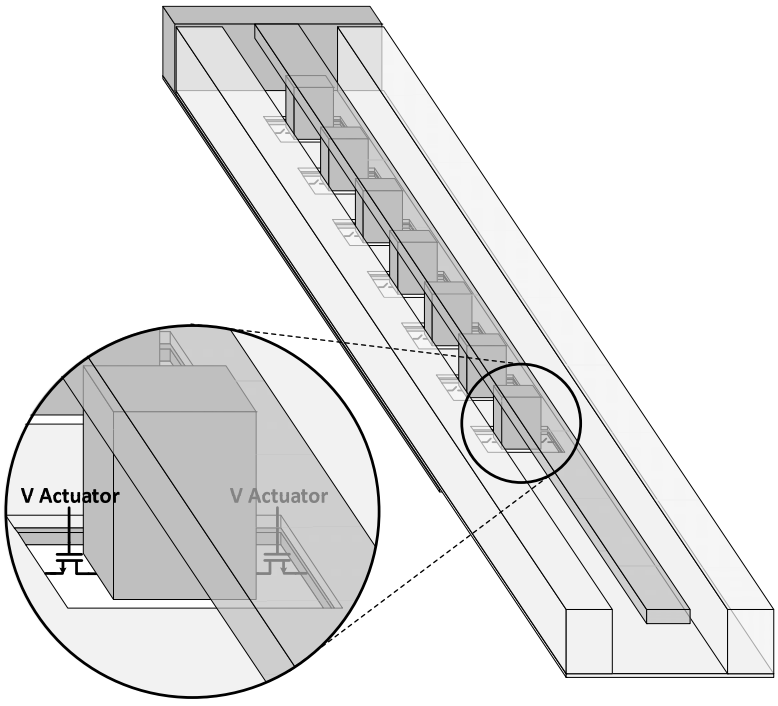

Figure 7.6: Tunable transmission line stub has switches placed at various points along the line to short out the signal line to the ground, changing the effective length of the stub.

The on resistance of the switches are the dominant source of added loss due to the actuators, which means that keeping it very low keeps the performance cost of the actuators low. Turning all of the switches off increases the effective length of the transmission line to a maximum while turning them all on moves to the minimum effective length. Because the voltage swing on the line when all of the switches are off increases the closer it gets to the stub's junction with the rest of the matching network, actuation switches are placed from the shorted side of the stub to about 60% of the way to the junction.

Large actuation ranges can be achieved simply by placing actuation switches toward the end of the line, which due to their proximity to the short circuit at the end of the line, keeps the voltage swing low when the switches are off. On top of that, because the switches are being placed at regular intervals within the transmission line, the off capacitance can be considered distributed, and can be absorbed into the transmission line model and included in the match. This means that much larger transistors can be used and low on resistances can be achieved. These stubs can be considered as shunt inductors at their connection to the rest of the output network, and simulations of one of these tunable transmission line stubs shown in Figure 7.7 show a tuning range from 25 pH to 71 pH with just 8 switches, showing the large actuation space that these actuators can achieve. Because these actuators are already digital switches, they can be set directly from the algorithm block without the need for any DACs.

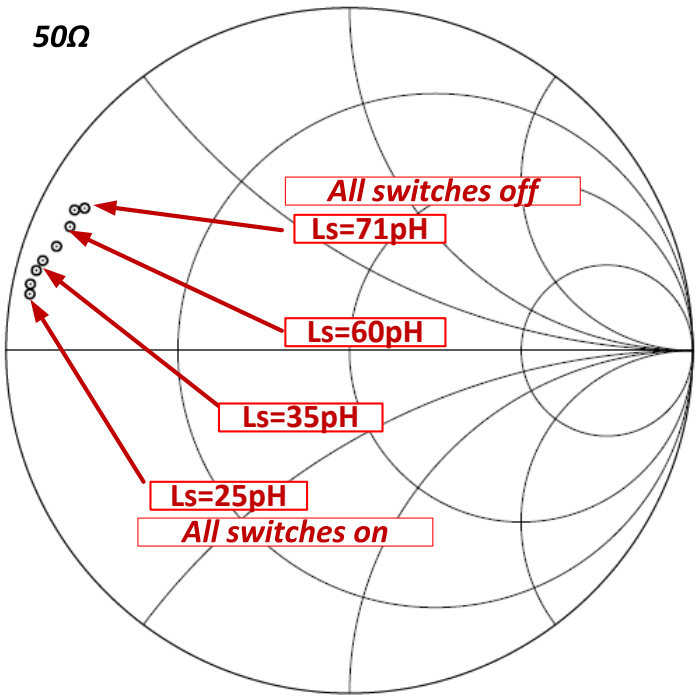

Figure 7.7: Measurement of a single tunable transmission line stub and its effective inductance at each state.

The self-healing PA utilizes these tunable transmission line stubs on all three of the output power combin-

ing matching network stubs. They were not used on the input matching network or the interstage matching networks because the expected variation in impedances of the loads for those networks (gates of the first and second amplifying stages) was small enough that the advantage of the matching networks did not warrant the additional complexity and loss the actuators come with. The output power combining matching network is shown in Figure 7.5, with the nominal impedance mapping from the 50 Ω load to the optimal impedance of the output stage with typical transistors, which for this example is $(7 + 7j)$ Ω. One additional benefit of the matching network being tunable is that the two output amplifiers did not need to be isolated. If there is mismatch between them, it can be minimized by tuning the output matching network.

There is another trade-off between the achievable resolution of the effective length of the stub and the performance cost of the actuators, and for the described design, a series of 7 actuation transistors along the line and thus 8 different effective stub lengths produces adequate resolution to cover the actuation space while having minimal impact on the amplifier output power. An ideal shorted stub transmission line that is less than 1/4 of a wavelength can be modeled as a parallel inductor to ground, and thus the actuation of each stub can be looked at as creating a variable inductor. Measurements were taken on a breakout cell of one of the tunable stubs, and show a wide range of tuning from an effective inductance of 25 pH up to 71 pH, as shown in Figure 7.7.

When all three tunable stubs are incorporated into the entire matching network, they enable two-dimensional control of the network, allowing matching from 50 Ω to an actuation space around the typical optimal impedance, matching from a variety of output loads in the case of load impedance mismatch back to the typical optimal impedance, or more practically some combination of both. Figure 7.8 shows the simulated actuation space of impedances looking into the combiner due to the tunable transmission line stubs with a 50 Ω output load. Thus if there are transistor variations that cause the optimal impedance for the amplification stages to move within that space, the actuators will be able to provide that impedance to deliver maximum power. Looking at the inverse case, Figure 7.9 is the simulated actuation space of output load impedances that can still be transformed back to the optimal impedance of (7+7j) Ω of the output stage with typical transistors.

7.3.3 Supply Voltage Actuators

Another way to reduce the DC power consumption of the PA is to reduce the supply voltage to the amplifiers. Simply lowering the voltage on-chip by dumping the power into a transistor or resistor does not actually decrease the DC power, so efficient tunable DC-DC conversion is required for this type of actuator to be effective. Since these types of converters fall out of the scope of this work, this type of actuator is not considered further, but should be kept in the designer's mind, especially if such a block is already being implemented on-chip for other reasons and can be taken advantage of by the self-healing system.

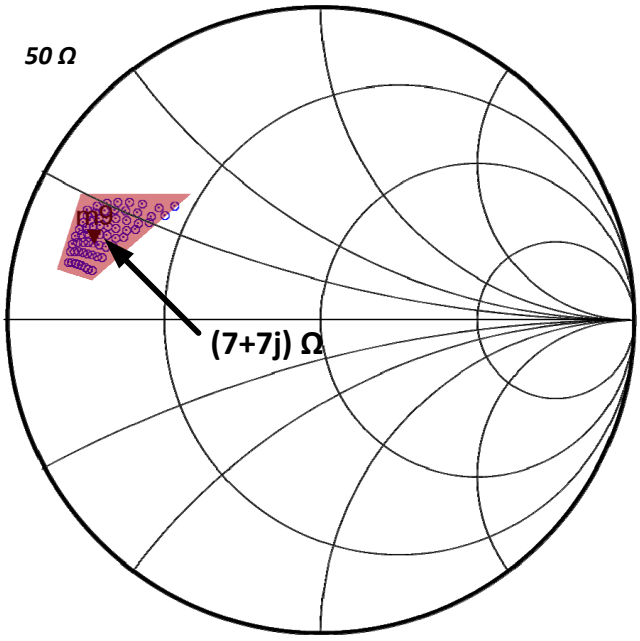

Figure 7.8: Actuation space of impedances that can be matched to a load impedance of 50 Ω.

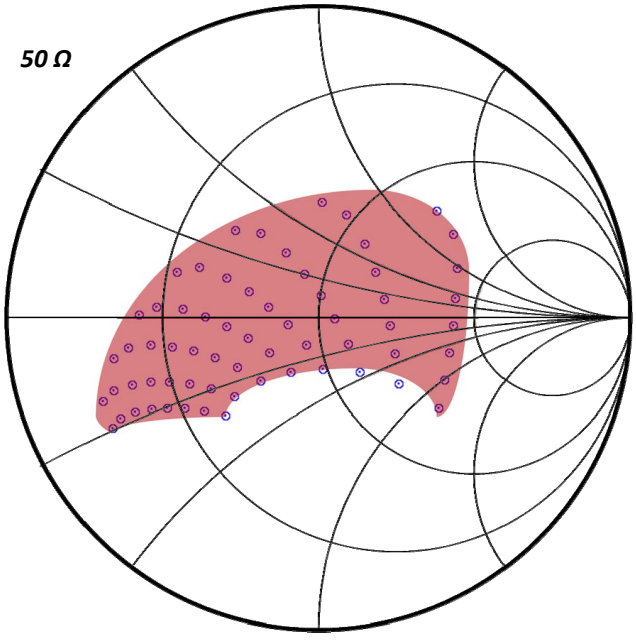

Figure 7.9: Actuation space of load impedances that can be transformed to the typical optimal impedance of (7+7j) Ω.

7.3.4 Transistor Architecture Actuators

Changing the size or architecture of the transistors themselves is also a possibility for actuators for mm-wave PAs. One can imagine that an easy way to deal with saturated output RF power degradation is simply to increase the size of the amplifying transistors. Switching fingers of transistors in and out of the circuit, however, can cause significant losses in performance from two dominant sources. The first is that the switches used to switch the transistors in or out will either have higher on resistance or off capacitance, and both will cause signal degradation. Secondly, even if perfect switches were available, changing the size of the transistors will significantly change the desired matching network, and put much more demand on the tunable matching networks, resulting either in more loss in the matching networks or in mismatch for some transistor sizes. For the example PA, it was determined that the best performance could be achieved by making the transistors large enough to cover the desired saturated output power and to use the gate bias actuators to achieve higher efficiencies while in back off, so transistor architecture actuators were not used.

7.4 Sensors, Data Converters, and Digital Algorithm

This section will provide a brief overview of the sensors, data converters and digital algorithm blocks to give context to the system level design and measurements to follow.

7.4.1 Sensors

Three types of sensors were implemented for this PA: a RF power sensor, a DC current sensor and a thermal sensor designed to sense the temperature at the transistor junctions. These sensors are presented in depth in [53, 104].

The RF sensor detects true RF power by using a directional coupler to tap off a small amount of power from the transmission line that feeds voltage sensors at both the coupled and isolated ports. Sensing both ports is critical as it enables true power sensing, regardless of load impedance mismatch, as the forward traveling wave shows up at the coupled port, and the wave reflected from the load shows up on the isolated port. The voltage sensors are implemented as transistor rectifiers followed by two stages of current amplification through current mirrors [105–109]. RF sensors are then placed at both the input and the output of the PA, with the coupling coefficients of the directional couplers scaled for each based upon expected power at the input and output of the PA.

The DC current sensor detects the DC current going to the amplifying transistors. It was incorporated into a voltage regulator that uses two operational amplifiers to set the chip VDD level to a given reference voltage regardless of the supply VDD [61]. Keeping the required headroom to a minimum was critical, and headrooms of 10-30 mV were achieved. The op-amps were then repurposed to also mirror 1% of the current flowing through the main path through an auxiliary path that was converted to a voltage through a resistance.

The thermal sensor operates by interdigitating diodes within the amplifying transistors themselves, and comparing those diodes to 'cold' diodes placed around 100 μm away on the chip. This gives a very localized temperature reading that corresponds to the amount of head dissipated in each transistor, a good proxy for the amount of DC power burned in each transistor. It is a good example of sensing performance metrics in other domains, in this case sensing electronic power in the thermal domain. The main drawback of these sensors is that they can only operate as fast as the transistors can heat up, which is determined by thermal time constants that are prohibitively long. Thus, the thermal sensors were not used directly during the healing process, and were instead used as a verification of the DC current sensors after the healing had taken place.

An overview of the three sensors is shown in Figure 7.10, and a performance summary is shown in Table 7.1

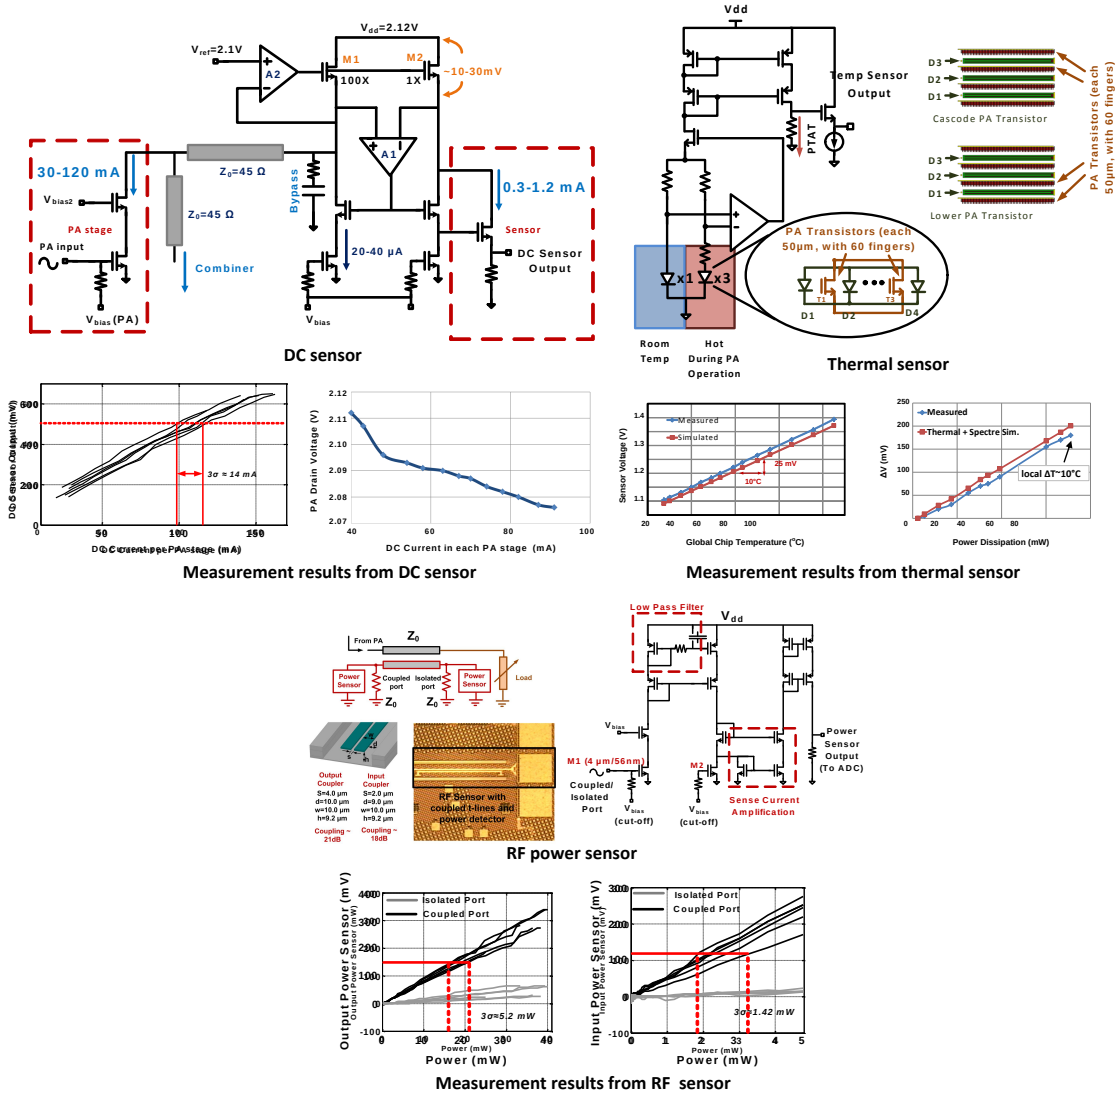

Figure 7.10: Summary of self-healing sensors, with schematics and measurements of the DC current sensor (upper left), the thermal sensor (upper right) and the RF power sensor (bottom)

Table 7.1: Performance Summary of on-chip sensors implemented for the self-healing PA.

Sensors	Measured Entities	Responsivity	Range	Sensor 1-bit resol.
True RF Power	In. Power	54 mV/mW	0-10 mW	55 μW
	Op. Power	8.3 mV/mW	01-100 mW	300 μW
DC Sensor	DC drawn by			
	Ip. Stage	8.5 mV/mA	0-60 mA	280 μA
	Op. Stage	4.2 mV/mA	0-120 mA	560 μA
Thermal Sensor	Power dissipated			
	Ip. Stage	4.0 mV/mA	0-130 mW	0.75 mW
	Op. Stage	2.0 mV/mA	0-260 mW	1.5 mW

7.4.2 Data Converters

Fully integrated self-healing systems require both integrated DACs and integrated ADCs. Though there are many architectures that can be used for such a DAC [110, 111], in this system, two 6-bit DACs are implemented using binary weighted switchable transistor-based current sources through a load resistor. One is used to bias the cascode transistor's gate of the amplifiers and has a range from 1.1 V to 1.95 V. The other is used to bias the common source transistor's gate and has a range of 450 mV to 1.05 V, and includes an extra bit to set the voltage to 0 V to completely turn off the amplifying stage.

The ADCs are implemented as 8-bit successive-approximation-register (SAR) ADCs. The data from all 12 of the sensors (4 from RF sensors, 4 from thermal sensors, and 4 from DC sensors) were multiplexed into 3 of the ADCs on chip. The output of the ADCs are serial digital lines fed to the digital core. An R-2R based DAC is used to enable the successive approximation. The block diagrams of the DACs and ADCs are shown in Figure 7.11.

7.4.3 Digital Algorithm

The digital algorithm was coded in VHDL and synthesized on the chip. While this algorithm utilizes a direct mapping of the sensor results to the performance metrics of interest, indirect mapping can also be used [112,113]. Another option would have been to use a micro controller, but higher performance and lower area overhead are required when the widget is coded directly. The algorithm uses a global state machine that controls individual state machines for reading the sensors, loading the actuators, and running the optimization.

Two optimizations were implemented for this chip. The first one simply tries to maximize the output RF power, and does a bulk search of the actuation space of 262,144 states to find it. The second one takes a desired output power as a user defined input, and tries to meet that desired power while consuming the minimum amount of DC power. It searches the actuation space by starting off at the lowest DC power levels, when the biasing of the common source transistors is lowest, and does a bulk search of the cascode biasing and the matching network tuning, and stops once it detects that it has achieved the desired output power.

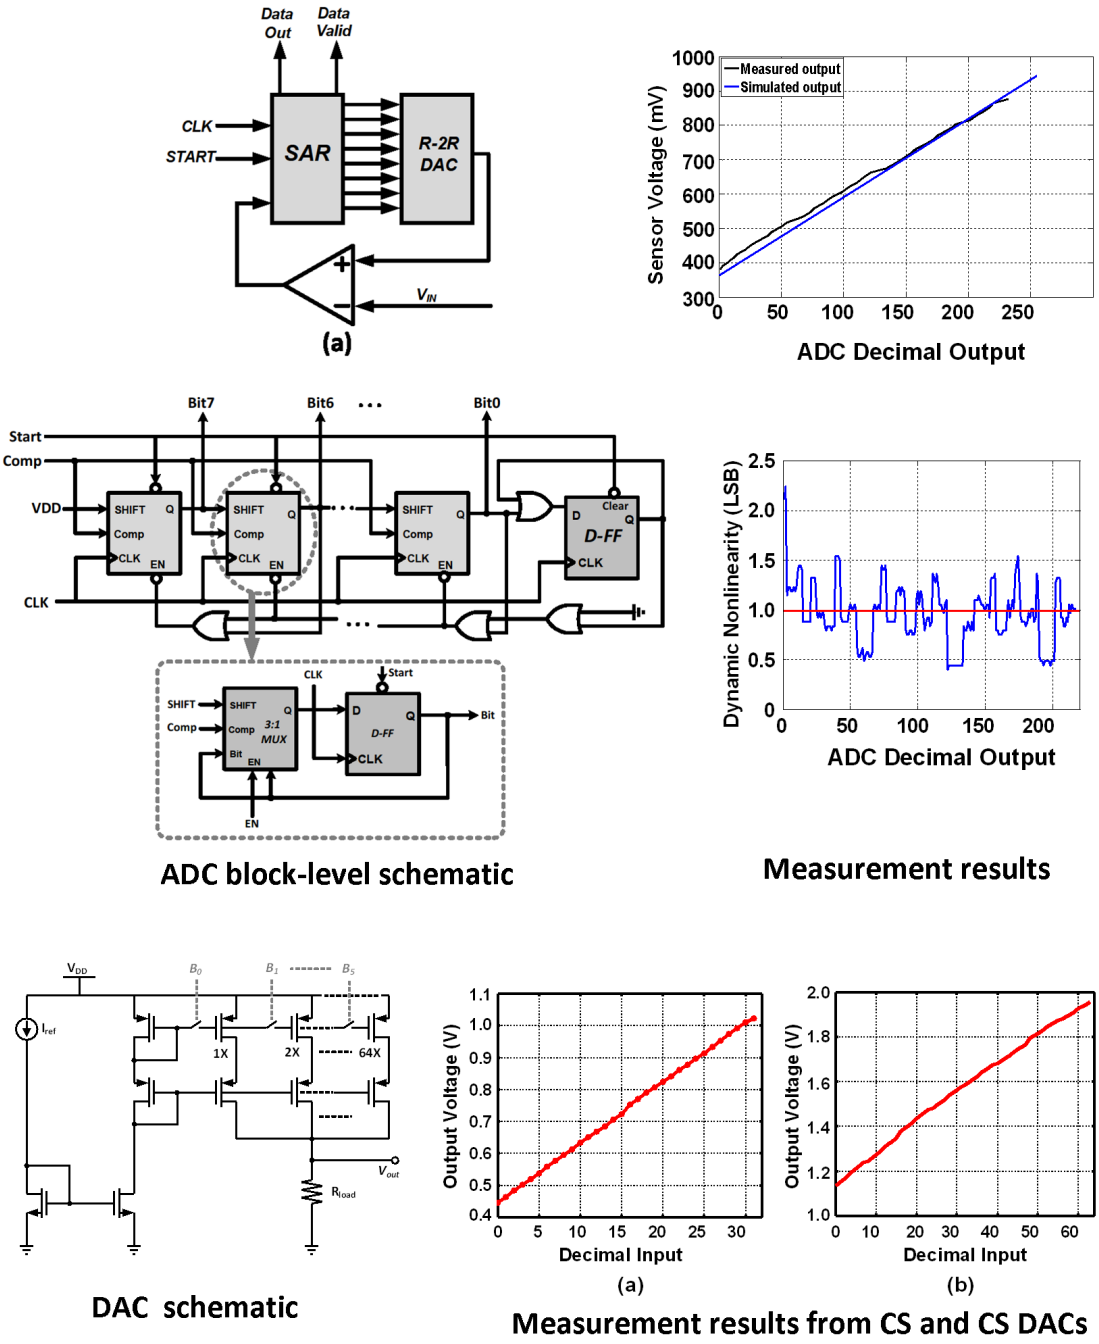

Figure 7.11: Summary of self-healing data converters, with the block diagram of the SAR ADC (upper left), ADC response to sensor voltage and DNL measuremetns (upper right), DAC schematic (lower left) and measured DAC output voltages (lower right)

Because of the modular design of the coding, more efficient algorithms can be coded and integrated into the existing code base with minimal additional effort. A flow chart of the optimization algorithm is shown in Figure 7.12.

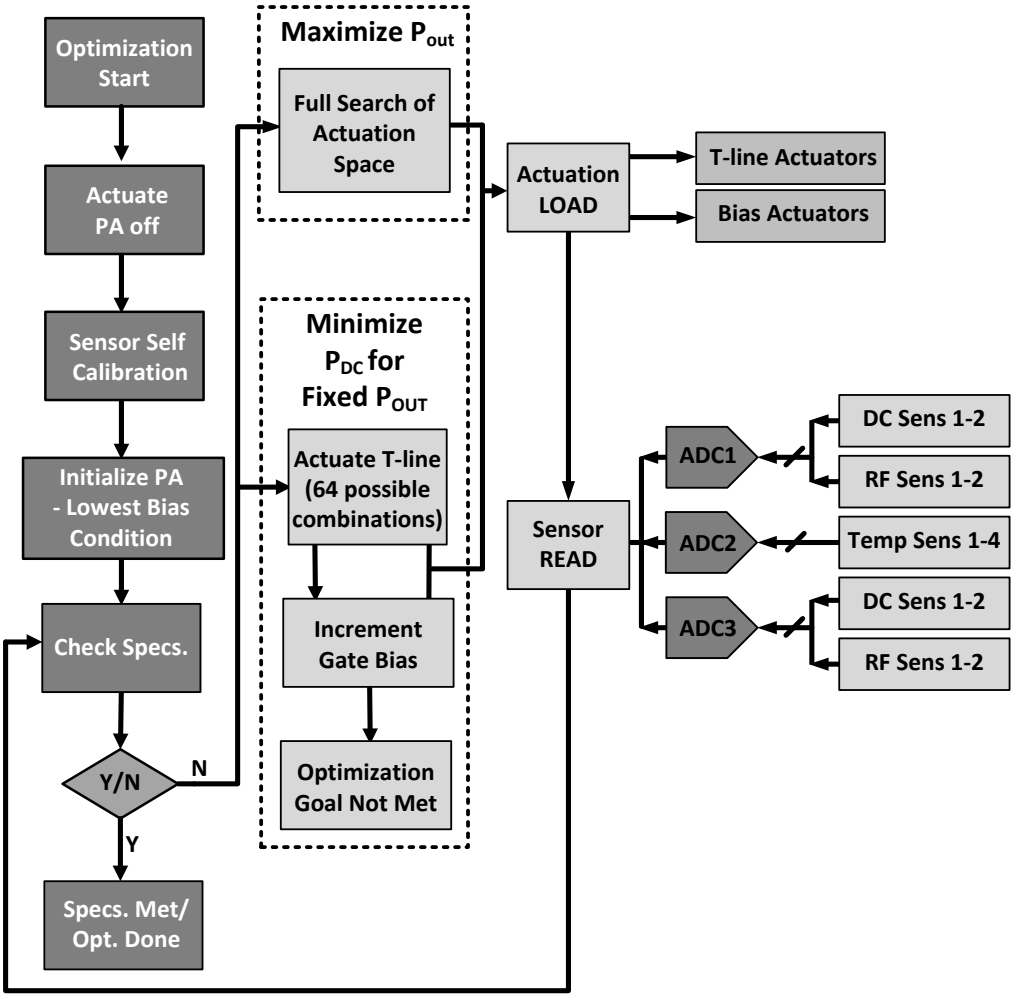

Figure 7.12: Flowchart showing details of optimization component.

Both algorithms start off with an offset calibration step where the PA is turned off digitally and all the sensor readings are recorded via the ADCs. This information is used as a self-calibration step that removes any offsets in the sensor/ADC chain. The algorithms then continue as they read data from the sensors, compute the RF output power and DC power consumption and change the actuation states to achieve the desired performance.

The digital core uses a test setup limited clock of 25 MHz and takes 3 μS per optimization iteration (set actuators, read all sensors, decide on next actuation state), which leads to a maximum healing time of 0.8 seconds in the case when the algorithm goes through every single actuation state.

7.5 System Measurements

To measure the entire PA self-healing system, the chip was mounted on a printed circuit board (PCB) and probed in the setup depicted in Figure 7.13. In order to evaluate the benefits of the self-healing, a nominal actuation state was selected as the default state without self-healing. That state was chosen to be the one that had best performance at saturation in simulation with typical transistors, and represents the optimal state with the information the designer had at the time of fabrication. The DC power consumption of all of the self-healing blocks is omitted when considering the performance of the default state, to obtain the best estimate of the performance of the amplifier without self-healing. Because the entire self-healing loop is integrated on the CMOS chip, to run the self-healing algorithm, a mode of operation, along with the desired output power (if necessary), and the 'run' command are the only things sent to the chip. The entire rest of the process is locally automated within the chip.

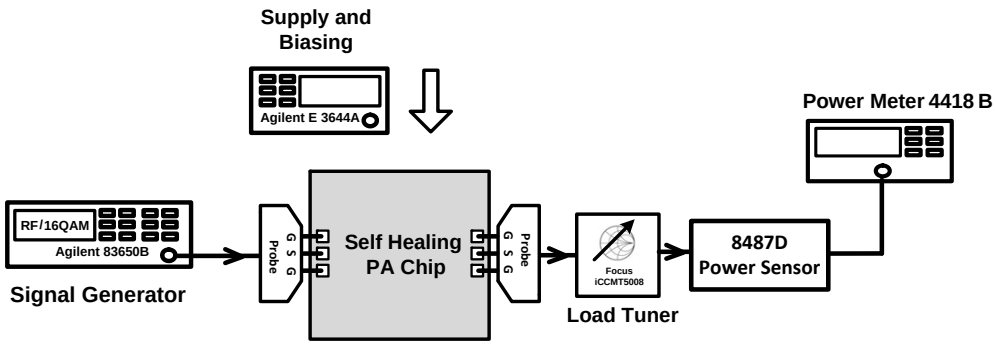

Figure 7.13: Measurement setup: PA chip is mounted on a printed circuit board (PCB) and probed. Calibrated mm-wave load tuner can set load impedances up to the 4-1 VSWR circle at the tips of the probe.

7.5.1 Self-Healing Measurements with 50 Ω Load

The healing ability of the amplifier for 50 Ω loads is presented first. Figure 7.14a shows the output power verses input power for an amplifier in its default state, as well as one that has been healed for maximum output power at low input power levels, and one that has been healed for maximum output power at the 1 dB compression point, the point where the gain has compressed 1 dB compared to the small signal gain. This plot shows the improvement in output power that can be achieved using self-healing, but also shows that there is no one optimal state for all input powers. The optimum load impedance for the maximum output power from the output amplifying transistors varies based upon the input power levels, and thus by tuning the matching network for the current power level, the corresponding optimal matching network can be found. This is an added benefit of the self-healing system, as it can be healed for the desired power level, and if that desired power level changes, it can be healed again. Near saturation where the default state was designed, it is close to the optimum, and thus there is not as much room for improvement. However, the default match at

small signal is farther from the optimum, and thus there self-healing can provide larger improvements to the performance. After healing, the PA achieves a small signal gain of 21.5 dB, with output powers of 16 dBm at saturation and 12.5 dBm at the 1 dB compression point while drawing 520 mW of DC power at 28 GHz.

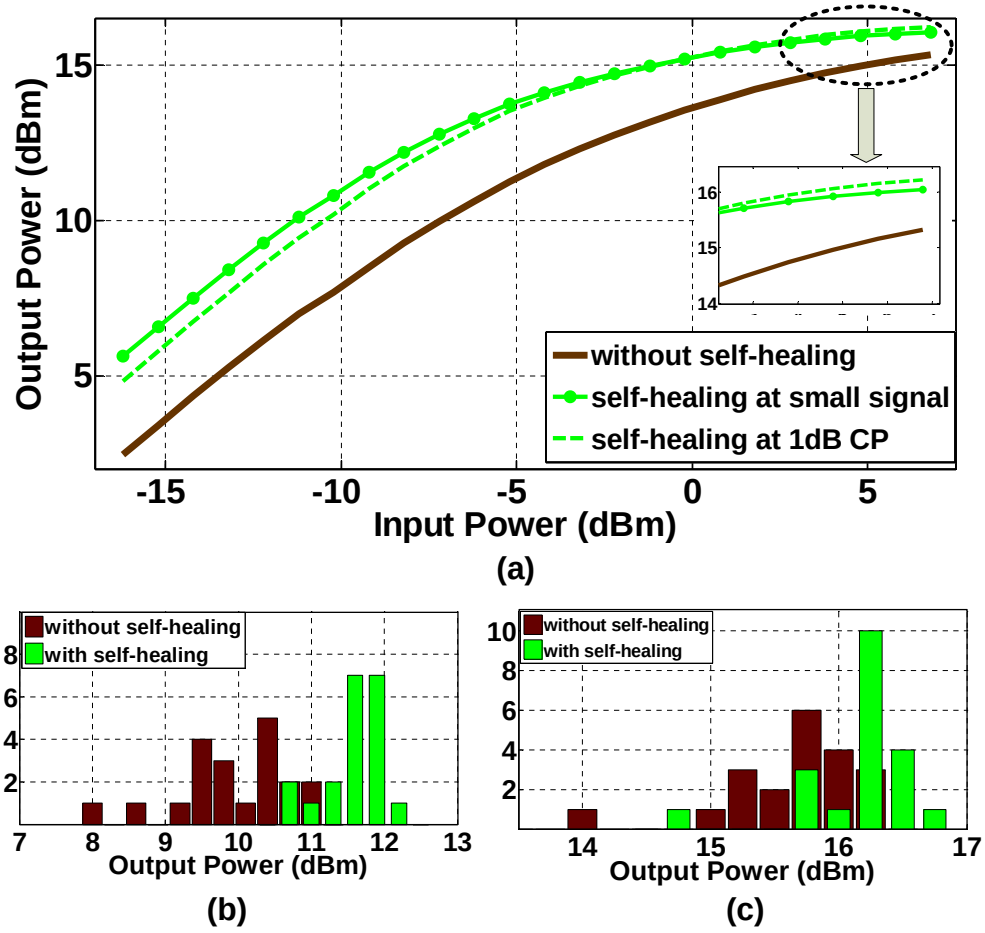

Figure 7.14: (a) Self-healing for maximum output power at small signal and near 1 dB compression point. Histograms of 20 chips in (b) back-off and (c) near 1 dB compression show improvement in output power as well as decrease in variation between chips with healing.

Because the default state without self-healing is designed for an amplifier near saturation, the default state is closer to optimal near saturation, and thus there is not as much room for improvement. The measurements show that in the back-off region, where the default state is further from optimal, the ability of the system to self-heal the PA increases. To evaluate the viability of the system over process variation, 20 chips were measured, and histograms, showing operation at small signal and near the 1 dB compression point, are plotted in Figure 7.14c.

The second algorithm to minimize the DC power while maintaining a desired output power is shown next. Figure 7.15a shows the DC power consumption of 20 chips for various output power levels, with a histogram

cross section of that plot for 12.5 dBm output power, near the 1 dB compression point, shown in Figure 7.15b. Because the state is not changing, the DC power levels without self-healing for each chip are relatively flat until the transistors really start to saturate and the power increases slightly. Once self-healing is turned on, there is still high DC power required to achieve very high output powers, but once the desired output power becomes even a couple dB below the saturated power, significant reduction in DC power consumption is observed. The DC power required to produce 12.5 dBm output power sees a 47% reduction in average power level over 20 chips, with a 78% decrease in the standard deviation between chips. This means that self-healing is both improving the performance, and also making it much more consistent across chips than in the default case. Once the power is near small signal levels, reductions of greater than 50% for every single chip measured are achieved.

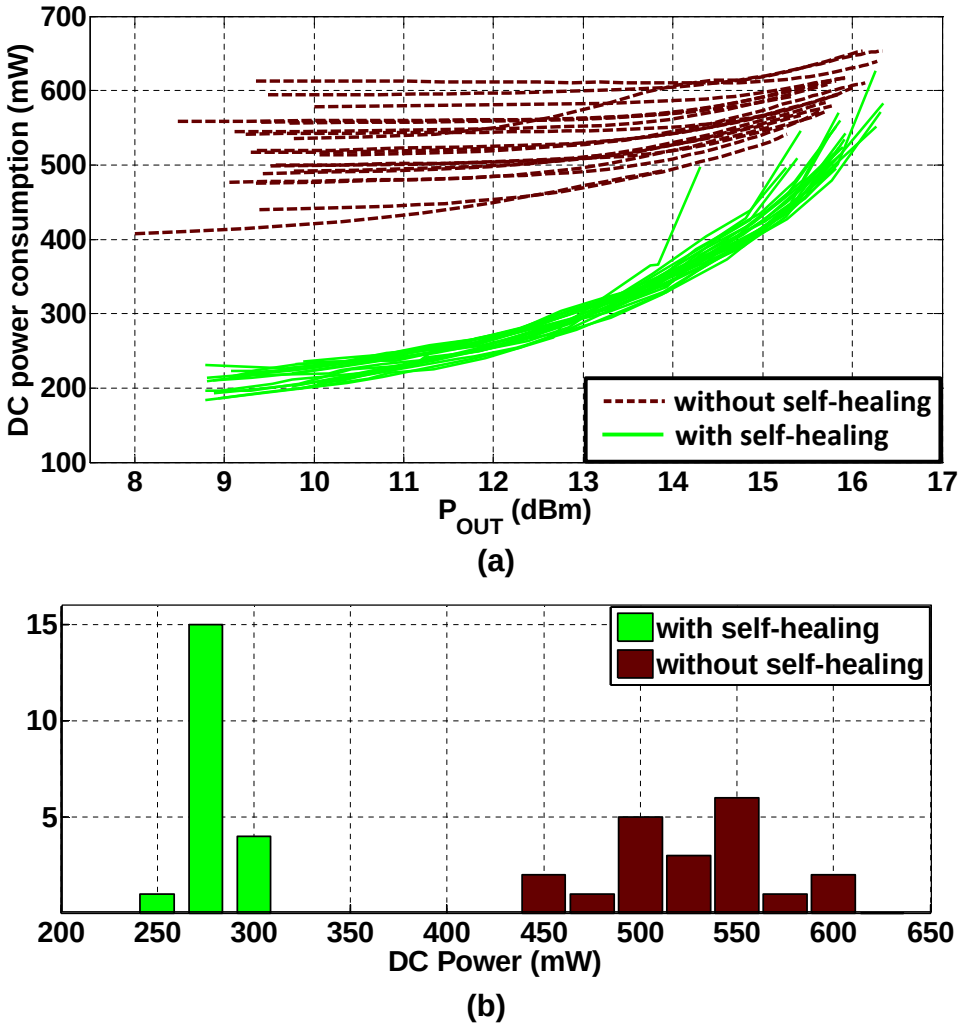

Figure 7.15: (a) Self-healing to minimize DC power while maintaining specified output power for 20 chips, with (b) a histogram when P_{OUT}=12.5 dBm.

7.5.2 Self-Healing Under Load Mismatch

To imitate the effects of antenna load impedance mismatch, a focus microwaves mm-wave load tuner [114, 115] capable of producing loads within the 12-1 VSWR circle was used. When the measurement setup is calibrated with the loss of the connecting cable and probe, it is capable of producing impedances at the probe tips within the 4-1 VSWR circle. Because the RF power sensors are capable of detecting the power delivered to the load regardless of its impedance, the algorithm does not need to determine what the load impedance is, as long as it can optimize the desired performance metric. Measurements were taken by sweeping through the achievable load impedances, and at each impedance running self-healing for each algorithm. The results for the algorithm to maximize output power are presented as a contour plot on Smith charts before and after self-healing in Figure 7.16. There is improvement across the entire 4-1 VSWR circle. For example the 15.5 dBm contour just barely shows up before healing, but occupies a large area of the Smith chart when self-healing is turned on.

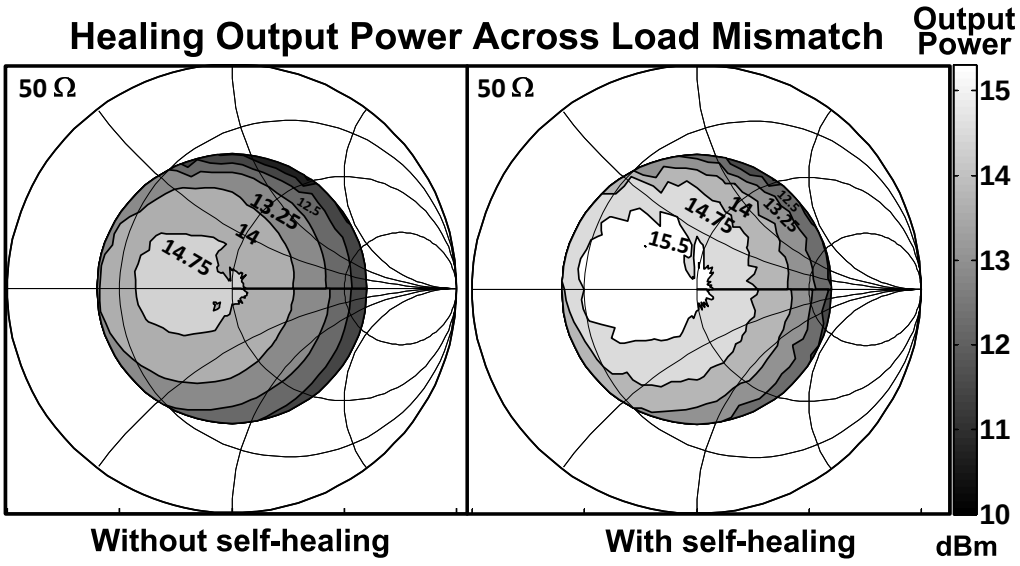

Figure 7.16: Contours before and after self-healing for maximum output power under load impedance mismatch.

To quantify self-healing for load impedance mismatch over process variation, 10 chips were measured, and two histograms are plotted in Figure 7.17. The first is taken with a load impedance of 33 Ω, where the output power is at a maximum, and the second is on the edge of the 4-1 VSWR circle at (23-44j) Ω.

The measurement results of the second algorithm to minimize DC power while maintaining a specified power are shown on the two Smith charts in Figure 7.18. The specified output power was set to 12.5 dBm, and the outermost contour of each plot shows the boundary where that power specification is met in each case. Because there is only one actuation state without self-healing, the DC power level does not change with the load impedance, but when self-healing is applied, substantial savings of up to 35% are achieved especially

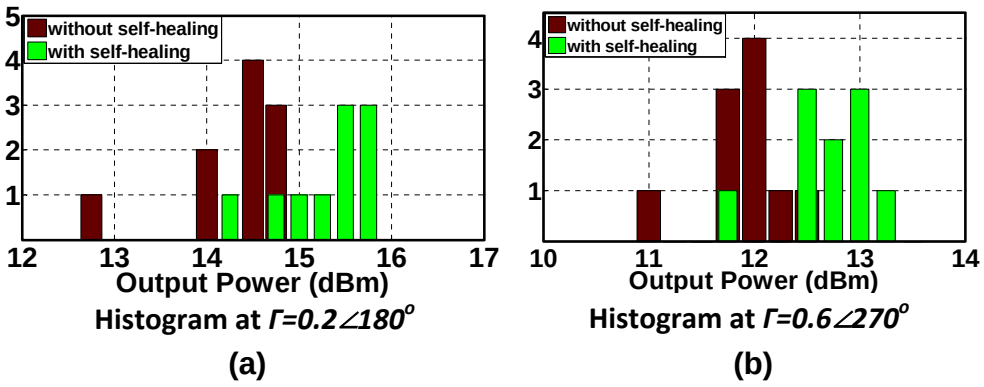

Figure 7.17: Histograms of 10 chips verify self-healing across multiple devices. One is near the optimum load (a) and the other is on the edge of the 4-1 VSWR circle (b).

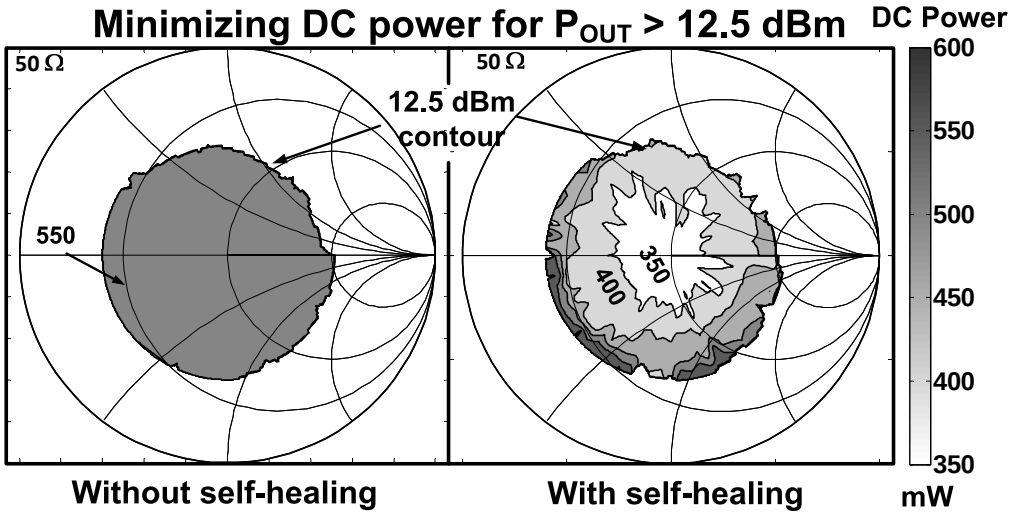

Figure 7.18: Contours before and after self-healing for minimum DC power while maintaining a specified 12.5 dBm output power.

near load impedances around 50 Ω, the impedance the match was designed for.

7.5.3 Power Amplifier Linearity Measurement

One reason that a linear PA was chosen for this design was to enable amplification of signals with non-constant envelope modulation. To verify the linearity of the PA and its operation with non-constant envelope modulation, the error vector magnitude (EVM) of a 100 ksps 16 quadrature amplitude modulation (QAM) signal was measured. A histogram of the 10 chips measured shows EVM when the PAs are outputting 12.5 dBm showing a reduction in the average value of the EVM from 5.9% down to 4.2% (Figure 7.19). The improvement in the linearity after self-healing is attributed to being able to provide the same output power without forcing the PA as far into saturation.

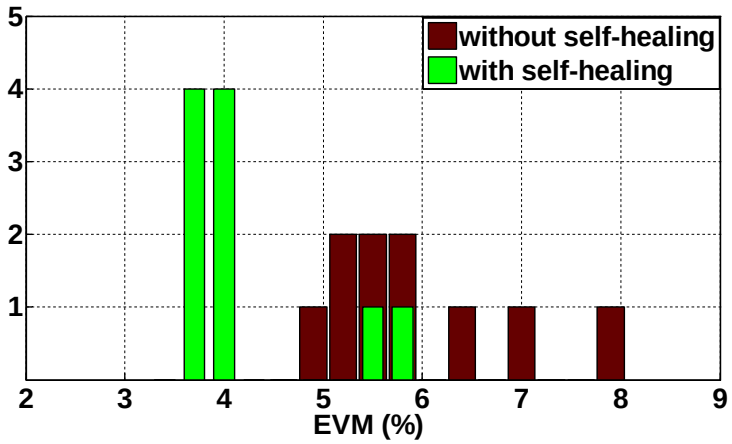

Figure 7.19: Histogram showing linearity improvement measuring the error vector magnitude (EVM) for 10 chips at 12.5 dBm output power and 100 ksps 16 quadrature amplitude modulation (QAM) before and after self-healing.

7.5.4 Self-Healing Laser-Induced Transistor Failure

Partial or total transistor failure can be caused by aging, transistor stress such as from voltage spikes, or other phenomenon. To show the self-healing system's ability to heal for partial and total transistor failure, a laser trimmer was used to blast away various parts of one of the output stage transistors. The output stage was chosen as the transistors to cut, as they are the ones that are pushed closest to breakdown and are likely going to be the first ones to fail. Only one of the two output stages was cut to cause a worst case scenario from a mismatch standpoint. As a reference, the amplifier before any laser blasting is shown in Figure 7.20a. Measurements taken in the default state and after healing for maximum output power with half of the common source transistor cut out are plotted in Figure 7.20b. Figure 7.20c shows the results with and without self-healing when half of the cascode stage was additionally cut out. Finally Figure 7.20d shows the results once the entire output stage is blasted away. This means that the matching network that was expecting to have two

similar drives at both inputs now has only a single input, and then a large stub where the other input used to be, destroying the original match that was shown in Figure 7.20.

The default case at small signal loses 7.2 dB in output power from when the output stage is whole to when it is completely cut out. 3 dB of that is due to having only one of the two output stages providing power, but another 4.2 dB is caused by mismatch of the matching network. Once healing is applied, the loss due to cutting out the output stage is only 3.3 dB, which when taking into account the 3 dB loss from only having a single output stage means that the tunable matching network was able to heal back to the point where there was only 0.3 dB additional loss from this catastrophic event. This is one of the very strong points of self-healing, as under nominal conditions, where the default design is close to optimum, self-healing can only possibly improve the circuit at most the deviation from the optimum, but in cases such as this where the default falls far from the optimum point, and would normally register as a total failure of the entire circuit, self-healing can provide very significant gains and keep the circuit operational even under these types of extreme conditions. A die photo of the entire chip with closeup images before and after the laser blasting is shown in Figure 7.21.

7.5.5 Yield Improvement and Performance Summary

In order to look at yield improvement, metrics and yield specifications must be defined. For this PA, the metrics are saturated output power, small signal gain, power added efficiency, and 4-1 VSWR tolerance, which is defined in dB as the maximum output power falloff for any load impedance within the 4-1 VSWR circle compared to the 50 Ω nominal load. Demonstrative yield specifications are set to show healing improvement for saturated output power, small signal gain and power added efficiency at $>$15.5 dBm, $>$20 dB, and $>$6% respectively, with self-healing resulting in a yield improvement from 20% to 90%, 20% to 100%, and 5% to 100% for each metric over 20 chips tested, and best achieved results of 16.5 dBm and 23.7 dB, and 7.2% respectively. For 4-1 VSWR tolerance, the yield specification was set to $<$3 dB, resulting in a yield improvement from 0% to 80% with self healing and a best achieved result of 2.28 dB over 10 chips measured. This results in an aggregate yield improvement due to self healing from 0% to 80%.

To the best of our knowledge, there is no stand-alone integrated self-healing PA to compare with this work. To do a fair comparison of this self-healing PA to other PAs without self-healing, additional data, such as variation between multiple chips, tolerance to VSWR events, efficiency in the back-off region, and amplifier operation in the case of transistor failure would be necessary and has not been reported in those publications.

7.6 Conclusion

Self-healing can enable significant improvements in performance if a mm-wave design has become susceptible to performance degradation caused by process or environmental variation. The effectiveness of the

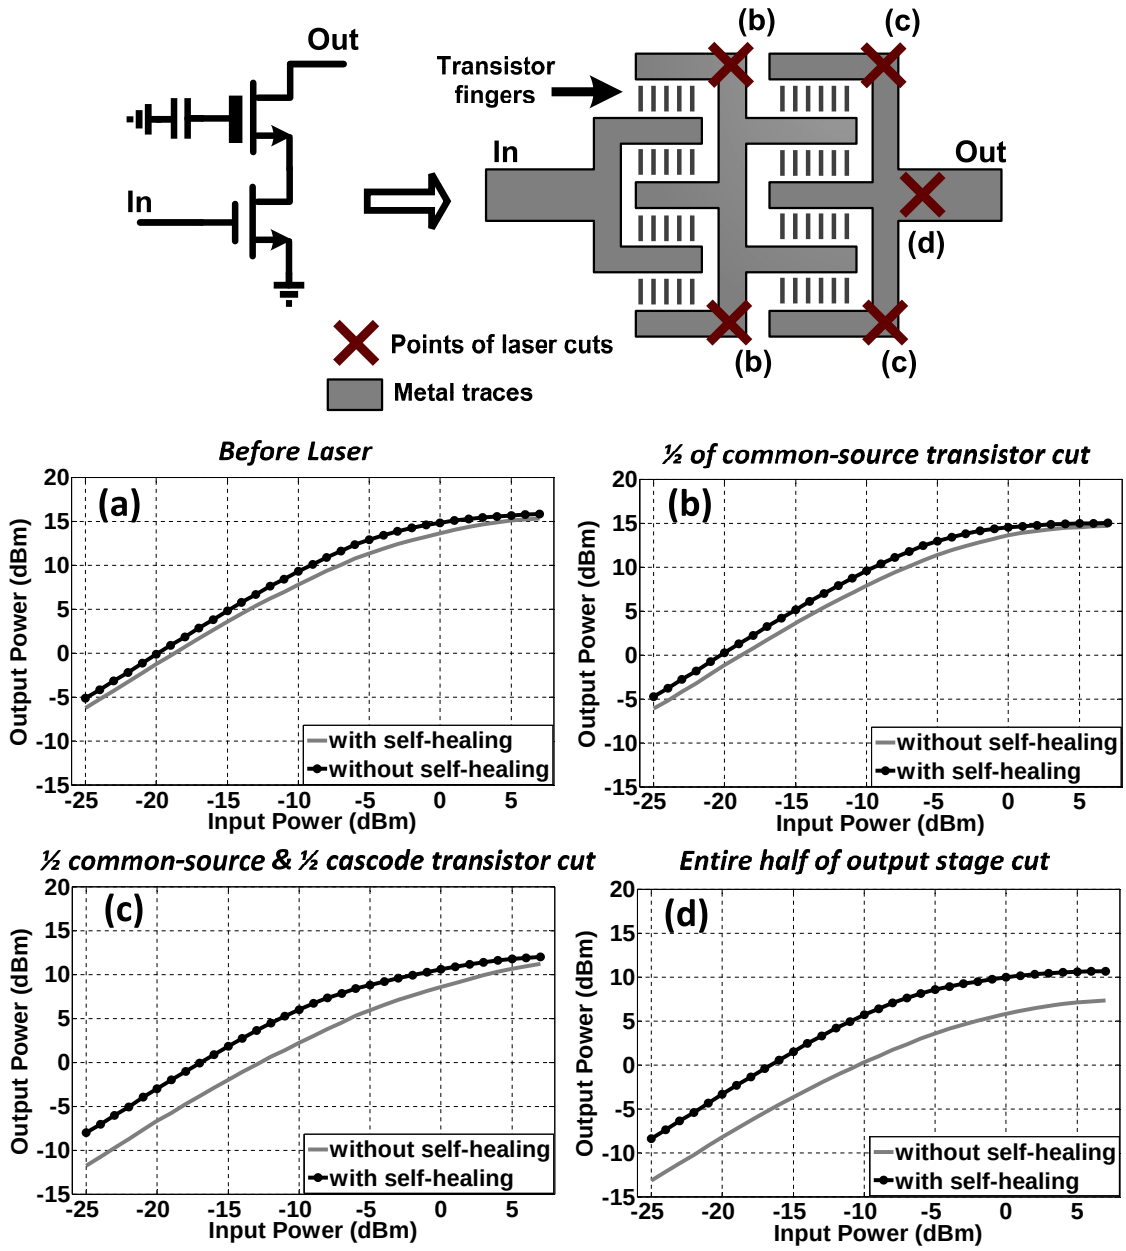

Figure 7.20: Schematic and layout location of laser trim points, and measurements before and after self-healing for maximum output power at various stages of transistor failure due to laser blasting show more than 5 dB improvement when self-healing is used in the worst case scenario of an entire output stage failing.

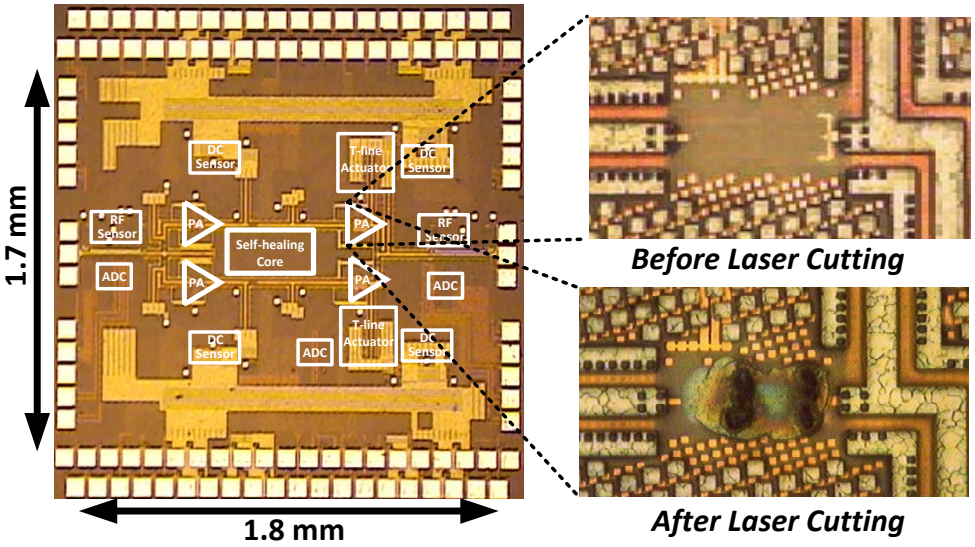

Figure 7.21: Chip photograph showing functional blocks, and closeup of the output stage before and after laser cutting.

self-healing system is dependent on how far from optimal a non-healing design becomes. By programming different algorithms into the digital core, the amplifier can become dynamic, and tradeoffs that traditionally had to be made at the design level before fabrication can be adjusted based upon the current environment and use case of the mm-wave circuit.

Self-healing enables aggressive mm-wave designs even in process nodes where process variation and mismatch can cause significant variation in transistor operation. By using robust design to sense the performance of the mm-wave circuit during the self-healing optimization process, the core mm-wave circuit can take advantage of that aggressive design and push performance levels while maintaining yield. It is expected that as the minimum feature size continues to decrease, process variation will continue to increase, and the need for self-healing or other reconfigurable circuit techniques will increase.

Environmental variation can also be mitigated using self-healing. Variation in PA performance due to load impedance mismatch caused by VSWR events can be handled by making the output matching network tunable, so that the output amplifier stages still see their optimum impedance for maximum output power. Transistor degradation and failure due to effects such as aging can also be counteracted by adjustment of the DC operation point of the transistors and by adjusting the matching networks as the optimal impedances change.

A 28 GHz power amplifier was presented as a case study of how such an integrated self-healing PA could be implemented. Measurements of multiple chips demonstrate the viability of an integrated self-healing system that requires no external calibration of any kind. Integrating the sensors, actuators, digital algorithm and data converters on a single chip allows for a completely automated healing system that improves aggre-

gate yield from several performance specifications from 0% up to 80%, while healing for process variation, mismatch, load impedance mismatch, and partial and total transistor failure.

al electronic

Chapter 8

Conclusion

This thesis has advanced the state of the art in both integrated radiators as well as reliable power generation. The concept of multi-port driven (MPD) radiators was introduced, opening up an entirely new design space that takes advantage of the free transistors and connections that are available when the antenna is integrated onto the same substrate as the driver circuitry. By placing transistors next to the antennas themselves, and driving them at several ports, an MPD radiator can eliminate several lossy blocks including those for impedance matching, power transfer, and power combination. A proof of concept 160 GHz 8-spoke radial MPD radiator was analyzed, designed, fabricated and measured. The analysis gave insights into the important design parameters of the radial MPD antenna, and showed the reason for the wide-band nature of the structure. The fabricated chip was measured to have a +4.6 dBm effective isotropic radiated power (EIRP), with 0.63 mW total radiated power at 161 GHz, and closely matched simulations. This EIRP and output power were the highest EIRP per element and highest radiated power of any reported integrated radiator on silicon without external modifications at the time of its publication.

The radial MPD concept was expanded to include dynamic polarization control (DPC), which enables one or more antennas to dynamically adjust their radiated polarization entirely electronically, without any mechanical reconfiguration. This can enable a single transmitter to transmit to any receive antenna regardless of its polarization or orientation in space, and can even track the polarization of that receive antenna as it moves, always maintaing a polarization match that enables the strongest possible coupling between antennas. DPC can even be used to modulate the radiated signal's polarization, providing a way to have up to 4-dimensional data transfer that cannot be fully detected by a receive antenna with only a single polarization. A 2x1 radial MPD array with DPC at 105 GHz was designed, fabricated and measured. When producing linear polarization, the polarization angle was able to be set to any value in the full range from 0° though 180° while maintaining an axial ratio above 10dB. Control of the axial ratio was also demonstrated, and ratios between 2.4 dB through 14 dB were shown while keeping the polarization angle constant. The array is capable of beam steering, and has a measured EIRP of 7.8 dBm and 0.9 mW total radiated power, which is currently the highest EIRP per element and total radiated power of any reported integrated radiator on silicon without external modifications.

The concept of the MPD antenna was also implemented in a silicon photonics process and replaced the driver transistor circuitry with silicon optical waveguides and photodiodes. By beaming two lasers with slightly different frequencies, THz signals at 350 GHz can be produced, and fed into MPD antennas on the substrate. A key benefit of producing the THz signal this way is that the silicon optical waveguides that input the signal and the metal antennas that output the signal are independent from each other and do not couple, making the design of the input distribution much simpler than in the all-electronic versions. Three versions of 2x2 optical MPD radiator arrays have been designed, one that does beam forming, one that does beam steering, and one that does DPC. They have simulated gains of 12.1 dBi for the first two chips and 11.9 dBi for the third chip, and simulated EIRPs of -2 dBm and -3 dBm respectively.

Self-healing was also used in two designs to improve the reliability of the radiation and power generation. First, a partially integrated self-healing MPD radiator was shown that uses substrate mode pickup sensors to infer the far field radiation of the chip, and DC current sensors to sense the DC power consumption. If performance degradation is observed, the system is designed to actuate the amplitudes and phases of the driver stages to rectify any issues and return the performance to the optimal state.

A fully-integrated and measured self-healing power amplifier (PA) system at 28 GHz was also presented. It uses DC current, radio frequency (RF) power, and thermal sensors to detect the output power, gain and efficiency of the PA, and sends that data to a digital core that is implemented on-chip. The core runs an optimization algorithm that can then actuate the PA by tuning the output power combining matching network through tunable transmission line stubs, or by adjusting the DC operating points of the amplifying stages through bias control of both the common source and cascode transistors. Healing was shown for variation in process, transistor mismatch for 20 measured chips, showing greater than 50% savings in DC power consumption at back off, and a significant decrease in variation between chips, as well as load impedance mismatch and linearity for 10 measured chips. Healing was even shown in the case of laser-induced partial and total transistor failure, and improvements of up to 3.9 dB in output power were observed. A representative yield specification was set, and the yield was improved from 0% up to 80%.

Appendix A

Mapping Current on a Plane to Broadside Far-Field Electric Field

This appendix will derive from Maxwell's equations an equation that will map the currents on a plane to the broadside electric field in the far field. The first part of the analysis follows the one found in [1], and then approximations based upon the specific environment of mapping broadside radiation from currents on a plane are made to simplify the final equation that is used as the starting point for the analysis in Chapter 3 to gain intuition for how to design on chip-antennas.

A.1 General Mapping of Currents to Electric Fields in the Far Field

The first set of equations to derive are a mapping from electric currents on a plane (chip) to electric fields in the far-field at an angle normal to the plane containing the currents. Start with Maxwell's equations:

$$\nabla \cdot \mathbf{B} = 0, \tag{A.1.1}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \tag{A.1.2}$$

and

$$\nabla \times \mathbf{B} = \mu \left(\mathbf{J} + \epsilon \frac{\partial \mathbf{E}}{\partial t} \right) \tag{A.1.3}$$

We are interested in time-harmonic fields that take the form

$$\mathbf{F}(t) = F_0 e^{j\omega t} \hat{\mathbf{u}} \tag{A.1.4}$$

where ω is the angular frequency, and thus can write

$$\frac{\partial \mathbf{F}}{\partial t} = j\omega \mathbf{F} \tag{A.1.5}$$

and can rewrite equation A.1.2 as

$$\nabla \times \mathbf{E} = -j\omega\mathbf{B}. \tag{A.1.6}$$

Because the divergence of $\mathbf{B} = 0$, and we are considering linear homogeneous materials, there exists fields H and A, such that

$$\mathbf{B} = \mu\mathbf{H} = \nabla \times \mathbf{A}. \tag{A.1.7}$$

Thus

$$\nabla \times \mathbf{E} = -j\omega\nabla \times \mathbf{A}, \tag{A.1.8}$$

and

$$\nabla \times [\mathbf{E} + j\omega\mathbf{A}] = 0. \tag{A.1.9}$$

Using (A.1.9) and the vector identity

$$\nabla \times (-\nabla\phi_e) = 0, \tag{A.1.10}$$

there exists ϕ_e such that

$$\mathbf{E} = -\nabla\phi_e - j\omega\mathbf{A}. \tag{A.1.11}$$

Taking the curl of both sides of (A.1.7) yields

$$\nabla \times \mu\mathbf{H} = \nabla \times \nabla \times \mathbf{A}, \tag{A.1.12}$$

and with the vector identity

$$\nabla \times \nabla \times \mathbf{A} = \nabla(\nabla \cdot \mathbf{A}) - \nabla^2\mathbf{A} \tag{A.1.13}$$

and the fact that the medium is homogeneous,

$$\mu\nabla \times \mathbf{H} = \nabla(\nabla \cdot \mathbf{A}) - \nabla^2\mathbf{A}. \tag{A.1.14}$$

Using (A.1.3), time-harmonic fields, and homogeneous materials,

$$\nabla \times \mathbf{H} = \mathbf{J} + j\omega\epsilon\mathbf{E} \tag{A.1.15}$$

and thus

$$\mu\mathbf{J} + j\omega\mu\epsilon\mathbf{E} = \nabla(\nabla \cdot \mathbf{A}) - \nabla^2\mathbf{A}. \tag{A.1.16}$$

Then, inserting (A.1.11), we get

$$\mu\mathbf{J} + j\omega\mu\epsilon(-\nabla\phi_e - j\omega\mathbf{A}) = \nabla(\nabla \cdot \mathbf{A}) - \nabla^2\mathbf{A}. \tag{A.1.17}$$

Next, define k such that $k^2 = \omega^2\mu\epsilon$, and rearranging terms yields

$$\nabla^2\mathbf{A} + k^2\mathbf{A} = -\mu\mathbf{J} + \nabla(\nabla\cdot\mathbf{A} + j\omega\mu\epsilon\phi_e) \tag{A.1.18}$$

To simplify the right side of the equation, we will now define ϕ_e to make

$$\nabla\cdot\mathbf{A} = -j\omega\mu\epsilon\phi_e, \tag{A.1.19}$$

or

$$\phi_e = -\frac{\nabla\cdot\mathbf{A}}{j\omega\mu\epsilon} \tag{A.1.20}$$

which results in

$$\nabla^2\mathbf{A} + k^2\mathbf{A} = -\mu\mathbf{J}. \tag{A.1.21}$$

Also, (A.1.11) now reduces to

$$\mathbf{E} = -\nabla\phi_e - j\omega\mathbf{A} = -j\omega\mathbf{A} - j\frac{1}{\omega\mu\epsilon}\nabla(\nabla\cdot\mathbf{A}) \tag{A.1.22}$$

and

$$\mathbf{H} = \frac{1}{\mu}\nabla\times\mathbf{A} \tag{A.1.23}$$

With equation A.1.21, we can start with a current density, and integrate to find $\mathbf{A}$, then with equations A.1.22 and A.1.23, we can take derivitives of $\mathbf{A}$ to find $\mathbf{E}$ and $\mathbf{H}$.

A.2 Solution for Point Sources

Next, we will solve out equations A.1.21 - A.1.23 for a point source, where a current density J_z is located at the origin and points along the z-axis. Thus

$$\nabla^2 A_z + k^2 A_z = -\mu J_z \tag{A.2.1}$$

at the origin, and at points away from the origin, $J_z = 0$, so

$$\nabla^2 A_z + k^2 A_z = 0 \tag{A.2.2}$$

away from the origin.

Since it is a point source, A_z will only be a function of r, and not direction (θ or ϕ), so we can write

$$\nabla^2 A_z(r) + k^2 A_z(r) = \frac{1}{r^2}\frac{\partial}{\partial r}\left[r^2\frac{\partial A_z(r)}{\partial r}\right] + k^2 A_z(r) = 0 \tag{A.2.3}$$

which reduces to

$$\frac{d^2A_z(r)}{dr^2} + \frac{2}{r}\frac{dA_z(r)}{dr} + k^2A_z(r) = 0 \tag{A.2.4}$$

There are two solutions,

$$A_{z1} = C_1\frac{e^{-jkr}}{r} \tag{A.2.5}$$

and

$$A_{z2} = C_2\frac{e^{-jkr}}{r} \tag{A.2.6}$$

We want the solution with outwardly traveling wave ($A_z = A_{z1}$). Also we notice that the static case ($\omega = 0$) $A_z = \frac{C_1}{r}$, so we can solve the static case, and then multiply by e^{-jkr}.

Thus

$$\nabla^2 A_z = -\mu J_z \tag{A.2.7}$$

which is Poisson's equation, whose solution is

$$A_z = \frac{\mu}{4\pi}\iiint_V \frac{J_z}{r}dv \tag{A.2.8}$$

and thus for the time varying solution

$$A_z = \frac{\mu}{4\pi}\iiint_V J_z\frac{e^{-jkr}}{r}dv \tag{A.2.9}$$

Similarly,

$$A_x = \frac{\mu}{4\pi}\iiint_V J_x\frac{e^{-jkr}}{r}dv \tag{A.2.10}$$

and

$$A_y = \frac{\mu}{4\pi}\iiint_V J_y\frac{e^{-jkr}}{r}dv \tag{A.2.11}$$

which leads to

$$\mathbf{A} = \frac{\mu}{4\pi}\iiint_V \mathbf{J}\frac{e^{-jkr}}{r}dv \tag{A.2.12}$$

For electric currents, the volume integral of (A.2.12) reduces to the line integral

$$\mathbf{A} = \frac{\mu}{4\pi}\int_C \mathbf{I}\frac{e^{-jkR}}{R}dl \tag{A.2.13}$$

A.3 Far-Field Approximation

Starting with the general form of $\mathbf{A}$

$$\mathbf{A} = A_r(r,\theta,\phi)\hat{\mathbf{u}}_\mathbf{r} + A_\theta(r,\theta,\phi)\hat{\mathbf{u}}_\theta + A_\phi(r,\theta,\phi)\hat{\mathbf{u}}_\phi \tag{A.3.1}$$

If we are far enough away from the source, all of the $1/r^n$ terms where $n > 1$ decay become negligible compared to n=1, and no terms can exist for n¡1 as it violates conservation of energy. Thus we can conclude that A will decay by $1/r$ in the far field, and will only have the phase shift e^{-jkr} found in (A.2.12), leading to

$$\mathbf{A} = [A'_r(\theta,\phi)\hat{\mathbf{u}}_\mathbf{r} + A'_\theta(\theta,\phi)\hat{\mathbf{u}}_\theta + A'_\phi(\theta,\phi)\hat{\mathbf{u}}_\phi]\frac{e^{-jkr}}{r} \qquad r \to \infty \tag{A.3.2}$$

and to solve equation A.1.22, we first calculate the $\nabla(\nabla\cdot\mathbf{A})$,

$$\begin{aligned}\nabla(\nabla\cdot\mathbf{A}) = \nabla\frac{1}{r^2}\frac{\partial}{\partial r}\left[r^2\frac{e^{-jkr}}{r}A'_r(\theta,\phi)\right] + \nabla\frac{1}{r\sin\theta}\frac{\partial}{\partial\theta}\left[\sin\theta\frac{e^{-jkr}}{r}A'_\theta(\theta,\phi)\right] + \\ \nabla\frac{1}{r\sin\theta}\frac{\partial}{\partial\phi}\left[\frac{e^{-jkr}}{r}A'_\phi(\theta,\phi)\right].\end{aligned} \tag{A.3.3}$$

First,

$$\begin{aligned}\nabla\frac{1}{r^2}\frac{\partial}{\partial r}\left[r^2\frac{e^{-jkr}}{r}A'_r(\theta,\phi)\right] &= \nabla\frac{1}{r^2}A'_r(\theta,\phi)\frac{\partial}{\partial r}\left[r\frac{e^{-jkr}}{r}\right] \\ &= \nabla\frac{1}{r^2}A'_r(\theta,\phi)\left(e^{-jkr} - jkre^{-jkr}\right) \\ &= \nabla A'_r(\theta,\phi)\left(\frac{e^{-jkr}}{r^2} - \frac{jke^{-jkr}}{r}\right)\end{aligned} \tag{A.3.4}$$

$$\begin{aligned}= A'_r(\theta,\phi)\left[-\frac{k^2e^{-jkr}}{r} - \frac{2e^{-ikr}}{r^3}\right]\hat{\mathbf{u}}_\mathbf{r} + \\ \frac{1}{r}\left(\frac{e^{-jkr}}{r^2} - \frac{jke^{-jkr}}{r}\right)\frac{\partial}{\partial\theta}A'_r(\theta,\phi)\hat{\mathbf{u}}_\theta + \\ \frac{1}{r\sin\theta}\left(\frac{e^{-jkr}}{r^2} - \frac{jke^{-jkr}}{r}\right)\frac{\partial}{\partial\phi}A'_r(\theta,\phi)\hat{\mathbf{u}}_\phi\end{aligned}$$

If we drop all $1/r^n$ terms where $n > 1$, as $r \to \infty$, and remember that $k^2 = \omega^2\mu\epsilon$, (A.3.4) reduces to

$$\nabla\frac{1}{r^2}\frac{\partial}{\partial r}\left[r^2\frac{e^{-jkr}}{r}A'_r(\theta,\phi)\right] = -\frac{\omega^2\mu\epsilon e^{-jkr}}{r}A'_r(\theta,\phi)\hat{\mathbf{u}}_\mathbf{r} \tag{A.3.5}$$

The second part of (A.3.3) is simplified as

$$\nabla \frac{1}{r\sin\theta}\frac{\partial}{\partial\theta}\left[\sin\theta\frac{e^{-jkr}}{r}A'_\theta(\theta,\phi)\right] = \nabla\frac{e^{-jkr}}{r^2\sin\theta}\frac{\partial}{\partial\theta}\left(\sin\theta A'_\theta(\theta,\phi)\right)$$

$$\begin{aligned} &= \frac{e^{-jkr}(-2-jkr)}{r^3\sin\theta}\frac{\partial}{\partial\theta}\left(\sin\theta A'_\theta(\theta,\phi)\right)\hat{\mathbf{u}}_\mathbf{r} + \\ &\quad \frac{e^{-jkr}}{r^3}\frac{\partial}{\partial\theta}\left[\frac{1}{\sin\theta}\frac{\partial}{\partial\theta}\left(\sin\theta A'_\theta(\theta,\phi)\right)\right]\hat{\mathbf{u}}_\theta + \\ &\quad \frac{e^{-jkr}}{r^3\sin^2\theta}\frac{\partial^2}{\partial\phi\partial\theta}\left(\sin\theta A'_\theta(\theta,\phi)\right)\hat{\mathbf{u}}_\phi \end{aligned} \tag{A.3.6}$$

Neglecting the $1/r^n$ terms for $n > 1$,

$$\nabla \frac{1}{r\sin\theta}\frac{\partial}{\partial\theta}\left[\sin\theta\frac{e^{-jkr}}{r}A'_\theta(\theta,\phi)\right] = 0. \tag{A.3.7}$$

Finally, the third part of (A.3.3) is simplified as

$$\nabla \frac{1}{r\sin\theta}\frac{\partial}{\partial\phi}\left[\frac{e^{-jkr}}{r}A'_\phi(\theta,\phi)\right] = \nabla\frac{e^{-jkr}}{r^2\sin\theta}\frac{\partial}{\partial\phi}\left(A'_\phi(\theta,\phi)\right)$$

$$\begin{aligned} &= \frac{e^{-jkr}(-2-jkr)}{r^3\sin\theta}\frac{\partial}{\partial\theta}\left(A'_\phi(\theta,\phi)\right)\hat{\mathbf{u}}_\mathbf{r} + \\ &\quad \frac{e^{-jkr}}{r^3}\frac{\partial}{\partial\theta}\left[\frac{1}{\sin\theta}\frac{\partial}{\partial\theta}\left(A'_\phi(\theta,\phi)\right)\right]\hat{\mathbf{u}}_\theta + \\ &\quad \frac{e^{-jkr}}{r^3\sin^2\theta}\frac{\partial^2}{\partial\phi\partial\theta}\left(A'_\phi(\theta,\phi)\right)\hat{\mathbf{u}}_\phi \end{aligned} \tag{A.3.8}$$

Again, neglecting the $1/r^n$ terms for $n > 1$,

$$\nabla \frac{1}{r\sin\theta}\frac{\partial}{\partial\phi}\left[\frac{e^{-jkr}}{r}A'_\phi(\theta,\phi)\right] = 0. \tag{A.3.9}$$

Plugging equations A.3.5,A.3.7 and A.3.9 back into equation A.3.3, we see

$$\nabla(\nabla\cdot\mathbf{A}) = -\frac{\omega^2\mu\epsilon e^{-jkr}}{r}A'_r(\theta,\phi)\hat{\mathbf{u}}_\mathbf{r} \tag{A.3.10}$$

and plugging that into (A.1.22) yields

$$\mathbf{E} = -j\omega\mathbf{A} + \frac{j\omega e^{-jkr}}{r}A'_r(\theta,\phi)\hat{\mathbf{u}}_\mathbf{r} = \frac{-jwe^{-jkr}}{r}\left[A'_\theta(\theta,\phi)\hat{\mathbf{u}}_\theta + A'_\phi(\theta,\phi)\hat{\mathbf{u}}_\phi\right] \tag{A.3.11}$$

A.4 Solving the broadside, far-field radiation from currents on a plane.

Next, A.2.12 and A.3.11 are combined and simplified as we are considering the broadside radiation produced by currents all contained within a plane, such as radiation from an integrated circuit, or printed circuit board (PCB). The first assumption made is that is that the currents are on wires that are all in the same plane, so the volumetric integral becomes a line integral. Second, we will assume that for broadside radiation, the observation point is far enough away that the observation distance r=R_0 is constant across the current plane, and thus can be simplified as

$$\mathbf{A} = \frac{\mu}{4\pi}\iiint_V \mathbf{J}\frac{e^{-jkR_0}}{R_0}dv = \frac{\mu}{4\pi}\int_C \mathbf{I}\frac{e^{-jkR_0}}{r}dl = \frac{\mu}{4\pi}\frac{e^{-jkR_0}}{r}\int_C \mathbf{I}dl \tag{A.4.1}$$

Note that since the currents are all in the plane perpendicular to r, $A'_r(\theta,\phi) = 0$, yielding

$$\mathbf{A} = \frac{e^{-jkR_0}}{r}\left[A'_\theta(\theta,\phi)\hat{\mathbf{u}}_\theta + A'_\phi(\theta,\phi)\hat{\mathbf{u}}_\phi\right] \tag{A.4.2}$$

and plugging into (A.2.12), the final mapping from currents on a plane to the broadside far-field electric field is

$$\mathbf{E} = \frac{-j\omega\mu}{4\pi}\frac{e^{-jkR_0}}{r}\int_C \mathbf{I}dl \tag{A.4.3}$$

Appendix B

Calculation of Polarization Ratio

All far-field electric fields can be represented with elliptical polarizations, with the two extreme cases of linear and circular polarization being the special cases when the length minor axis of the ellipse is 0 or equal to the length of the major axis respectively. This elliptical polarization can be expressed as the sum of two orthogonal polarizations, for example with linearly polarized fields $\mathbf{E_x}$ and $\mathbf{E_y}$. Similarly, circular polarizations $\mathbf{E_{CW}}$ and $\mathbf{E_{CCW}}$, the electric fields in the clockwise (CW) and counterclockwise (CCW) polarizations respectively, form an orthogonal basis, which means that any far-field electric field can be broken down into a CW and CCW component. The polarization ratio is defined as the ratio of the magnitude of the electric field in the CW polarization divided by the magnitude of the electric field in the CCW polarization. This means that fields with polarization ratios near 0 are CCW circularly polarized, ratios near 1 are linearly polarized, and ratios much greater than 1 are CW circularly polarized. Thus any far-field electric field can be written as

$$\mathbf{E_{Total}} = \mathbf{E_{CW}} + \mathbf{E_{CCW}} \tag{B.0.1}$$

with phasors $\mathbf{E_{CW}}$ and $\mathbf{E_{CCW}}$ defined as

$$\mathbf{E_{CW}} = A_{CW}(\hat{\mathbf{u}}_{\mathbf{x}'} + j\hat{\mathbf{u}}_{\mathbf{y}'}) \tag{B.0.2}$$

$$\mathbf{E_{CCW}} = A_{CCW}(\hat{\mathbf{u}}_{\mathbf{x}'} - j\hat{\mathbf{u}}_{\mathbf{y}'}) \tag{B.0.3}$$

for some cartesian coordinates $\hat{\mathbf{u}}_{\mathbf{x}'}$ and $\hat{\mathbf{u}}_{\mathbf{y}'}$ (Fig. B.0.1), and phasor reference time t' such that the time of the maximum instantaneous magnitude of the electric field when $\mathbf{E_{CW}}$ and $\mathbf{E_{CCW}}$ add occurs at phases 0 and π and is pointed in the $\hat{\mathbf{u}}_{\mathbf{x}'}$ direction for phase 0. This also implies that the minimum magnitude of the field when $\mathbf{E_{CW}}$ and $\mathbf{E_{CW}}$ subtract occur at a phase $\pi/2$ and $3\pi/2$ and will point in the positive or negative $\hat{\mathbf{u}}_{\mathbf{y}'}$ direction, depending on the relative values of A_{CW} and A_{CCW}. Thus the polarization ratio is just A_{CW}/A_{CCW}.

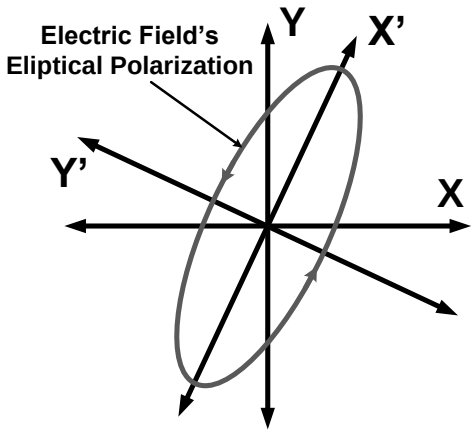

Figure B.0.1: Polarization diagram of a general electric field showing the change of coordinates from X and Y to X' and Y' aligned on the major and minor axes respectively.

The fields can also be written as a sum of linearly polarized phasor fields, $\mathbf{E}_{\mathbf{x}'}$ and $\mathbf{E}_{\mathbf{y}'}$ as

$$\mathbf{E}_{\mathbf{x}'} = (A_{CW} + A_{CCW})\hat{\mathbf{u}}_{\mathbf{x}'}, \text{and} \tag{B.0.4}$$

$$\mathbf{E}_{\mathbf{y}'} = j(A_{CW} - A_{CCW})\hat{\mathbf{u}}_{\mathbf{y}'}. \tag{B.0.5}$$

The linearly polarized receive antenna measures the linearly polarized projection of the electric field onto its major axis. By rotating the chip around the broadside axis, the receive antenna picks up the linearly polarized projection of the electric field at every angle. The maximum power received corresponds to the antenna aligning with the $\hat{\mathbf{u}}_{\mathbf{x}'}$ direction while the minimum corresponds to the antenna aligning with the $\hat{\mathbf{u}}_{\mathbf{y}'}$ direction. Those maximum and minimum received powers, $P_{x'}$ and $P_{y'}$ are proportional to the square of the electric field, and thus

$$\frac{P_{x'}}{P_{y'}} = \frac{(A_{CW} + A_{CCW})^2}{(A_{CW} - A_{CCW})^2}. \tag{B.0.6}$$

Solving for A_{CW}/A_{CCW} yields

$$\text{Polarization Ratio} = \frac{A_{CW}}{A_{CCW}} = \frac{\sqrt{\frac{P_x'}{P_y'}} + 1}{\sqrt{\frac{P_x'}{P_y'}} - 1}. \tag{B.0.7}$$

Appendix C

Tunable Transmission Lines

The ability to tune a passive matching network after fabrication can enable significant increases in performance and can combat degradation due to many factors including process variation, mismatch, aging effects and environmental variations such as load mismatch, to name a few. One possible way to tune a matching network is to tune the parameters of the transmission lines themselves.

A physical transmission line on an integrated circuit (IC) can be modeled well by its intrinsic impedance (Z_0, in Ω), its effective dielectric constant (ϵ_e), its physical length (d, in m), and its attenuation (A, in dB/m). Another parameter that is often used is the dielectric loss tangent (TanD), though for transmission lines on integrated circuits, the attenuation is the overwhelming source of loss in the circuit, and the TanD is neglected.

The goal of the design is to minimize the loss, so the attenuation will always be kept to a minimum. This means that actuators can be designed to tune Z_0, ϵ_e, and d. Tuning of the physical length is performance prohibitive for series lines, and tuning of the lengths of the stubs is discussed in detail in Chapter 7. This appendix will focus on ways to tune the other two parameters of the transmission lines, Z_0 and ϵ_e.

These parameters are related to the distributed inductance (L) and capacitance (C) through the following relations [22]:

$$Z_0 = \sqrt{\frac{L}{C}}, \tag{C.0.1}$$

and

$$\epsilon_e = \frac{LC}{\epsilon_0 \mu} \tag{C.0.2}$$

where ϵ_0 is the free space dielectric constant of $8.85 \times 10^{-12} F/m$ and μ is the permeability of the material. For the dielectrics within the scope of this chapter employed on standard IC processes, is equal to the permeability of free space of $4\pi \times 10^{-7} fracVsAm$.

This means that if we can independently control the inductance and capacitance of the transmission line, we can control both the intrinsic impedance and the effective dielectric constant independently.

C.1 Inductance Tuning of Transmission Lines

The inductance of a transmission line is heavily dependent on the distance between the signal current and its return path. By forcing the return current farther away, the inductance is increased, while allowing the current to return closer to the signal will decrease that current. This is the principle behind slotted slow wave transmission lines, where the return current of a micro-strip line is forced by slots in the ground plane to travel further from the signal than a normal micro-strip, while the capacitance is maintained by strips of metal between the slots. This has the effect of increasing the inductance while maintaining the capacitance, which from Equation C.0.2, increases the effective dielectric constant. We also can see from Equation C.0.1 that it will change the intrinsic impedance of the line, but in general it is easier to modify the impedance of the line as it is based on its physical dimensions, whereas the dielectric constant is a material property.

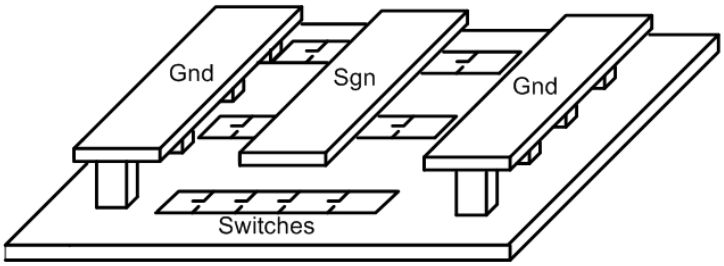

Figure C.1.1: 3D cartoon of the tunable slow wave transmission line, where slots in the ground plane force the return current to run through the side grounds unless switches within the slots are shorted, enabling the return current to flow directly below the signal line.

One possible way to make the inductance of the transmission line tunable is to utilize a slotted slow wave transmission line, but to put digital switches between the slots that, when on, allow the current to flow like a normal microstrip line, but that when off, switch the transmission line into a slow wave transmission line, shown in Figure C.1.1. This is a tricky proposition, as the opposite sides of the digital switch are both hooked up to the ground plane, so the switch needs to provide a lower impedance short circuit when on than the metal path does going around the slot. The physical size of the transistor switch cannot be larger than the slot itself, or the metal resistances and inductances routing to the transistor will be larger than the path going around the slot. Thus, transistors are chosen to entirely fill the physical size of the slot, yielding the least inductance and resistance when the switch is on. By switching on different numbers of transistors, various slots can be shorted out, and thus various inductance values can be achieved.

C.2 Capacitance Tuning of Transmission Lines

The capacitance of a transmission line is shunt capacitance; that is, it is a capacitance between the signal and ground that is distributed along the length of the line. To tune the value of this capacitance, additional capacitors can be switched in along the line. It is important that the capacitors be separated so that the bulk of the return current flows through the same path regardless of the capacitor tuning value, and thus the inductance remains relatively unchanged.

One way of adding this capacitance is to add large capacitors that tap the transmission line at evenly spaced points at large intervals down the line. This can be useful as the inductance of the tap line reduces the self-resonant frequency of the shunt capacitor, so if that inductance is fixed, larger capacitors will have higher self-resonance frequencies. However, the bulk of this inductance is because these capacitors are placed outside of the transmission line itself, and thus can be reduced greatly by placing the capacitors within the transmission line itself. This means that the transistors must be smaller, and to get the same tuning range, must be placed near each other.

In a microstrip line, the capacitance can be modified by switching islands of metal that are between the signal and the ground lines to the ground connection as shown in Figure C.2.1. This locally increases the capacitance, but because the islands on the middle metal layer are not connected, the bulk of the return current still flows through the lower ground metal, maintaining the inductance. Similar to the tunable inductance case, different numbers of capacitors can be switched on, and thus various distributed capacitances along the line can be achieved.

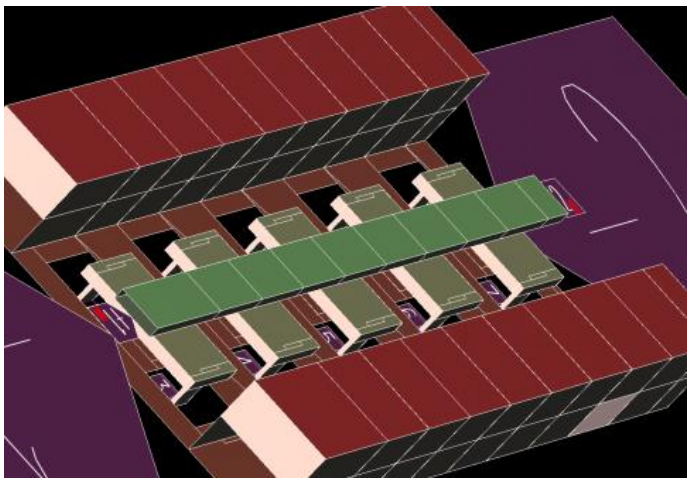

Figure C.2.1: 3-D rendering of the tunable capacitance transmission line, where islands of floating metal within the transmission line can be switched in to modify the distributed capacitance on the line.

Both of these methods are compatible with each other, and when combined, can offer two-dimensional tuning of both the capacitance and inductance of a transmission line.

C.3 Tunable Transmission Line Test Structures

Test structures of the tunable transmission lines were created in both a 65 nm bulk CMOS as well as a 45 nm SOI CMOS process. In the case of the 65 nm CMOS process, both inductance and capacitance tuning transmission lines are implemented. The lines are 'tub' type hybrid coplaner/microstrip lines, that either have a switchable slotted slow wave transmission line for inductance tuning, or switchable metal islands between the signal and ground lines for capacitance tuning. The lower ground is implemented on the 1st and 2nd metal layers, with the signal on the 9th (top) metal layer.

The capacitors are implemented half way between the signal and ground lines on the 7th metal layer, They are attached to the drain of a 20 μm width, minimum length transistor, whose source is tied to the local ground, and gate is controlled by the tuning line. The layout for these test structures is shown in Figure C.3.1.

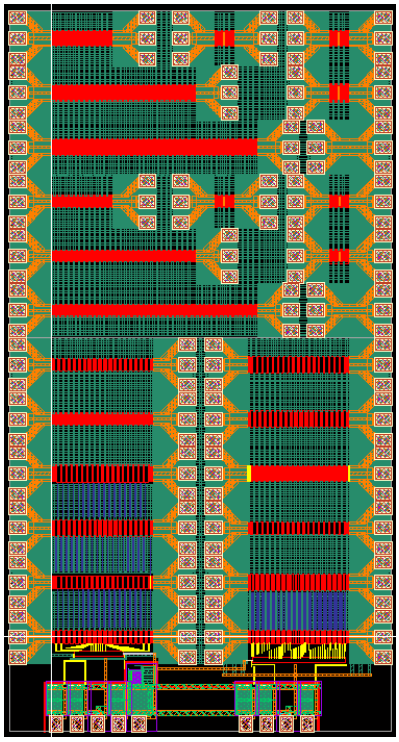

Figure C.3.1: Layout of the tunable transmission line test structure block, with calibration structures at the top, and the tunable transmission lines in the middle, and digital shift registers to control the transmission lines at the bottom.

A second set of capacitive tuning transmission line test structures are implemented in 45 nm SOI CMOS process. These lines are 8k μm wide, with a space between the signal and ground of 10 μm. The capacitor islands are 8 μm by 18 μm, placed every 20 μm, and switched to ground with a 2q width minimum length transistor. A layout of the tunable capacitance transmission line test structure is shown in Figure. C.3.2

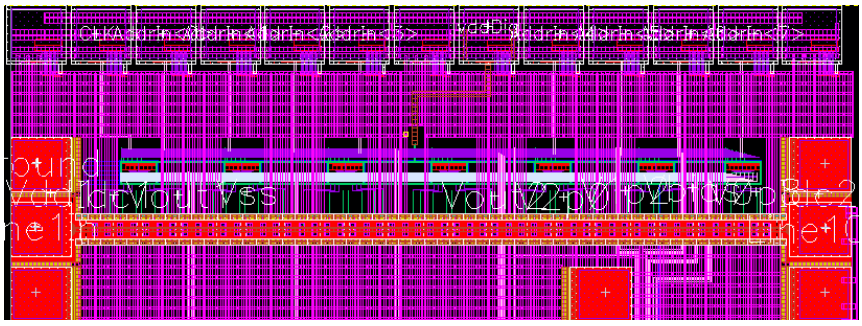

Figure C.3.2: Layout of the 1.1 mm tunable capacitance transmission line test structure. The digital switches are controlled by 8 address blocks that are located next to the transmission line.

The lines are probed from 100 MHz through 110 GHz, and attached to a network analyzer to measure the 2 port S parameters. These S parameter files are then put into Agilent's Advance Design System (ADS) software, and a fit to a transmission line model is made. This model allows for a fitting of intrinsic impedance, dielectric constant, and attenuation. The length of the line is fixed at the actual physical length.

For the capacitive case, as more switches are turned on, the capacitance increases, and for the inductive case, as more switches are turned on, the inductance will decrease. Figure C.3.3 shows measurements of the intrinsic impedance of the three transmission lines against the percentage of switches turned on. This shows that the intrinsic impedance will go down as the number of switches turned on increases in both cases, as can be seen in the plot. The effective dielectric constant, however, is proportional to both L and C, so as a higher percentage of switches turns on, the dielectric constant goes down for inductive tuning, and up for capacitive tuning. This can be seen in the measurement results in Figure C.3.4. The tuning range of the impedance is up to 15%, and the dielectric constant can be tuned by up to 30%.

Lastly, the attenuation must be discussed. Transistor switches are not ideal and thus inherently have loss in the switches. In the case of the 65 nm bulk CMOS process, the loss increases from around 600 dB/m to over 1200 dB/m at 28 GHz in both tuning cases. For the 45nm SOI CMOS process, the loss increased from 730 dB/m to 1450 dB/m. As part of an output power combining matching network, this type of increase in attenuation is unacceptable, so while reasonable tuning of both the inductance and capacitance is shown, the cost in attenuation is prohibitive, and this type of transmission line tuning was omitted from the final design of the self-healing power amplifier.

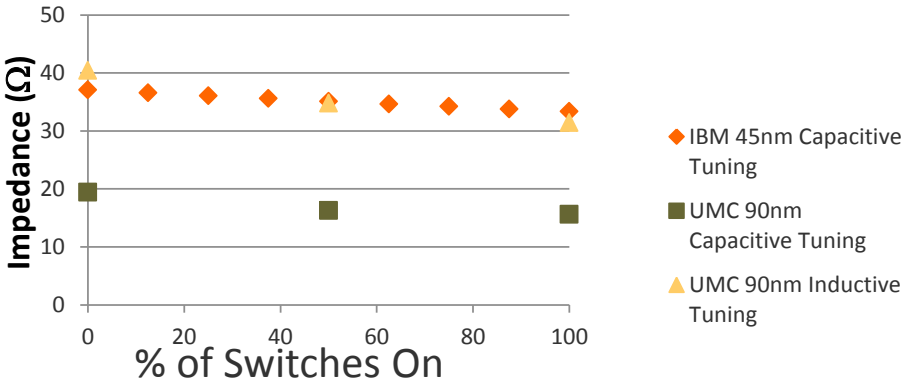

Figure C.3.3: Measured transmission line impedance as a function of the number of switches turned on for the three tunable transmission line test structures.

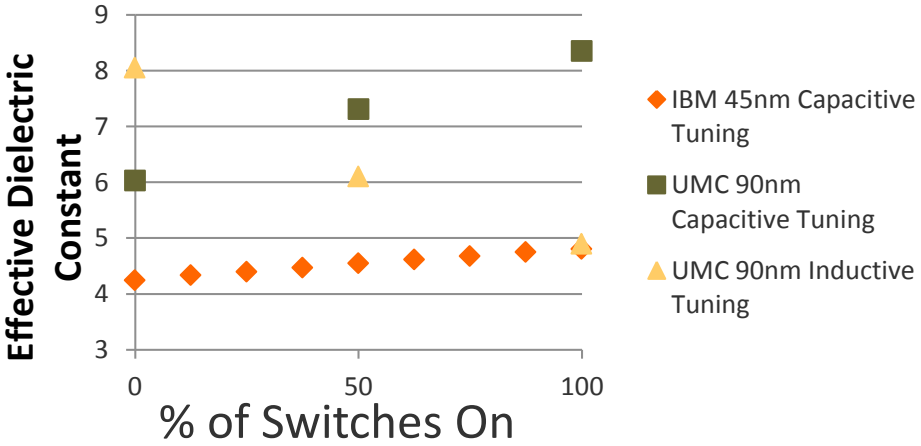

Figure C.3.4: Measured effective dielectric constant as a function of the number of switches turned on for the three tunable transmission line test structures.

Bibliography

[1] C. A. Balanis, *Antenna Theory: Analysis and Design*. Wiley-Interscience, 2005.

[2] D. B. Rutledge, D. P. Neikirk, and D. P. Kasilingam, "Integrated-circuit antennas," in *Infrared and Millimeter-Waves*. New York: Academic, 1983, pp. 1–90.

[3] G. Rebeiz, "Millimeter-wave and terahertz integrated circuit antennas," *Proceedings of the IEEE*, vol. 80, no. 11, pp. 1748–1770, 1992.

[4] A. Natarajan, A. Komijani, X. Guan, A. Babakhani, and A. Hajimiri, "A 77-GHz phased-array transceiver with on-chip antennas in silicon: Transmitter and local LO-path phase shifting," *IEEE J. of Solid-State Circuits*, vol. 41, no. 12, pp. 2807–2819, 2006.

[5] A. Babakhani, X. Guan, A. Komijani, A. Natarajan, and A. Hajimiri, "A 77-GHz phased-array transceiver with on-chip antennas in silicon: Receiver and antennas," *IEEE J. of Solid-State Circuits*, vol. 41, no. 12, pp. 2795–2806, 2006.

[6] E. Seok, C. Cao, D. Shim, D. Arenas, D. Tanner, C.-M. Hung, and K. O, "A 410GHz CMOS push-push oscillator with an on-chip patch antenna," in *IEEE Int. Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, 2008, pp. 472–629.

[7] S. Pan, L. Gilreath, P. Heydari, and F. Capolino, "Investigation of a wideband BiCMOS fully on-chip-band bowtie slot antenna," *IEEE Antennas and Wireless Propagation Letters*, vol. 12, pp. 706–709, 2013.

[8] J.-D. Park, S. Kang, S. Thyagarajan, E. Alon, and A. Niknejad, "A 260 GHz fully integrated CMOS transceiver for wireless chip-to-chip communication," in *Symp. On VLSI Circuits (VLSIC)*, 2012, pp. 48–49.

[9] Y.-C. Ou and G. Rebeiz, "Differential microstrip and slot-ring antennas for millimeter-wave silicon systems," *IEEE Trans. on Antennas and Propagation*, vol. 60, no. 6, pp. 2611–2619, 2012.

[10] S. Nicolson, A. Tomkins, K. Tang, A. Cathelin, D. Belot, and S. Voinigescu, "A 1.2V, 140GHz receiver with on-die antenna in 65nm CMOS," in *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, 2008, pp. 229–232.

[11] Z. Chen, C.-C. Wang, H.-C. Yao, and P. Heydari, "A BiCMOS W-band 2x2 focal-plane array with on-chip antenna," *IEEE J. of Solid-State Circuits*, vol. 47, no. 10, pp. 2355–2371, 2012.

[12] N. Alexopoulos, P. Katehi, and D. Rutledge, "Substrate optimization for integrated circuit antennas," in *IEEE Trans. on Microwave Theory and Techniques*, 1982, pp. 190–192.

[13] M. Kominami, D. M. Pozar, and D. Schaubert, "Dipole and slot elements and arrays on semi-infinite substrates," *IEEE Trans. on Antennas and Propagation*, vol. 33, no. 6, pp. 600–607, 1985.

[14] D. M. Pozar, "Considerations for millimeter wave printed antennas," *IEEE Trans. on Antennas and Propagation*, vol. 31, no. 5, pp. 740–747, 1983.

[15] W. Dou and Z. Sun, "Surface wave fields and power in millimeter wave integrated dipole antennas," *International Journal of Infrared and Millimeter Waves*, vol. 18, no. 3, pp. 711–721, 1997. [Online]. Available: http://dx.doi.org/10.1007/BF02678007

[16] G. Eleftheriades and M. Qiu, "Efficiency and gain of slot antennas and arrays on thick dielectric substrates for millimeter-wave applications: a unified approach," *IEEE Trans. on Antennas and Propagation*, vol. 50, no. 8, pp. 1088–1098, 2002.

[17] P. Katehi and N. Alexopoulos, "On the effect of substrate thickness and permittivity on printed circuit dipole properties," *IEEE Trans. on Antennas and Propagation*, vol. 31, no. 1, pp. 34–39, 1983.

[18] R. Han and E. Afshari, "A 260ghz broadband source with 1.1mw continuous-wave radiated power and eirp of 15.7dbm in 65nm cmos," in *IEEE International Solid-State Circuits Conference Digest*, 2013, pp. 138–139.

[19] D. Filipovic, S. Gearhart, and G. Rebeiz, "Double-slot antennas on extended hemispherical and elliptical silicon dielectric lenses," *IEEE Trans. on Microwave Theory and Techniques*, vol. 41, no. 10, pp. 1738–1749, 1993.

[20] W. Shin, B.-H. Ku, O. Inac, Y.-C. Ou, and G. Rebeiz, "A 108-114 ghz 4x4 wafer-scale phased array transmitter with high-efficiency on-chip antennas," *Solid-State Circuits, IEEE Journal of*, vol. 48, no. 9, pp. 2041–2055, 2013.

[21] K. Sengupta, *Silicon-based terahertz circuits and systems*. Dissertation (Ph.D.), California Institute of Technology., 2012. [Online]. Available: http://resolver.caltech.edu/CaltechTHESIS:06112012-145654043

[22] D. M. Pozar, *Microwave Engineering*. Wiley, 3rd edition, 2004.

[23] J. D. Krauss, *Antennas*. McGraw-Hill, New York, 1988.

[24] S. Drabowitch, A. Papiernik, H. D. Griffiths, J. Encinas, and B. L. Smith, *Modern Antennas*. Springer, 2nd edition, Netherlands, 2005.

[25] B. Razavi, "A millimeter-wave cmos heterodyne receiver with on-chip lo and divider," *IEEE J. of Solid-State Circuits*, vol. 43, no. 2, pp. 477 –485, Feb. 2008.

[26] M. Varonen, M. Karkkainen, M. Kantanen, and K. Halonen, "Millimeter-wave integrated circuits in 65-nm cmos," *IEEE J. of Solid-State Circuits*, vol. 43, no. 9, pp. 1991 –2002, Sept. 2008.

[27] I. Akoi, S. D. Kee, D. B. Rutledge, and A. Hajimiri, "Distributed active transformer-a new power-combining and impedance-transformation technique," *IEEE Trans. on Microwave Theory and Techniques*, vol. 50, no. 1, pp. 316 –331, Jan 2002.

[28] I. Aoki, S. D. Kee, D. B. Rutledge, and A. Hajimiri, "Fully integrated CMOS power amplifier design using the distributed active-transformer architecture," *IEEE J. of Solid-State Circuits*, vol. 37, no. 3, pp. 371 –383, Mar 2002.

[29] A. Komijani, A. Natarajan, and A. Hajimiri, "A 24-GHz, +14.5-dBm fully integrated power amplifier in 0.18- μm CMOS," *IEEE J. of Solid-State Circuits*, vol. 40, no. 9, pp. 1901–1908, 2005.

[30] K. Sengupta and A. Hajimiri, "A compact self-similar power combining topology," in *IEEE MTT-S International Microwave Symposium Digest (MTT)*, May 2010, pp. 244 –247.

[31] R. Gupta, B. Ballweber, and D. Allstot, "Design and optimization of CMOS RF power amplifiers," *IEEE J. of Solid-State Circuits*, vol. 36, no. 2, pp. 166 –175, Feb 2001.

[32] G. Liu, T.-J. K. Liu, and A. Niknejad, "A 1.2V, 2.4GHz fully integrated linear CMOS power amplifier with efficiency enhancement," in *IEEE Custom Integrated Circuits Conference*, Sept. 2006, pp. 141 –144.

[33] K. Sengupta and A. Hajimiri, "Sub-THz beam-forming using near-field coupling of distributed active radiator arrays," in *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, 2011, pp. 1–4.

[34] K. Sengupta and A. Hajimiri, "A 0.28 THz power-generation and beam-steering array in cmos based on distributed active radiators," *IEEE J. of Solid-State Circuits*, vol. 47, no. 12, pp. 3013–3031, 2012.

[35] ——, "Distributed active radiation for thz signal generation," in *IEEE International Solid-State Circuits Conference Digest (ISSCC)*, 2011, pp. 288–289.

[36] ——, "A 0.28thz 4x4 power-generation and beam-steering array," in *IEEE International Solid-State Circuits Conference Digest (ISSCC)*, 2012, pp. 256–258.

[37] Z. Wang, P.-Y. Chiang, P. Nazari, C.-C. Wang, Z. Chen, and P. Heydari, "A 210GHz fully integrated differential transceiver with fundamental-frequency VCO in 32nm SOI CMOS," in *IEEE Int. Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, 2013, pp. 136–137.

[38] E. Laskin, K. Tang, K. H. K. Yau, P. Chevalier, A. Chantre, B. Sautreuil, and S. Voinigescu, "170-GHz transceiver with on-chip antennas in SiGe technology," in *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, 2008, pp. 637–640.

[39] S. Nassif, N. Mehta, and Y. Cao, "A resilience roadmap," in *Design Automation Test in Europe Conference Exhibition*, 2010, pp. 1011–1016.

[40] K. Bernstein, D. J. Frank, A. E. Gattiker, W. Haensch, B. L. Ji, S. R. Nassif, E. J. Nowak, D. J. Pearson, and N. J. Rohrer, "High-performance CMOS variability in the 65-nm regime and beyond," *IBM J. of Research and Development*, vol. 50, no. 4.5, pp. 433 –449, July 2006.

[41] P. Stolk and D. Klaassen, "The effect of statistical dopant fluctuations on MOS device performance," in *International Electron Devices Meeting, IEDM*, Dec. 1996, pp. 627 –630.

[42] S. Borkar, "Designing reliable systems from unreliable components: the challenges of transistor variability and degradation," *IEEE Micro*, vol. 25, no. 6, pp. 10–16, 2005.

[43] K. Kuhn, "CMOS transistor scaling past 32nm and implications on variation," in *IEEE/SEMI Advanced Semiconductor Manufacturing Conference*, 2010, pp. 241–246.

[44] A. Asenov, "Simulation of statistical variability in nano MOSFETs," in *IEEE Symposium on VLSI Technology*, 2007, pp. 86–87.

[45] T. Mizuno, J. Okumtura, and A. Toriumi, "Experimental study of threshold voltage fluctuation due to statistical variation of channel dopant number in MOSFET's," *IEEE Trans. on Electron Devices*, vol. 41, no. 11, pp. 2216 –2221, Nov 1994.

[46] M. Ruberto, O. Degani, S. Wail, A. Tendler, A. Fridman, and G. Goltman, "A reliability-aware RF power amplifier design for CMOS radio chip integration," in *IEEE International Reliability Physics Symposium,*, 2008, pp. 536–540.

[47] O. Hammi, J. Sirois, S. Boumaiza, and F. Ghannouchi, "Study of the output load mismatch effects on the load modulation of doherty power amplifiers," in *IEEE Radio and Wireless Symposium*, 2007, pp. 393–394a.

[48] A. Keerti and A.-V. H. Pham, "RF characterization of SiGe HBT power amplifiers under load mismatch," *IEEE Trans. on Microwave Theory and Techniques*, vol. 55, no. 2, pp. 207 –214, Feb. 2007.

[49] H. Hashemi, X. Guan, A. Komijani, and A. Hajimiri, "A 24-GHz SiGe phased-array receiver-LO phase-shifting approach," *IEEE Trans. on Microwave Theory and Techniques*, vol. 53, no. 2, pp. 614 –626, Feb. 2005.

[50] A. Natarajan, A. Komijani, and A. Hajimiri, "A 24 GHz phased-array transmitter in 0.18 μm CMOS," in *IEEE Int. Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb. 2005, pp. 212 –594 Vol. 1.

[51] S. Alalusi and R. Brodersen, "A 60GHz phased array in CMOS," in *IEEE Custom Integrated Circuits Conference (CICC)*, 2006, pp. 393–396.

[52] A. Valdes-Garcia, S. Nicolson, J.-W. Lai, A. Natarajan, P.-Y. Chen, S. Reynolds, J.-H. Zhan, D. Kam, D. Liu, and B. Floyd, "A fully integrated 16-element phased-array transmitter in SiGe BiCMOS for 60-GHz communications," *IEEE J. of Solid-State Circuits*, vol. 45, no. 12, pp. 2757–2773, 2010.

[53] S. M. Bowers, K. Sengupta, K. Dasgupta, B. D. Parker, and A. Hajimiri, "Integrated self-healing for mm-wave power amplifiers," *IEEE Trans. on Microwave Theory and Techniques*, vol. 61, no. 3, pp. 1301–1315, 2013.

[54] F. Bohn, K. Dasgupta, and A. Hajimiri, "Closed-loop spurious tone reduction for self-healing frequency synthesizers," in *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, June 2011, pp. 1 –4.

[55] A. Tang, F. Hsiao, D. Murphy, I.-N. Ku, J. Liu, S. D'Souza, N.-Y. Wang, H. Wu, Y.-H. Wang, M. Tang, G. Virbila, M. Pham, D. Yang, Q. Gu, Y.-C. Wu, Y.-C. Kuan, C. Chien, and M. Chang, "A low-overhead self-healing embedded system for ensuring high yield and long-term sustainability of 60GHz 4Gb/s radio-on-a-chip," in *IEEE Int. Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb. 2012, pp. 316 –318.

[56] J.-C. Liu, R. Berenguer, and M. Chang, "Millimeter-wave self-healing power amplifier with adaptive amplitude and phase linearization in 65-nm CMOS," *IEEE Tran. on Microwave Theory and Techniques*, vol. 60, no. 5, pp. 1342–1352, 2012.

[57] J.-C. Liu, A. Tang, N. Wang, Q. Gu, R. Berenguer, H. Hsieh, P. Wu, C. Jou, and M. C. F. Chang, "A V-band self-healing power amplifier with adaptive feedback bias control in 65 nm CMOS," in *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, 2011, pp. 1–4.

[58] S. Sleiman, A. Akour, W. Khalil, and M. Ismail, "Millimeter-wave bist and bisc using a high-definition sub-ranged detector in 90nm cmos," in *IEEE International Midwest Symposium on Circuits and Systems (MWSCAS),*, 2010, pp. 477–480.

[59] S. Y. Kim, O. Inac, C.-Y. Kim, and G. Rebeiz, "A 76-84 ghz 16-element phased array receiver with a chip-level built-in-self-test system," in *Radio Frequency Integrated Circuits Symposium (RFIC), 2012 IEEE*, 2012, pp. 127–130.

[60] E. Cohen, M. Ruberto, M. Cohen, O. Degani, S. Ravid, and D. Ritter, "A cmos bidirectional 32-element phased-array transceiver at 60ghz with ltcc antenna," in *Radio Frequency Integrated Circuits Symposium (RFIC), 2012 IEEE*, 2012, pp. 439–442.

[61] Y. C. Huang, H. Hsieh, and L. Lu, "A low-noise amplifier with integrated current and power sensors for RF BIST applications," in *25th IEEE VLSI Test Symposium*, 2007, pp. 401–408.

[62] S. M. Bowers, K. Sengupta, K. Dasgupta, and A. Hajimiri, "A fully-integrated self-healing power amplifier," in *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, June 2012, pp. 221 –224.

[63] K. Sengupta and A. Hajimiri, "A compact self-similar power combining topology," in *IEEE MTT-S International Microwave Symposium Digest*, 2010, pp. 244–247.

[64] C. Law and A.-V. Pham, "A high-gain 60GHz power amplifier with 20dBm output power in 90nm CMOS," in *IEEE Int. Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, 2010, pp. 426–427.

[65] R. Wanner, R. Lachner, and G. Olbrich, "A monolithically integrated 190-GHz SiGe push-push oscillator," *IEEE Microwave and Wireless Components Letters*, vol. 15, no. 12, pp. 862–864, 2005.

[66] S. Shahramian, Y. Baeyens, N. Kaneda, and Y.-K. Chen, "A 70-100 GHz direct-conversion transmitter and receiver phased array chipset demonstrating 10 Gb/s wireless link," *IEEE J. of Solid-State Circuits*, vol. 48, no. 5, pp. 1113–1125, 2013.

[67] T. Krems, W. Haydl, H. Massler, and J. Rudiger, "Millimeter-wave performance of chip interconnections using wire bonding and flip chip," in *IEEE MTT-S International Microwave Symposium Digest*, vol. 1, 1996, pp. 247–250 vol.1.

[68] Z. Feng, W. Zhang, B. Su, K. Gupta, and Y. Lee, "Rf and mechanical characterization of flip-chip interconnects in cpw circuits with underfill," *IEEE Trans. on Microwave Theory and Techniques*, vol. 46, no. 12, pp. 2269–2275, 1998.

[69] A. Tang and M.-C. Chang, "183ghz 13.5mw/pixel cmos regenerative receiver for mm-wave imaging applications," in *IEEE International Solid-State Circuits Conference Digest (ISSCC)*, 2011, pp. 296–298.

[70] R. Han and E. Afshari, "A 260GHz broadband source with 1.1mW continuous-wave radiated power and EIRP of 15.7dBm in 65nm CMOS," in *IEEE Int. Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, 2013, pp. 138–139.

[71] Y. Zhao, J. Grzyb, and U. Pfeiffer, "A 288-GHz lens-integrated balanced triple-push source in a 65-nm CMOS technology," in *Proceedings of the ESSCIRC*, 2012, pp. 289–292.

[72] A. Hajimiri, "mm-wave silicon ICs: An opportunity for holistic design," in *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, 2008, pp. 357–360.

[73] S. M. Bowers and A. Hajimiri, "An integrated multi-port driven radiating source," in *IEEE MTT-S International Microwave Symposium Digest*, 2013.

[74] G. Rebeiz, "Millimeter-wave and terahertz integrated circuit antennas," *Proceedings of the IEEE*, vol. 80, no. 11, pp. 1748–1770, 1992.

[75] O. Momeni and E. Afshari, "High power terahertz and millimeter-wave oscillator design: A systematic approach," *IEEE J. of Solid-State Circuits*, vol. 46, no. 3, pp. 583–597, 2011.

[76] N. Behdad and K. Sarabandi, "A wide-band slot antenna design employing a fictitious short circuit concept," *IEEE Trans. on Antennas and Propagation*, vol. 53, no. 1, pp. 475–482, 2005.

[77] S. Raman and G. M. Rebeiz, "Single- and dual-polarized millimeter-wave slot-ring antennas," *IEEE Trans. on Antennas and Propagation*, vol. 44, no. 11, pp. 1438–1444, 1996.

[78] S. M. Bowers and A. Hajimiri, "Multi-port driven radiators," *paper accepted in IEEE Trans. on Microwave Theory and Techniques*.

[79] B. Jalali and S. Fathpour, "Silicon photonics," *Journal of Lightwave Technology*, vol. 24, no. 12, pp. 4600–4615, 2006.

[80] R. A. Soref and J. Larenzo, "All-silicon active and passive guided-wave components for l = 1.3 and 1.6um," *Quantum Electronics, IEEE Journal of*, vol. 22, no. 6, pp. 873–879, 1986.

[81] P. D. Trinh, S. Yegnanarayanan, and B. Jalali, "Integrated optical directional couplers in silicon-on-insulator," *Electronics Letters*, vol. 31, no. 24, pp. 2097–2098, 1995.

[82] C. Gunn, "Cmos photonics for high-speed interconnects," *Micro, IEEE*, vol. 26, no. 2, pp. 58–66, 2006.

[83] M. Hochberg and T. Baehr-Jones, "Towards fabless silicon photonics," *Nature Photonics*, vol. 4, pp. 492–494, 2010.

[84] Y. Liu, R. Ding, M. Gould, T. Baehr-Jones, Y. Yang, Y. Ma, Y. Zhang, A.-J. Lim, T.-Y. Liow, S.-G. Teo, G.-Q. Lo, and M. Hochberg, "30ghz silicon platform for photonics system," in *Optical Interconnects Conference, 2013 IEEE*, 2013, pp. 27–28.

[85] R. Ding, T. Baehr-Jones, T. Pinguet, J. Li, N. Harris, M. Streshinsky, L. He, A. Novack, A.-J. Lim, T. Liow, S.-G. Teo, G.-Q. Lo, and M. Hochberg, "A silicon platform for high-speed photonics systems," in *Optical Fiber Communication Conference and Exposition (OFC/NFOEC), 2012 and the National Fiber Optic Engineers Conference*, 2012, pp. 1–3.

[86] M. Hochberg, N. Harris, R. Ding, Y. Zhang, A. Novack, Z. Xuan, and T. Baehr-Jones, "Silicon photonics: the next fabless semiconductor industry," *Solid-State Circuits Magazine, IEEE*, vol. 5, no. 1, pp. 48–58, 2013.

[87] T. T. H. Eng, S. Sin, S. C. Kan, and G. K. L. Wong, "Surface-micromachined movable soi optical waveguides," in *The 8th International Conference on Solid-State Sensors and Actuators, 1995.*, vol. 1, 1995, pp. 348–350.

[88] Y. Vlasov and S. McNab, "Losses in single-mode silicon-on-insulator strip waveguides and bends," *Opt. Express*, vol. 12, pp. 1622–1631, 2004.

[89] P. J. Foster, J. Doylend, P. Mascher, A. Knights, and P. G. Coleman, "Optical attenuation in defect-engineered silicon rib waveguides," *Journal of Applied Physics*, vol. 99, no. 7, pp. 073 101–073 101–7, 2006.

[90] K. Giboney, M. J. W. Rodwell, and J. Bowers, "Traveling-wave photodetector theory," *Microwave Theory and Techniques, IEEE Transactions on*, vol. 45, no. 8, pp. 1310–1319, 1997.

[91] M. Hochberg, T. Baehr-Jones, G. Wang, J. Huang, P. Sullivan, L. Dalton, and A. Scherer, "Towards a millivolt optical modulator with nano-slot waveguides," *Opt. Express*, vol. 15, pp. 8401–8410, 2007.

[92] D. Taillaert, W. Bogaerts, P. Bienstman, T. Krauss, P. van Daele, I. Moerman, S. Verstuyft, K. De Mesel, and R. Baets, "An out-of-plane grating coupler for efficient butt-coupling between compact planar waveguides and single-mode fibers," *Quantum Electronics, IEEE Journal of*, vol. 38, no. 7, pp. 949–955, 2002.

[93] U. Pfeiffer and D. Goren, "A 20 dBm fully-integrated 60 GHz SiGe power amplifier with automatic level control," *IEEE J. of Solid-State Circuits*, vol. 42, no. 7, pp. 1455–1463, 2007.

[94] E. Maricau and G. Gielen, "Transistor aging-induced degradation of analog circuits: Impact analysis and design guidelines," in *Proceedings of ESSCIRC*, Sept. 2011, pp. 243 –246.

[95] R. Kumar and V. Kursun, "Voltage optimization for temperature variation insensitive CMOS circuits," in *48th Midwest Symposium on Circuits and Systems,*, 2005, pp. 476–479 Vol. 1.

[96] S. Sakurai and M. Ismail, "Robust design of rail-to-rail CMOS operational amplifiers for a low power supply voltage," *IEEE Journal of Solid-State Circuits*, vol. 31, no. 2, pp. 146–156, 1996.

[97] K. Siwiec, T. Borejko, and W. Pleskacz, "PVT tolerant LC-VCO in 90 nm CMOS technology for GPS/galileo applications," in *IEEE International Symposium on Design and Diagnostics of Electronic Circuits Systems*, 2011, pp. 29–34.

[98] D. Sylvester, D. Blaauw, and E. Karl, "ElastIC: an adaptive self-healing architecture for unpredictable silicon," *IEEE Design Test of Computers*, vol. 23, no. 6, pp. 484–490, 2006.

[99] A. Goyal, M. Swaminathan, A. Chatterjee, D. Howard, and J. Cressler, "A new self-healing methodology for RF amplifier circuits based on oscillation principles," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 20, no. 10, pp. 1835–1848, 2012.

[100] S. Yaldiz, V. Calayir, X. Li, L. Pileggi, A. Natarajan, M. Ferriss, and J. Tierno, "Indirect phase noise sensing for self-healing voltage controlled oscillators," in *IEEE Custom Integrated Circuits Conference (CICC)*, 2011, pp. 1–4.

[101] K. Jayaraman, Q. Khan, B. Chi, W. Beattie, Z. Wang, and P. Chiang, "A self-healing 2.4GHz LNA with on-chip S11/S21 measurement/calibration for in-situ PVT compensation," in *IEEE Radio Frequency Integrated Circuits Symposium*, 2010, pp. 311–314.

[102] H. Wang, K. Dasgupta, and A. Hajimiri, "A broadband self-healing phase synthesis scheme," in *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, June 2011, pp. 1 –4.

[103] T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits, 2nd Ed.* Cambridge University Press, 2003.

[104] K. Sengupta, K. Dasgupta, S. M. Bowers, and A. Hajimiri, "On-chip sensing and actuation methods for integrated self-healing mm-wave CMOS power amplifier," in *IEEE MTT-S International Microwave Symposium Digest (MTT)*, June 2012, pp. 1 –3.

[105] Q. Yin, W. Eisenstadt, R. Fox, and T. Zhang, "A translinear RMS detector for embedded test of RF ICs," *IEEE Trans. on Instrumentation and Measurement*, vol. 54, no. 5, pp. 1708–1714, 2005.

[106] C. de La Cruz-Blas, A. Lopez-Martin, A. Carlosena, and J. Ramirez-Angulo, "1.5-V current-mode CMOS true RMS-DC converter based on class-AB transconductors," *IEEE Trans. on Circuits and Systems II : Express Briefs*, vol. 52, no. 7, pp. 376–379, 2005.

[107] A. Valdes-Garcia, R. Venkatasubramanian, J. Silva-Martinez, and E. Sanchez-Sinencio, "A broadband CMOS amplitude detector for on-chip RF measurements," *IEEE Transactions on Instrumentation and Measurement*, vol. 57, no. 7, pp. 1470–1477, 2008.

[108] S. Cripps, *RF power amplifiers for wireless communications*. Atrech House Microwave Library, 2006.

[109] Y. Huang, H. Hsieh, and L. Lu, "A build-in self-test technique for RF low-noise amplifiers," *IEEE Trans. on Microwave Theory and Techniques*, vol. 56, no. 5, pp. 1035–1042, 2008.

[110] F. Cheung and P. Mok, "A monolithic current-mode CMOS DC-DC converter with on-chip current-sensing technique," *IEEE J. of Solid-State Circuits*, vol. 39, no. 1, pp. 3–14, 2004.

[111] S. Lee and D. Allstot, "Electrothermal simulation of integrated circuits," *IEEE J. of Solid-State Circuits*, vol. 28, no. 12, pp. 1283–1293, 1993.

[112] S. Devarakond, V. Natarajan, S. Sen, and A. Chatterjee, "BIST-assisted power aware self healing RF circuits," in *IEEE International Mixed-Signals, Sensors, and Systems Test Workshop*, 2009, pp. 1–4.

[113] S. Akbay and A. Chatterjee, "Built-in test of RF components using mapped feature extraction sensors," in *Proceedings of IEEE VLSI Test Symposium*, 2005, pp. 243–248.

[114] "Basics on load pull and noise measurements," in *Focus Microwaves Application Note 8*, Jun 1994.

[115] C. Tsironis, R. Meierer, B. Hosein, T. Beauchamp, and R. Jallad, "MPT, a universal multi-purpose tuner," in *65th ARFTG Conference Digest*, June 2005, pp. 113 – 117.