

Scalable Arrays From Millimeter-Wave Sensing to Microwave Wireless Power Transfer

Thesis by
Alex Eben Ayling

In Partial Fulfillment of the Requirements for the
degree of
Doctor of Philosophy



CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2026
Defended March 3rd, 2026

© 2026

Alex Eben Ayling

ORCID: 0009-0008-5440-7785

All rights reserved except where otherwise noted

ACKNOWLEDGEMENTS

First, I would like to thank my advisor, Ali Hajimiri. Ali was not only a wonderful advisor, but a good friend. Ali always pushed me to think about big problems and in doing so to push myself. He gave me his support and confidence in tackling those problems, including ones where success was quite important, even though I had never demonstrated the ability to handle them before. That allowed me to grow as a researcher and as a human, something else he cares deeply about, and I'm grateful to have had those opportunities in his lab.

I would also like to thank my undergraduate advisor, Keshab Parhi, and my M.S. advisor, Elyse Rosenbaum. In working with both, I took the first steps to learn the ropes of research. That gave me the foundation that I needed before starting at Caltech.

I would like to thank all the CHIC lab members I have interacted with, an exhaustive list being: Austin Fikes, Ailec Wu, Oren Mizrahi, Raha Riazati, Jesse Brunet, Samir Nooshabadi, Debjit Sarkar, Volkan Gurses, David Baum, Zachary Szekeres, Suraj Samaga, Xueru Zhou, Imaad Syed, Jatin Matur, Craig Ives, Matan Gal-Katziri, M. Reza Hashemi, Elliot Williams, Aroutin Khachaturian, Reza Fatemi, and Parham Porsandeh. It has been an honor to work with such a brilliant group of people, and to call you my friends. To all the CHIC lab members, thank you for putting up with my banter throughout the years. It's always nice to unwind a bit with some laughs.

I thank Professor Sergio Pellegrino and his group members Charles Sommer, Alan Truong, and Terry Gdoutos. Additionally, I want to thank Professor Harry Atwater and his group members Michael Kelzenberg and Phil Jahelka. I always learn from your talks during the biweekly SSPP meetings.

I had the pleasure of working with Damon Russell and Sharmila Padmanabhan, both at JPL at the time of our collaboration. They were always willing to go out of their way and help me, and I have found that to be the case with all the JPLers I have interacted with.

To my family Mama, Papa, Blake, Claire, Lindsay, Sammie, Johnny, and Grandma, thank you for supporting me during the last 8 years of grad school. Without your love I would not have made it.

To Amma, Nanna and all of my extended Indian relatives, being accepted into your

family has been one of the greatest honors of my life. I always feel at home with you.

Finally, I want to thank the most important person in my life, my wife Apurva. I met Apurva twelve years ago at the University of Minnesota, on the first day of my undergraduate studies in my first lecture. We have been together ever since. Apurva, you have been with me every step of this PhD, encouraging me to apply to Caltech and supporting me when times were tough. You are the kindest, sweetest, funniest wife I could ever ask for. Every day with you is a treasure, and marrying you is what I am most proud of in my life. I dedicate this PhD to you.

ABSTRACT

Wireless power transfer at a distance, long relegated to the realm of science fiction, has seen a resurgence in recent years. Chief among its promises is Space-Based Solar Power (SBSP), an ambitious project to deploy kilometer scale photovoltaic arrays in space and beam its power down to Earth using a complementary microwave phased array. The building blocks of the array are phased array tiles, which can be instantiated to produce larger apertures.

The tile must be simultaneously lightweight and flexible for deployment in space, low-cost, high-performance, and scalable. First, the results of the MAPLE mission, which tested wireless power transfer in space using custom flexible arrays, are presented. Using the results of that mission, the design and testing of next-generation, fully flexible 8x8 element phased array tile are presented. The tile is driven by a custom 22-nm CMOS FDSOI RFIC that achieves record efficiency and performance. These results represent not only a step forward toward practical microwave wireless power transfer but offer new directions in communications and sensing driven by flexible arrays.

Additionally, topics on maximum power point tracking in SBSP systems, transmitarrays for SBSP, and the design of a fully-integrated, scalable, and low-cost D-band (110-170GHz) radiator tile are discussed.

PUBLISHED CONTENT AND CONTRIBUTIONS

- [1] Alex Ayling and Ali Hajimiri. “Flip-Chip Aperture Coupled D-Band Active Radiator Tiles in 22-nm CMOS FDSOI”. In: *IEEE Journal of Solid-State Circuits* 60.6 (2025), pp. 1932–1946. DOI: 10.1109/JSSC.2024.3485062. Contribution: A. Ayling did all of the design work, measurements, and manuscript writing.
- [2] Alex Ayling et al. “Wireless power transfer in space using flexible, lightweight, coherent arrays”. In: *Acta Astronautica* 224 (2024), pp. 226–243. ISSN: 0094-5765. DOI: <https://doi.org/10.1016/j.actaastro.2024.08.006>. URL: <https://www.sciencedirect.com/science/article/pii/S009457652400434X>. Contribution: A. Ayling worked on conceptualization, data curation, formal analysis, investigation, methodology, resources, software, validation, visualization, writing – original draft, writing – review editing.
- [3] Jesse Brunet, Alex Ayling, and Ali Hajimiri. “Transmitarrays for Wireless Power Transfer on Earth and in Space”. In: *IEEE Journal of Microwaves* 4.4 (2024), pp. 934–945. DOI: 10.1109/JMW.2024.3459859. Contribution: A. Ayling did the hardware design, reviewed the cost analysis and assisted with manuscript writing.
- [4] Alex Ayling and Ali Hajimiri. “Fine Pitch D-Band Transmit Modules with Flip-Chip Aperture Coupled Antennas”. In: *2024 IEEE/MTT-S International Microwave Symposium - IMS 2024*. 2024, pp. 323–326. DOI: 10.1109/IMS40175.2024.10600369. Contribution: A. Ayling did all of the design work, measurements, and manuscript writing.
- [5] A. Fikes et al. “The Caltech Space Solar Power Demonstration One Mission”. In: *2022 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. 2022, pp. 18–22. DOI: 10.1109/WiSEE49342.2022.9926883. Contribution: A. Ayling assisted with the design of MAPLE and reviewed the manuscript.
- [6] Alex Ayling, Ailec Wu, and Ali Hajimiri. “Mutual Power Optimization of Photovoltaics and Wireless Power Transfer for Space Based Solar Power”. In: *2022 IEEE/MTT-S International Microwave Symposium - IMS 2022*. 2022, pp. 202–205. DOI: 10.1109/IMS37962.2022.9865358. Contribution: A. Ayling worked on the hardware design, measurements, and manuscript writing.
- [7] Alex Ayling et al. “A Flexible, Tileable Phased Array Enabled by CMOS Beamformers”. In: *Accepted to 2026 IEEE/MTT-S International Microwave Symposium - IMS 2026*. 2026.

Contribution: A. Ayling was the project lead, designed the VCDL and PA, and contributed to manuscript writing.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	v
Published Content and Contributions	vi
Table of Contents	vii
List of Illustrations	x
List of Tables	xviii
Nomenclature	xix
Chapter I: Introduction	1
Chapter II: Microwave Array for Power-transfer Low-orbit Experiment (MAPLE)	4
2.1 Introduction	4
2.2 System Design	7
2.3 In-Orbit: Experimental Results	20
2.4 Beam-to-Earth Experiments	25
2.5 Anomalous Results	35
2.6 Conclusion	43
Chapter III: Next Generation Flexible Wireless Power Transfer Array Using a Fully Integrated 22-nm CMOS FDSOI RFIC	45
3.1 Introduction	45
3.2 Specifications	46
3.3 22-nm CMOS FDSOI RFIC Design	49
3.4 RFIC Test Structure Measurements	59
3.5 Connectorized Testboard Measurements	60
3.6 8x8 Flexible Array Prototype	63
3.7 Conclusion	66
Chapter IV: Mutual Power Optimization	69
4.1 Introduction	69
4.2 Maximum Power Point Tracking Using a DC-DC Converter	70
4.3 PA Duty Cycling	73
4.4 Hardware Implementation	75
Chapter V: Transmitarrays for Wireless Power Transfer	80
5.1 Introduction	80
5.2 Cost Analysis	82
5.3 Transmitarray Prototype	91
5.4 Conclusion	92
Chapter VI: Scalable D-band Active Radiator Tiles in 22-nm CMOS FDSOI	96
6.1 Introduction	96
6.2 Antenna Design	100
6.3 Electronics Design	108
6.4 Tripler-Based Phase Shifter	113

6.5 Measurement Results	119
6.6 Conclusion	128
Bibliography	129
Appendix A: Timing Requirements for Large Arrays	139
Appendix B: Phase-Shifter Quantization Loss for Large Arrays	146
Appendix C: MPO Timing Requirements	149

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
2.1 The Caltech SBSP concept [1]: A PV-WPT array converts solar energy to DC power and DC power to RF power which is then beamed to Earth and collected by a ground station.	5
2.2 (a) CAD rendering of MAPLE with parts labeled, and (b) assembly. Dimensions of the cubesat are 10cm x 34.6cm x 22.6cm.	8
2.3 MAPLE block diagram, showing various components and redundancy. Flight computer is not part of MAPLE, but has parallel communication lines and power lines to both motherboards on the AP.	9
2.4 65-nm CMOS RFIC block diagram.	10
2.5 65-nm CMOS RFIC die photo.	11
2.6 Flexible array with popup dipoles (engineering model depicted), fixed to an aluminum frame for heatsinking and mechanical integration.	12
2.7 5×5 rectenna array of dual-polarized patch antenna receivers.	14
2.8 MAPLE integration on Vigoride-5 spacecraft from Momentus Inc. [36].	16
2.9 RFIC thermal simulation. Heat flows from the chip through the aluminum frame to the temperature boundary at the bottom face. Heatsink stackup is shown in the lower left (dashed box not displayed in other images.)	17
2.10 Texas A&M Radiation Effects Facility proton beam.	18
2.11 (a) Rectenna #2 (broadside) optimization on 2023/03/03 - 23:51 UTC, and (b) Rectenna #1 (sidewall) optimization on 2023/03/14 - 03:19 UTC.	21
2.12 Rectenna #2 focusing with only 16 elements controlled by (a) DB1 on 2023/03/21 - 23:16 UTC, and (b) DB2 on 2023/03/22 - 00:47 UTC.	22
2.13 Interior of MAPLE during a power beaming experiment on March 15th, 2023 – 4:40 AM UTC. Lit LED (center) indicates power transmission to the (sidewall) rectenna.	23

2.14	Raw rectified power from both on-Earth (black) and in-orbit (red) interference tests. On-Earth curves (black) represent the mean of six tests performed before launch; in-orbit curves (red) represent the first test performed in orbit. The label on the upper left indicates the “coordinates” of the core: (RFIC, quad, stack, core). The y-axes have limits of [0 20] mW.	24
2.15	Tracking long-term performance degradation by comparing two benchmark experiment runs eight months apart. Left: Power rectified during different phases of the benchmark experiment. “O” = optimization, “R” = recall, “I” = interference. “O1 (peak)” denotes the peak power during the first optimization. Right: A comparison of individual core amplitudes (A_i) at the beginning (x-axis) and end (y-axis) of the ten month deployment. Each data point corresponds to a core, labeled using the core number ($\in [1, 32]$) for the first core in the stack. Degraded cores circled in red.	25
2.16	MAPLE beam-to-earth far-field pattern	27
2.17	Dish antenna equatorial mount.	28
2.18	Ground station receiver block diagram. There is an additional 3dB of loss from the cabling between the LNA output and the mixer input.	29
2.19	Receiver electronics enclosure.	29
2.20	Receiver calibration setup. The receiver is located on the roof of Caltech’s Moore Laboratory. The target is located on the top floor of Caltech Hall (formerly Millikan Library).	31
2.21	Dish antenna with optical sight fixed at the top.	31
2.22	BTE #1 (5/23/2023) results. IF signal power in the upper right, Doppler shift on the lower right. Snapshots of MAPLE’s beam projected on the ground (orange ellipse) shown on the left with corresponding time.	33
2.23	BTE #2 (6/30/2023) results. IF signal power in the upper right, Doppler shift on the lower right. Snapshots of MAPLE’s beam projected on the ground (orange ellipse) shown on the left with corresponding time.	34
2.24	BTE #3 setup on Moore Laboratory rooftop.	35

2.25	BTE #3 (7/28/2023) results. PA current draw (top), IF signal power (middle), and Doppler shift plotted vs. time. Snapshots of MAPLE's beam projected on the ground (orange ellipse) shown at the bottom with corresponding time.	36
2.26	(a) On-Earth optimization sequence performed directly prior to launch, and (b) In-orbit sequence performed on March 3rd, 2023.	37
2.27	Supply current drop anomaly on DB1. Left: Supply current for RFIC1 (lines A and B) measured over time. Right: Supply current plotted against board temperature. t_{knee} , T_{knee} , and I_{equil} labeled. . . .	38
2.28	(a) On-Earth chip-chip interference test performed directly prior to launch, and (b) in-orbit test performed on March 21, 2023.	40
2.29	Long-term system degradation as indicated by two metrics, T_{knee} and I_{eq}	41
3.1	Available DC power per element vs. frequency.	48
3.2	Flexible wireless power transfer array driven by fully integrated CMOS RFICs.	50
3.3	PLL1 block diagram (a) and (b) PLL2 block diagram.	51
3.4	VCDL Block diagram.	53
3.5	VCDL delay-cell circuit schematic.	53
3.6	VCDL output phase vs. control voltage over temperature (a), and (b) across process skews.	54
3.7	VCDL phase shifter reference voltage generation with DLL (left) distributed to VCDL channel phase shifters (right.)	55
3.8	Self-calibrated DLL phase-shift curve vs. code (a) and minimum delay deviation over temperature (b).	57
3.9	Inverter-based LO distribution. Only the half-circuit is displayed in this figure, in actuality the distribution is differential.	57
3.10	PA and driver schematic. Driver is the common-source amplifier on the left and PA the cascode amplifier on the right.	58
3.11	Cascode core load-pull power contours (a) and efficiency contours (b). The red dot corresponds to the nominal impedance from the matching network at the drain.	58
3.12	Test structure chip die photo. Chip dimensions are 1.62mm x 1.1mm.	60
3.13	PTAT schematic (left) and PTAT voltage vs. temperature (right). . . .	61

3.14	PLL phase-noise measurement vs. simulation. Loop bandwidth setting refers to the digital code for adjusting loop-filter for optimal phase-noise.	61
3.15	Full chip die photo.	62
3.16	Interposer (left) and connectorized motherboard (right) photos. Interposer dimensions are 11mm x 11mm.	62
3.17	DLL phase shifter curve for channels 4-8.	63
3.18	(a) Phase noise at PA output. (b) Channel #4 VCDL RMS phase error with and without DLL calibration over temperature.	64
3.19	Plot of PA efficiency vs frequency. Prior art shown in blue (data from [50]), PA shown in red, and full RFIC in green.	64
3.20	Assembled 64-element flexible array front (left) and back (right). Active array size is 14.4cm x 14.4cm.	65
3.21	Azimuth and elevation far-field scans of 64-element array.	65
3.22	Power beaming in concave, flat, and convex shape configurations. . .	66
3.23	Near-field scanner power plot cuts at varying distances from array. The 3dB contour (in y-direction) is shown in the bottom plot.	67
3.24	Near-field scanner power plots of refocused beams. The first percentage is the refocusing efficiency with the DLL and the second without.	68
4.1	Typical IV curve of a photovoltaic cell. The point (V_{MPP} , I_{MPP}) is the maximum power point of the cell.	69
4.2	Maximum power point tracking with a buck converter.	70
4.3	PV IV curve with output resistor load-line. The gray shaded region of the plot represents the achievable impedances using a buck-converter. . .	71
4.4	22-nm CMOS FDSOI PA output power and efficiency vs. V_{dd}	72
4.5	22-nm CMOS FDSOI PA output power equivalent supply impedance vs V_{dd}	72
4.6	(a) Inductorless switching converter, and (b) equivalent circuit with very large C_d	73
4.7	(a) PA duty cycling equivalent supply impedance, and (b) incorporation in the PA.	75
4.8	Generating Unit (GU) and Receiving Unit (RU) block diagram. . . .	76
4.9	(a) PV emulator schematic, (b) measured IV curves of PV emulator, and (c) experimental setup.	76
4.10	MPO laboratory experiment results.	77

4.11	MPO outdoor experiment setup.	78
4.12	MPO laboratory experiment results.	78
5.1	Transmitarray wireless power transfer system. The transmitarray, acting as a lens, refocuses the divergent beam from the transmitter. . .	80
5.2	Cartoon showing (a) power transfer when transmitarray is aligned with direction of power transfer, (b) power transfer when transmitarray is not well-aligned with the direction of power transfer, and (c) power transfer with a transmitarray constellation—the active array can now beam through a transmitarray that is both closer and better aligned with the direction of power transfer.	83
5.3	Cartoon showing a full constellation including active arrays, transmitarrays, and rectenna arrays. As the transmitarrays move through their orbits, the active arrays continuously adjust their beams to focus on the nearest transmitarray.	84
5.4	Diffraction patterns of the active array at the transmitarray plane (middle) and of the transmitarray at the rectenna plane (bottom) for five points as the transmitarray moves through its orbit (top). The part of the active array's pattern that falls outside the transmitarray is greyed out; The power in the circle is reradiated by the transmitarray towards the rectenna array. In this case, $N = 9$ so that each transmitarray has an effective orbital span of 40°	87
5.5	Cost analysis results. A.A. = active array. T.A. = transmitarray. For each A.A. diameter we plot (a) the minimum achievable LCOE both with and without T.A.'s and (b) LCOE percent increase/decrease relative to the LCOE without T.A.'s.	90
5.6	The transmitarray (a) and rectenna array (b) diameters at minimum LCOE in Fig. 5.5a.	91
5.7	Transmitarray prototype stack-up. The outer layers are patch antennas with a slot in the middle for coupling.	92
5.8	Fabricated transmitarray prototype: individual layers (left) and assembled array (right)	93
5.9	Prototype transmitarray measurement setup.	93
5.10	Normalized power densities (a) without the transmitarray and (b) with the transmitarray. Intensity values are normalized to the maximum intensity measured <i>with</i> the transmitarray. Black ellipses correspond to the -3db contours of each plots respective peak.	94

5.11	Rectenna output voltage while focusing. The rectenna DC output is terminated by a 54Ω load in both cases.	95
6.1	Hemispherical Dielectric lens mounted to a silicon RFIC (left) and ray-tracing diagram (right) of slot antenna on the RFIC.	97
6.2	Aperture coupled integrated radiator.	97
6.3	Chip-tile and PCB assembly with flip-chip aperture coupled antennas. Four differential PAs excite four sets of differential slots on-chip. The slots drive off-chip patch antennas at $\lambda/2$ pitch fabricated on the PCB. Chips are tiled with narrow spacing to form large, dense arrays. . . .	99
6.4	Flip-chip aperture coupled antennas side-view.	100
6.5	Flip-chip aperture coupled antennas stackup.	100
6.6	2x2 radiator E-plane and H-plane gain simulations at 140GHz. . . .	101
6.7	2x2 radiator gain and front-side radiation vs. silicon substrate thickness at 140GHz.	103
6.8	2x2 radiator gain and radiation efficiency vs. bump height at 140GHz.	103
6.9	2x2 radiator gain and radiation efficiency vs. PCB substrate thickness at 140GHz.	104
6.10	2x2 radiator gain and efficiency vs slot length (a) and Smith chart impedance (b).	104
6.11	2x2 radiator gain degradation for varying x and y translations at 140GHz.	105
6.12	Flip-chip aperture coupled antenna with underfill.	106
6.13	2x2 radiator gain vs. patch length at 140GHz (a) and (b) gain vs. frequency (patch length = $610\mu\text{m}$).	106
6.14	Flip-chip aperture coupled antenna with underfill, stacked patch antennas, and backside reflector. Slot dimensions and patch width are the same as Fig. 6.3.	107
6.15	2x2 stacked patch radiator pattern (a) and gain vs. frequency (b). . . .	107
6.16	Chip block diagram and circuit schematics.	109
6.17	Channel block diagram and power breakdown.	109
6.18	PA input distribution and antenna feed.	110
6.19	PA output power contours (in dBm) load-pull simulation.	110
6.20	Tripler #1 and tripler #2 output power vs. tripler #1 gate bias. Simulation done using full RF chain in Fig. 6.16.	111
6.21	Full channel simulated PA output power and phase shift vs. tripler #1 gate bias (a), and output phase corner simulations (b).	111

6.22	Element isolation simulation.	112
6.23	Transistor schematic for drain current harmonic analysis.	114
6.24	Masking function modulates $V_{ov}(t)$ to produce the rectified drain current.	114
6.25	I_{int} and I_{up} magnitude (a) and phase (b) for $\theta = 0^\circ$, 45° , and 90° . $V_{AC} = 400\text{mV}$, $V_s = 50\text{mV}$, and θ displayed in graph.	116
6.26	Third harmonic magnitude (left) and phase (right) of the schematic in Fig. 6.23. $V_{AC} = 400\text{mV}$, $V_s = 50\text{mV}$, and θ displayed in graph. . .	117
6.27	Maximum phase shift over $V_g = [0, 500]\text{mV}$ (a) and tuning voltage required to achieve 10% to 90% of the maximum phase shift (b). . .	118
6.28	Drain current third harmonic ripple over $V_g = [0, 500]\text{mV}$	118
6.29	Frequency tripler schematic with full degeneration (a) and partial degeneration (b).	119
6.30	Frequency tripler circuit phase range (a), 10% to 90% phase shift tuning voltage (b), maximum conversion gain (c), and output voltage ripple (d).	120
6.31	Chip die photograph (a) and PCB (b).	121
6.32	Far-field measurement setup.	121
6.33	Spectrum analyzer output using VDI WR6.5 SAX-M frequency extender.	122
6.34	Two element interference pattern vs. tripler #1 gate bias voltage at 140 GHz (a), and at 135GHz, 140GHz, and 145GHz (b).	122
6.35	Single chip-tile boresight EIRP vs. frequency.	123
6.36	Single chip-tile H-plane (a) and E-plane (b) patterns.	124
6.37	Single chip-tile H-plane steering (a) and E-plane steering (b) patterns. .	125
6.38	Far-field received power spectrum and harmonic rejection ratios. Harmonic rejection is referenced to the EIRP of the carrier at 139GHz. .	126
6.39	Single chip-tile with underfill boresight EIRP vs. frequency.	126
6.40	Two-chip 2x4 array boresight EIRP vs. frequency.	127
6.41	Two-chip 2x4 array H-plane (a) and E-plane (b) boresight patterns. .	127
6.42	Two-chip 2x4 array H-plane steering patterns.	128
A.1	2D-planar NxN array	139
A.2	Uniform continuous probability density function for Δ_i	143
A.3	Coherent gain with uniform phase error distribution.	144
A.4	Coherent gain with Gaussian phase error.	145
B.1	Gaussian phase offset wrapping.	147

C.1	$\frac{P_r}{P_{max}}$ vs. delay spread for different duty-cycle values.	152
C.2	Delay range at $\frac{P_r}{P_{max}} = 95\%$ vs. duty-cycle.	152

LIST OF TABLES

<i>Number</i>	<i>Page</i>
2.1 Link budget parameters.	30
3.1 Specifications for WPT RFIC in 22-nm CMOS FDSOI.	47
3.2 VCDL phase range across PVT corners. Phase range is the phase-shift from tuning $V_{cont} = 0V$ to $V_{cont} = 0.8V$	54
3.3 22-nm PA simulations results.	59
5.1 Unit-area cost of the transmitarray (T.A.) derived from the active array (A.A.) costs provided in [23].	86
6.1 Comparison with other works at D-band.	128
A.1 G_c vs. RMS phase jitter with Gaussian random variables.	145
B.1 G_c for different phase-shifter resolutions.	148

NOMENCLATURE

- ACS.** Attitude Control System.
- AP.** Aggregation Panel.
- BTE.** Beam-to-Earth.
- CMU.** Clock Multiplier Unit.
- COTS.** Commercial Off-the-Shelf.
- DB1.** Daughter-Board 1.
- DB2.** Daughter-Board 2.
- DLL.** Delay-Locked Loop.
- FA.** Flex Array.
- JPL.** Jet Propulsion Laboratory.
- LCOE.** Levelized Cost of Electricity.
- LNA.** Low-Noise Amplifier.
- LO.** Local Oscillator.
- MPO.** Mutual Power Optimization.
- MPPT.** Maximum power point tracking.
- PA.** Power Amplifier.
- PCB.** Printed Circuit Board.
- PLL.** Phase-Locked Loop.
- PV.** Photovoltaics.
- RFIC.** Radiofrequency Integrated Circuit.
- SBSP.** Space-based Solar Power.
- VCDL.** Voltage-Controlled Delay Line.
- VCO.** Voltage-Controlled Oscillator.
- WPT.** Wireless Power Transfer.

Chapter 1

INTRODUCTION

The ability to coherently control large numbers of radiating elements is fundamentally reshaping the landscape of wireless systems. Phased arrays, once confined to defense applications, are becoming central to wireless communication, sensing, imaging, and increasingly, wireless power transfer (WPT). As semiconductor technologies scale and integration density improves, it has become possible to realize highly complex beamforming systems on millimeter-scale, low-cost CMOS radio-frequency integrated circuits (RFICs)

At the heart of this transformation lies scalable arrays composed of compact, power-efficient, and phase-controllable unit cells. When replicated across large apertures, these unit cells form systems capable of directing electromagnetic energy with high precision. The same principles govern applications across vastly different regimes from millimeter-wave sensing at D-band frequencies to microwave wireless power transmission over kilometer-scale distances. The distinction lies not in the underlying physics, but in the constraints imposed by application, environment, and scale.

This dissertation explores scalable array architectures across this spectrum, with a particular emphasis on flexible, lightweight implementations suitable for space-based solar power (SBSP) and high-frequency integrated radiator tiles for millimeter-wave systems. While these applications may appear disparate, they share common challenges: coherent phase control, power efficiency, thermal management, integration density, and manufacturability at scale.

Wireless power transfer serves as a compelling testbed for scalable arrays. By coherently summing the radiation from many elements, a phased array can concentrate energy at distant receivers without mechanical steering. Unlike conventional power distribution, WPT enables dynamic energy routing to multiple spatially separated loads, offering flexibility unattainable with wired infrastructure. In terrestrial settings, this capability can enable applications where wired connections are not possible, such as persistent intelligence, surveillance, and reconnaissance using wirelessly powered drones. In space, space-based solar power envisions kilometer-scale photovoltaic arrays converting sunlight to DC power, which is then converted

to microwave radiation and transmitted to Earth. The feasibility of such a system depends critically on the development of lightweight, efficient, and scalable array tiles that can survive and operate in harsh environments.

A central contribution of this work is the demonstration and analysis of MAPLE (Microwave Array for Power-Transfer Low-Orbit Experiment), the first dedicated free-space wireless power transfer experiment using lightweight, flexible phased arrays deployed in orbit. Its successful operation provides not only validation of flexible array architectures in space, but also insight into long-term degradation mechanisms, environmental interactions, and system-level optimization strategies required for practical SBSP.

Building upon this foundation, the dissertation presents the design of next-generation fully flexible phased array tiles driven by custom 22-nm CMOS FDSOI RFICs. Using the lessons from MAPLE, the RFIC is completely redesigned from scratch. The chip integrates all of the functionality required for a phased array transmitter, including frequency synthesis, phase-shifting, and efficient conversion of DC to RF power using a power amplifier (PA). Record DC-RF conversion efficiency is demonstrated, and a fully-flexible, 64-element phased array tested in wireless power transfer scenarios.

Beyond the microwave regime, the same scaling philosophy is extended to D-band (110–170 GHz) active radiator tiles. At these frequencies, wavelength scaling permits dense integration of radiators and electronics, enabling compact high-gain apertures suitable for sensing and high-data-rate communication. However, the millimeter-wave frequency exacerbates losses and requires complex packaging, particularly for $\lambda/2$ spaced arrays. This work introduces the flip-chip aperture-coupled architecture that tightly integrates the CMOS RFIC with the antenna radiator using standard PCB fabrication techniques. The tiles achieve dense spacing, low-cost packaging, high performance, and scalability at D-band.

This dissertation aims to demonstrate that with careful, application-tailored co-design, arrays can be made at scale with high-performance and low-cost. The remainder of this dissertation is organized as follows. Chapter 2 presents the MAPLE mission, including system design, in-orbit results, and beam-to-Earth experiments. Chapter 3 describes a next-generation flexible 8×8 phased array tile enabled by a fully integrated 22-nm CMOS RFIC. Chapter 4 introduces mutual power optimization strategies for jointly maximizing photovoltaic and RF transmission efficiency for WPT systems. Chapter 5 explores transmitarrays as a cost-reduction strategy for

large-scale SBSP systems. Chapter 6 presents scalable D-band active radiator tiles in 22-nm CMOS FDSOI. The appendices provide analytical treatments of array timing requirements, phase quantization loss, and modulation timing requirements for large WPT arrays.

Together, these contributions advance the state of scalable electromagnetic arrays from millimeter-wave sensing to microwave wireless power transfer, illustrating both the opportunities and the challenges that lie ahead.

Chapter 2

MICROWAVE ARRAY FOR POWER-TRANSFER LOW-ORBIT EXPERIMENT (MAPLE)

2.1 Introduction

Wireless power transfer (WPT) at distance could transform the nature of energy distribution and access, both on Earth and in space. This has wide implications for everything from consumer electronic devices to large scale energy generation and distribution. It can make up for gaps due to the volatile nature of clean energy sources like wind and solar, expand access to locations where physical connections are non-existent, and enable entirely new applications. This can be achieved using phased arrays consisting of large numbers of small sources working coherently to focus transmitted energy on one or more receivers. By adjusting the relative timing of the individual elements, energy can be quickly dispatched to numerous select locations simultaneously without physically moving the transmitter.

The operational flexibility of WPT introduces a myriad of possibilities. While this is true for terrestrial WPT, it is especially relevant in space where physical connections between spacecraft, space stations, and ground stations are impractical if not impossible. In particular, WPT in space enables space-based solar power (SBSP), a concept first conceived by Isaac Asimov in his short story, "Reason" [2], where solar energy is converted to microwaves in space and transmitted to Earth. Caltech's vision for SBSP (Fig. 2.1) comprises the deployment of massive kilometer-scale arrays of photovoltaic (PV) cells in space [3]. The solar power collected by the solar cells is first converted to DC power by the photovoltaics and then locally to RF power. Using phased arrays, the RF power is then wirelessly transmitted to destinations on Earth or, as envisioned by Asimov, to other planets inhabited by humans. Although light of any wavelength may be used to transfer energy wirelessly (e.g., lasers), RF/microwave wavelengths in certain frequency ranges are particularly well-suited for SBSP due to lower atmospheric attenuation [4], permitting operation in all weather conditions. Also, the larger wavelengths in the microwave band (relative to lasers) combined with the physical diffraction limit can be used to guarantee that the power intensity on Earth will not exceed a predetermined maximum, alleviating safety concerns. A sufficiently large array will




Figure 2.1: The Caltech SBSP concept [1]: A PV-WPT array converts solar energy to DC power and DC power to RF power which is then beamed to Earth and collected by a ground station.

focus close to 90% of its power in the main focal point at the receiver, with close to 100% achievable by tapering the power towards the edges of the array [4]. At the destination, the RF power is received and converted to low frequency (AC or DC) power using rectifying antenna (rectenna) arrays. Rectenna arrays for RF-DC conversion are well-studied and offer high conversion efficiency, with some systems demonstrating efficiencies as high as 70-80% [5].

Unlike solar, wind, geothermal, or hydropower, SBSP has the unique advantage of providing continuous, stable, and predictable renewable energy [6] without imposing storage demands. Moreover, SBSP can dynamically dispatch power to nearly any desired location on Earth, including locations where the capacity or the infrastructure to satisfy the demand for electricity may not be available.

Significant work has been done on SBSP over the last 50+ years, as meticulously documented in [4][7]. Early work in the 1960s and 1970s by private and government

institutions [8, 9, 10, 11, 12, 13] culminated in NASA’s system report in 1979 [14]. Progress continued with Japan’s ISY-METS mission in 1993, which demonstrated sub-orbital WPT [15]. New SBSP concepts were developed by NASA [16], ESA [17], and JAXA [18, 19] in the mid 1990s and early 2000s. With the push to transition to renewable energy sources, SBSP proposals have seen a resurgence in recent years [20, 21, 22].

Despite the potential benefits of SBSP and the large body of existing research, the enormous launch costs associated with deploying a massive array in high-altitude orbits have, to date, prohibited practical implementations of any SBSP system. To reduce the overall system mass and cost, lightweight, flexible phased arrays powered by thin-film photovoltaic cells on one or both sides can be used. Radio frequency integrated circuits (RFICs) further reduce mass and size by combining onto a single millimeter-scale chip all the necessary digital and analog circuitry for DC to RF power conversion and focusing. These RFICs combined with thin, lightweight PV and flexible, lightweight RF arrays reduce launch costs and help bring SBSP closer to reality [23].

These lightweight systems, mostly from our group at Caltech, have been developed and tested on Earth [24, 25, 26, 27, 28, 29, 30]. However, to the best of our knowledge, a dedicated free-space WPT system using this or any other technology has not been deployed and demonstrated in space.

MAPLE (Microwave Array Power Transfer LEO Experiment) is an experiment to test these arrays in space and evaluate their strengths and weaknesses. This is a critical step in evaluating the feasibility of this technology when subjected to space’s extreme environmental conditions. Ultimately, the technology must function reliably for many years for SBSP to be economically practical, and the results from this mission will inform the next development phase. This is discussed in more detail in chapter 3 of this thesis.

The remainder of this chapter is organized as follows: Section 2.2 covers system design, including a discussion of the main hardware components, software design, mechanical and thermal considerations, and space qualification testing. Section 2.3 covers results from in-orbit experiments, including demonstrations of wireless power transmission and reception at multiple targets, focusing algorithm operation and performance, and a comparison to performance on Earth. Section 2.4 covers the system design and results for the “beam-to-Earth” experiment, during which power generated by MAPLE was successively focused at and detected on Earth. Section


2.5 examines anomalous results and presents a hypothesis about the likely causes. Section 2.6 concludes the paper and briefly discusses future work.

2.2 System Design


MAPLE, shown in Fig. 2.2a, is built around a 6U CubeSat chassis and serves to test wireless power transfer in space. MAPLE is mounted to Vigoride-5, a spacecraft that provided integration with the launch provider (SpaceX), communication, attitude control, and power supply infrastructure for MAPLE. Integration of MAPLE and Vigoride-5 is shown in Fig. 2.8, taken before launch. The inside of the CubeSat, shown in Fig. 2.2b, hosts a number of devices, including:

- Two custom 16-channel silicon RFICs responsible for frequency generation, power amplification, and phase control.
- An array of 32 flexible co-cured dipole antennas.
- A flexible printed circuit board (PCB) on which the RFICs and antennas are mounted ("flex array" or FA).
- A rigid PCB with voltage regulation, on-board telemetry, and redundant microcontrollers that coordinate control of all of MAPLE's functionality ("aggregation panel" or AP).
- Two rectenna arrays with the option to load power to a resistor, bright LED, or a widebeam LED.
- Two narrow field-of-view cameras.
- Two wide field-of-view cameras.
- An opening opposite the flexible PCB sealed with a sapphire window.

The payload is designed to maximize thermal and mechanical stability, given the requirement that it survive launch and months of thermal cycling in orbit. To this end, MAPLE was designed to sink heat from power sources to the flight deck while reflecting sunlight and hazardous ionizing radiation. Moreover, redundancy was built into every level of MAPLE. Every functionality was duplicated, allowing for the possibility of mission success despite possible component failure. This redundancy is evident and partially elucidated by Fig. 2.3, which shows all of MAPLE's major components and the communication lines between them. In theory,



(a)



(b)

Figure 2.2: (a) CAD rendering of MAPLE with parts labeled, and (b) assembly. Dimensions of the cubesat are 10cm x 34.6cm x 22.6cm.

most of the mission goals could be accomplished with the survival of only five out of 11 component blocks. Details on these designs are presented in the next subsections.




Figure 2.3: MAPLE block diagram, showing various components and redundancy. Flight computer is not part of MAPLE, but has parallel communication lines and power lines to both motherboards on the AP.

Custom 65-nm CMOS RFIC

The core functionality of the WPT array is provided by the custom RFIC. The chip, designed in a 65-nm CMOS process, provides the frequency synthesis, timing, beam forming, and DC to RF conversion for 16 individual channels. Modern CMOS scaling enables this complex system to be integrated into a small, extremely lightweight $2.8 \times 2.8\text{mm}$ die, a scale unattainable using discrete components. Two of these chips are used in MAPLE to drive the 32-element phased array.

A block diagram and die photo of the chip is shown in Fig. 2.4. A 48MHz reference provided by an off-chip crystal oscillator is routed to both chips in MAPLE. On chip, it feeds into a centralized, x52 multiplication ratio phase locked loop (PLL) which synthesizes a 2.496GHz signal that is distributed to four quadrants (quads). Each quad consists of four cores. Each core contains a power amplifier chain (PA) driven by a fixed x4 multiplication ratio clock multiplier unit (CMU) that synthesizes the 9.984GHz local oscillator (LO). Phase shifting is accomplished by

injecting a digitally programmable current into the loop filter of each core-level CMU [24]. Adjacent PAs are powered in a series-stacked configuration, enabling higher supply voltages [31]. Each differential PA output is coupled to a 50Ω microstrip transmission line on the flex-array via an on-chip balun, which terminates at the pop-up dipole antenna.




Figure 2.4: 65-nm CMOS RFIC block diagram.

Popup Dipole Antennas

MAPLE uses lightweight co-cured popup antennas previously discussed in [32, 33] as the radiator of a unit element in the array. Electromagnetically, the antennas behave as resonant ground plane-backed dipoles. Mechanically, they are co-cured with a weave of fiberglass that allows them to be packed and folded. Once unfolded,




Figure 2.5: 65-nm CMOS RFIC die photo.

the inherent mechanical stress returns them to their nominal position. More details are provided in [33].

The antennas were tested for electromagnetic and mechanical performance, especially as it relates to survival in harsh space environments. Results presented in [33] demonstrate the antennas possess all the typical advantages of a ground-backed dipole antenna (efficiency, bandwidth, simplicity, directivity, etc.) while also demonstrating thermal and mechanical durability. The ground plane-backed dipole antennas with a vertical feed that doubles as a single-ended to differential balun were mounted to a single-ended transmission line in MAPLE's flexible PCB using 3M DP2216 epoxy.

Flex Array

The RFICs and popup antennas are mounted on a lightweight, flexible, four-layer polyimide PCB, which together are referred to as the “flex array” or FA. The FA measures roughly $160 \times 90\text{mm}$ and, while physically a single board, effectively




Figure 2.6: Flexible array with popup dipoles (engineering model depicted), fixed to an aluminum frame for heatsinking and mechanical integration.

consists of two adjacent daughter boards, or DBs, each containing 16 antennas arranged in a 4×4 grid. Each daughter board is driven by a single RFIC, with digital control originating from an independent microcontroller soldered near the RFIC. By radiating coherently, the two daughter boards form a phased array of 32 elements in an 8×4 configuration (Fig. 2.6). Considered together, the two sub-arrays increase the aperture size and, therefore, the focusing performance of the array; considered separately, they introduce an additional layer of redundancy.

The four metal layers provide the following functionality:

1. Layer 1 provides component pads for the RFICs, microcontrollers, and associated passive circuitry.
2. Layer 2 contains power distribution lines for 3.3V, 1.8V and 1.2V supplies that power the digital, power amplifier, and analog circuitry, respectively.
3. Layer 3 consists of a low-mass ground plane for the microwave traces and

antennas. The mass reduction relative to a standard ground plane is accomplished using a low-fill factor, rectangular mesh.

4. Layer 4 routes 50Ω microstrip transmission lines to the antenna solder pads.

Aggregation Panel

The Aggregation Panel PCB ("Aggregation Panel" or AP) provides the electronic interface between devices in MAPLE and the flight computer (FC) on Vigoride-5. It is a six-layer, conformally-coated, rigid polyimide board.

At the core of the AP are two redundant mother boards (MBs) (again, physically on the same board), wired in parallel, and each capable of communicating and controlling all subsystems, including the DBs and RFICs. Each MB consists of an Atmel SAMD21 microcontroller, and other accompanying devices. The AP also integrates the necessary switches and sensors for system telemetry (current, voltage, temperature), connected to and controlled by the MBs.

The AP interfaces with the two rectennas via two separate voltage sensing lines for the two MBs and control lines for switching between the rectenna LEDs and loads. The AP also hosts two narrow-FOV cameras and communicates with the wide-FOV cameras located off-board. Redundant crystal oscillators mounted on the AP provide the RFICs with a low-frequency reference; oscillators are powered and switched by either MB.

Finally, the AP hosts the voltage regulation devices which source power from the FC and output supplies to all MAPLE devices. All supply lines from the FC are 2x redundant; diode OR-ing of the FC supplies allow for operation with either supply in the event one fails.

Rectennas

MAPLE contains two rectennas that receive and convert the FA's wirelessly transmitted RF power to DC. Two boards provide redundancy, as well as spatial diversity to permit demonstration of the RFIC's ability to quickly steer power to different locations. One rectenna board is positioned broadside the array, while the other is offset and mounted on a sidewall. Rectennas are multi-substrate four-layer boards using Rogers 4003C for the antennas, Rogers 4350 for the rectifier circuits, and an FR4 pre-impregnated layer in the middle.

Each rectenna comprises a 5×5 grid of rectenna elements (combined in DC), two




Figure 2.7: 5×5 rectenna array of dual-polarized patch antenna receivers.

redundant maximum power point tracking boost circuits, two 22Ω DC loads¹, and two white LEDs each controlled by one of the two MBs (Fig. 2.7.) Onboard MOSFETs allow a control signal from the AP to switch between the 22Ω load and the boost circuit that powers the LED. The load serves both as a stable, power-optimal load for optimization as well as a safety precaution to prevent damage from open circuit rectenna operation. Each rectenna element consists of a dual polarized, probe-fed patch antenna that directly feeds two identical rectifier circuits, designed based on [34]. For a nominal rectified power of 210mW, the rectenna array has a conversion efficiency of 40% with respect to the power incident on its aperture. As shown in [35], rectenna efficiencies at X-band can reach 70% under optimum load and power settings.

Cameras

Four redundant cameras (two wide FOV and two narrow FOV) are present on MAPLE. Wide FOV cameras have $\approx 170^\circ$ FOV fisheye lenses attached to Raspberry Pi High Quality Cameras and are mounted on MAPLE's sidewall. Narrow FOV cameras (Raspberry Pi Camera Module 2) are mounted directly to the AP. Each camera is connected to a Raspberry Pi Zero W which handles camera operations and saves pictures to SD cards. UART serial lines provide the control and data

¹Four individual, 22Ω , 0.25W resistors are wired in a series-parallel configuration to increase the load power rating to 1W and mitigate possible thermally-induced load variations.

transfer mechanism between the Raspberry Pis and AP.

Software

MAPLE utilizes an implementation of the algorithm presented in [34]. This algorithm is designed to bypass the trade-off between accuracy and runtime that typically plagues optimizations of high-dimensional systems². It does this by sweeping over successively small ranges of phases with a large number of elements. Groups of elements swept together (called “masks”) help ensure changes exceed the noise floor of the system while also accelerating the optimization process by parallelizing tests. Orthogonal mask sets are chosen such that different combinations of elements are varied in different sweeps, thus increasing the search space explored in few iterations.

The focusing algorithm is predicated on a closed-loop feedback configuration of a “generating unit” (GU) and “receiving unit” (RU), as depicted in Fig. 5 of [34]. In MAPLE, the FA (flexible PCB, antennas, and RFICs) constitutes the “GU array.” The array is controlled by a microcontroller on the MB which coordinates execution of the focusing algorithm by adjusting phases sent to the RFICs; the FA together with the MB constitute the GU. The MB also reads back voltage levels from the receiver array (RU) which is implemented as the rectenna discussed above.

MAPLE’s optimization algorithm begins with an initialization of the RFIC settings and phases. Then, the MB begins adjusting select phases, detecting the resultant change at the receiver, and updating phases based on the change in received power. The algorithm continues for different combinations of elements until a specified number of iterations have been completed.

A functional system running the focusing algorithm would produce the type of curve, relating rectified power at the receiver to elapsed time, shown in Fig. 2.12a: power increases rapidly at the beginning and stabilizes at, presumably, a maximum.

Additionally, the ability to upload an array of digital codes from MB memory to the RFICs, herein called “recall,” was implemented. Recall allows MAPLE to measure the power rectified at the receiver using phase settings that were either saved in (nonvolatile) program memory (for example, to compare performance on Earth and in orbit) or in volatile memory (for example, to start an optimization with the phase codes output as a local maximum by a previous optimization during

²For an array of 32-elements, each digitally controlled with a 9-bit phase code, there are $(2^9)^{32} = 512^{32} \approx 10^{87}$ unique phase configurations.

the same experiment.) The former enables MAPLE to make direct comparisons of system performance under identical phased array programming while the latter enables “cascaded” optimizations, where the final state of the optimization is used as the starting state of the subsequent one




Figure 2.8: MAPLE integration on Vigoride-5 spacecraft from Momentus Inc. [36].

Mechanical Enclosure

MAPLE’s enclosure is a standard 6U aluminum CubeSat frame. The 6U frame contains all of MAPLE’s electronics, as shown in Fig. 2.2. Cabling for power and external communications are all provided by a single Micro-D connector mounted on an aluminum back-panel fixed to the CubeSat frame. The peripheral cameras and rectennas are fixed to aluminum rings which then mount to the CubeSat frame. A sapphire viewing port placed downrange of the FA is optically transparent for pictures and provides a seal to protect against atomic oxygen. The sapphire was chosen to be 4.4mm ($\lambda/2$ at 10GHz) thick to maximize its RF transmittance when focusing power outside of MAPLE.

Vigoride-5 Integration

MAPLE was integrated on the Vigoride-5 spacecraft built by Momentus Inc. as shown in Fig. 2.8. Vigoride-5 provides the power, communications, thermal sink, and attitude control for MAPLE as well as ALBA and DOLCE, two additional instruments in the SSPD-1 mission [37].

Thermal Design

All components in MAPLE are thermally coupled to the deck of the Vigoride-5 spacecraft. The deck serves as an isothermal boundary in the accompanying analysis. Most components are thermally coupled to the deck via connections from the AP's ground plane to the CubeSat frame.




Figure 2.9: RFIC thermal simulation. Heat flows from the chip through the aluminum frame to the temperature boundary at the bottom face. Heatsink stackup is shown in the lower left (dashed box not displayed in other images.)

As the most concentrated sources of heat, the RFICs required special attention to ensure their temperatures did not rise above their effective operating point of 80°C. The design and simulation of the RFIC's heat sink solution is shown in Fig. 2.9. An aluminum frame that supports both the FA and AP includes two pedestals placed directly behind RFICs. The RFICs are bonded to these pedestals with Stycast 2850 thermal epoxy. Bond line thickness is controlled by glass beads mixed into the thermal epoxy. The heat from the chips couples to Vigoride-5 predominantly through this aluminum frame. Thermal simulations with a 50°C temperature boundary (worst-case) at the deck produces a max temperature of approximately 74°C at the RFICs.

In addition to the conductive heat sinking, MAPLE is enclosed by a 12-layer blanket

of multi layer insulation (MLI) to prevent excessive radiative cooling and shield it from direct sun exposure. The outer layers of the MLI stackup are aluminized polyimide, while the inner layers are low emissivity aluminum blankets with perforations for air egress during the launch.

Radiation Testing




Figure 2.10: Texas A&M Radiation Effects Facility proton beam.

Radiation exposure is one of the primary, unique environmental challenges in space missions. Before launch, it was important to verify that all chosen components would survive radiation exposure for prolonged periods of time.

Commercial off-the-shelf (COTS) electronics were preferred due their low-cost and ease of acquisition. Compendiums of COTS parts subject to radiation testing at NASA's Jet Propulsion Laboratory (JPL) guided the selection process to ensure components had moderate resilience to radiation events [38][39]. For parts, such as the RFICs without space or radiation testing legacy, extensive proton radiation testing was performed at the Texas A&M Radiation Effects Facility, whose setup is shown in Fig. 2.10. This testing ensured survivability under the mission's expected total ionizing dose. Parts were exposed to a 45 MeV proton beam for either one or three minutes at a time with a fluence of 10^{10} particles/cm². No physical or electronic damage was observed.

Vibration Testing

MAPLE was subjected to vibration testing as required by the launch provider, SpaceX, in accordance with their known launch acceleration spectrum. Vibration testing was performed a number of times over a two year period, both to individual components [33] and to the entire assembled payload. The following tests were performed:

- Shock: A time-domain test meant to simulate sudden accelerations associated with dropping or banging a component.
- Sine Burst: A time-domain test at a single (low) frequency meant to test high displacement vibration.
- Sine Vibe: A frequency-domain sweep in accordance with the rocket's spectrum.
- Random Vibe: A time-domain test which generates a random vibration waveform in accordance with the rocket's spectrum.

Tests were performed in all three Cartesian axes; MAPLE passed all testing without damage and with minimal variations to its mechanical spectrum. Tests did not alter electronic performance.

Temperature and Vacuum Testing

Thermal vacuum chamber (TVAC) testing verified MAPLE could survive the harsh thermal environment of space. Thermal testing was performed on individual components [33, 33] and the entire assembled payload, both in vacuum and at ambient pressure.

For TVAC testing, MAPLE was placed in a vacuum chamber pumped down to an internal pressure of less than 10^{-6} torr. Temperature control in the chamber was programmed to simulate the moderate and extreme temperature limits expected in orbit:

- Predicted (power-off) temperature limits [5, 15]°C
- Protoflight temperature limits [−10, 30]°C
- Extreme temperature limits [−30, 60]°C

MAPLE spent several hours off, "soaking" at the extreme minimum and extreme maximum limits. Then, several square wave cycles with a 95 minute period held

MAPLE at the protoflight limits. Normal operation command sequences were tested during both the hot and cold phases. Several final soaks at the extreme limits followed, ending with a 100 minute soak at room temperature.

The payload survived testing, as verified by visual inspection, electrical testing, and TVAC testing results.

2.3 In-Orbit: Experimental Results

MAPLE was launched into orbit on a SpaceX Falcon 9 rocket on January 3rd, 2023 UTC. MAPLE's orbital height was approximately 527 km with an inclination of 97.5° [40]. After initial qualification testing of Vigoride-5 completed, MAPLE's first experiment was conducted on March 3rd, 2023 at 11:51 PM UTC. MAPLE was deployed for a total of ten months; testing took place during the latter eight months.

In-Orbit Wireless Power Transfer Experimental Results

Effective execution of the focusing algorithm resulting in relatively high power received at the rectifier constitutes the single most important result from MAPLE. To successfully generate the desired optimization curve, all the software and hardware needs to work. Partial or complete failure of a component would affect several parameters of the measured optimization curves: the shape and/or the maximum asymptotic value.

The rectified power curves corresponding to optimizations performed in LEO are pictured in Fig. 2.11³. The flex array created coherent focal points on each of the rectennas. The peak rectified DC powers received by the the broadside and sidewall rectennas were 251mW and 231mW, respectively. Each of the 16 stacks, composed of two PA cores (Fig. 2.4), consumed about 230mA from a 1.73V supply. Once phases were optimized to maximize power rectified by the receiver, power was switched from a resistive load to a boost circuit which powers an LED. The intensity of the LED, shown in Fig. 2.13, demonstrates the magnitude of power transferred.

Power degraded towards the end of the focusing in Fig. 2.11a and 2.11b. This result is a deviation from ground testing. Separate in-orbit measurements focusing three times back-to-back using only 16 elements controlled by daughter-board 1 (DB1) (Fig. 2.12a) and daughter-board (DB2) (Fig. 2.12b) were conducted. The results indicate degradation in the recovered power from DB2 that is exacerbated in subsequent optimizations. DB1, in comparison, output power consistent with

³The curves correspond to the optimizations that resulted in the highest rectified power to the respective rectennas during the duration of the mission.




Figure 2.11: (a) Rectenna #2 (broadside) optimization on 2023/03/03 - 23:51 UTC, and (b) Rectenna #1 (sidewall) optimization on 2023/03/14 - 03:19 UTC.

ground tests and remained stable over time. Additionally, the 32 element power levels are restored once the system is restarted from an off state. Degradation in DB2 is the most likely cause of the drop in power seen in the full 32-element focusing experiments (Fig. 2.11). More discussion of this phenomenon is presented in Section 2.5.

Individual Core Performance

MAPLE is equipped with a function which generates two-core interference patterns at the receiver for the characterization of individual element (core) performance. Testing of each of individual cores proceeds as follow:

1. Addressing a single stack, one core is turned on with variable phase while the




Figure 2.12: Rectenna #2 focusing with only 16 elements controlled by (a) DB1 on 2023/03/21 - 23:16 UTC, and (b) DB2 on 2023/03/22 - 00:47 UTC.

other is turned on with constant phase.

2. The phase of the first core is swept digitally through some subset of the entire programmable range.
3. At each phase iteration, voltage at the broadside receiver is collected and recorded back to the flight computer (FC) operating MAPLE. Then, the phase of the second core is swept with an identical subset of codes while the first core's phase is held constant.
4. This is repeated for each of the 16 stacks across the two RFICs on MAPLE.




Figure 2.13: Interior of MAPLE during a power beaming experiment on March 15th, 2023 – 4:40 AM UTC. Lit LED (center) indicates power transmission to the (sidewall) rectenna.

An example of data generated from this procedure done on Earth is presented in Fig. 2.14, in black. The same procedure done in space is overlaid in red.

The rectified power vs. swept phase should produce a pseudo-sinusoidal interference pattern, which is seen in Fig. 2.14. Six on-Earth tests performed after integrating MAPLE onto the flight deck of Vigoride-5 but before launch, herein called “integration tests,” are the best indications of on-Earth performance for the purpose of comparison to in-orbit data. Integration tests were performed as close to launch as possible to provide the best possible baseline for in-orbit performance. From the raw data in Fig. 2.14, the phase-shifter curves can be extracted by fitting it to the model derived in [41]. Averaged over tests, cores, and phase codes, the absolute deviation from the mean CMU phase function for integration tests is small at 3.90° , demonstrating the CMU’s ability to consistently produce the same output phase for a given input code.

Three in-orbit tests performed over two days, two months after launch, can be compared to integration tests and among themselves to understand how the CMUs perform differently in orbit and over time. The ensemble standard deviation of CMU phase functions among in-orbit tests, averaged over cores and phase codes, was 5.87° , greater than that for the integration tests. This is indicative of decreased CMU phase-shifting performance stability in space, possibly caused, as discussed




Figure 2.14: Raw rectified power from both on-Earth (black) and in-orbit (red) interference tests. On-Earth curves (black) represent the mean of six tests performed before launch; in-orbit curves (red) represent the first test performed in orbit. The label on the upper left indicates the “coordinates” of the core: (RFIC, quad, stack, core). The y-axes have limits of [0 20] mW.

below, by increased RFIC temperatures owing to the vacuum in space. Compared to the mean CMU curve from integration testing, in-orbit CMU curves have mean absolute deviation of 20.43° but median absolute deviation of 9.72° . Thus, CMU phase-shifting stability is slightly worse in space, and the exact CMU curve changes as a result of the change in environment.

Long-Term Performance Degradation

MAPLE ran a benchmark experiment sequence in the first and last days of its ten month deployment period. These two runs were performed under nearly identical conditions with a nearly identical sequence in an effort to directly quantify long-term system degradation. The sequence consists of a number of single optimizations, successive optimizations (one optimization starting where the previous stabilized), recalls, chip-to-chip interference cycles, and a characterization of individual core performance via amplitudes of two-core interference patterns. Fig. 2.15 compares the rectified power during these tests. Considering a mean taken over these 18 benchmark values, power declined by 14.7%.

However, when individual cores were analyzed, we observe degradation in perfor-




Figure 2.15: Tracking long-term performance degradation by comparing two benchmark experiment runs eight months apart. Left: Power rectified during different phases of the benchmark experiment. “O” = optimization, “R” = recall, “I” = interference. “O1 (peak)” denotes the peak power during the first optimization. Right: A comparison of individual core amplitudes (A_i) at the beginning (x-axis) and end (y-axis) of the ten month deployment. Each data point corresponds to a core, labeled using the core number ($\in [1, 32]$) for the first core in the stack. Degraded cores circled in red.

mance in only three stacks (six cores), circled on the left of Fig. 2.15. The other 26 cores demonstrated a very high correlation between amplitude in the early and late tests, appearing virtually unaffected by time spent in space. CMU phase functions (not shown) were virtually unchanged over the deployment period.

Summing the individual amplitudes and assuming ideal coherence, degradation in these three stacks contributes to a total degradation in peak power of 18.25%. The similarity between measured decline in peak power and potential peak power based on individual core amplitudes (14.7% vs. 18.25%, respectively) suggests that power decrease is not owing to a broad, chip-level degradation in performance, but rather the effect of a number of cores that have individually degraded. Moreover, the system’s ability to reach 191mW (I (max), 2023/11/10), close to the original peak of 247mW (I (max), 2023/03/15) prior to long-term degradation, suggests that fundamental system performance is largely maintained.

2.4 Beam-to-Earth Experiments

MAPLE is primarily a self-contained instrument. That is, the vast majority of experiments are conducted internally in the 6U cubesat frame. Experiments can be performed under any conditions in orbit, with minor implications on the temperature depending on whether the spacecraft is in Earth’s shadow. Results are recorded on the flight computer and downlinked when in line-of-sight with ground stations located at the northern or southern poles.

However, there is one experiment that breaks this paradigm. MAPLE, although not explicitly designed for it, is capable of radiating outside of its box. By doing this at a time where it is within line-of-sight of a ground station at Caltech, and pointing its beam appropriately, it can be detected on the ground. It should be clearly emphasized that this is detection rather than power transfer, detection being defined as the recovered power being less than the power needed to receive it. MAPLE's aperture is simply too small to focus its energy to a practical footprint on Earth. However, detection on Earth validates in a complex and demanding situation the full functionality of the MAPLE by stretching it to its limit. These experiments are referred to as the "beam-to-earth" (BTE) experiments.

The critical components to the experiment, other than a functioning flex array, are the ability to direct MAPLE's beam outside of the cubesat, point the beam at a location on Earth using Vigoride-5's attitude control system (ACS), and finally a ground station to detect the signal with sufficient SNR. All of these are discussed in the following sections.

Three BTE experiments were done over a three month interval between May 2023 to July 2023. MAPLE's beam was successfully detected in all experiments, with the final providing the highest quality data. The results of each experiment will be discussed to provide insight on the challenges encountered during these experiments.

Antenna Pattern

MAPLE is enclosed in MLI as discussed in section 2.2. The MLI is mostly aluminum, and thus electromagnetically behaves as a metal cavity. The sapphire window, designed to be optically transparent while still providing protection against atomic oxygen in LEO, is also where RF power can leak out from MAPLE. Sapphire has a dielectric constant of 11.6 in the axis normal to the surface (in this case, C-axis) [42]. To maximize the transmission through the window, the thickness was chosen to be $\lambda/2$ at 10GHz (4.4 mm), where λ is the wavelength in sapphire. The vacuum-sapphire-vacuum interface has zero reflection for a normally incident wave when the sapphire thickness is $\lambda/2$. This can be understood with a transmission-line analogy, where the impedance looking into a $\lambda/2$ transmission line is always equal to the load impedance regardless of the line's characteristic impedance.

A closed-loop optimization with the ground station is not possible in space, so we rely on recall to produce a beam optimized for radiation outside of MAPLE. To do this, MAPLE was placed on a far-field range (on Earth) and power was optimized

at the far-field probe. These were repeated under varying array conditions and the optimal phase states saved in the microcontroller flash memory, allowing them to be recalled on-demand. The results from those experiments are shown in Fig. 2.16, where a tile is defined as a single RFIC and the accompanying 16-element array. The EIRP is approximately 30dBm with both tiles, and drops to about 27dBm when only one tile is on. Interestingly, this is only a 3dB drop, less than the expected 6dB for an array in free-space. This indicates the antenna gain is limited by the window aperture size.

The beam is centered at $+13^\circ$ off from the window's axis with both tiles on, $+20.5^\circ$ with only tile 1 on and $+6.5^\circ$ with only tile 2 on. Due to the limitations of tile 2 as discussed previously, only tile 1 was used for the beam-to-earth experiment. Thus, Vigoride-5's ACS system was configured to provide a 20.5° correction to point MAPLE's beam at the ground station.




Figure 2.16: MAPLE beam-to-earth far-field pattern

Link Budget and Ground Station Design

At a frequency of 10GHz and an orbital height of 527km, the free-space path loss is 167dB when Vigoride-5 is directly overhead of the ground station. The EIRP of the transmitter is 27dBm from Fig. 2.16. So, the combination of the the path loss and transmitter EIRP results in a signal power of -140 dBm at the receiver. To boost this value, we use a high-gain antenna at the ground station. In particular, a 67cm diameter dish antenna with a gain of a 33.5 dBi and a beamwidth of 2.8° was

selected. The narrow beamwidth of the receive antenna implies the dish must be steered such that it continuously tracks MAPLE's location during its orbital pass. This is done by mounting the dish antenna on a standard telescope equatorial mount, shown in Fig. 2.17.



Figure 2.17: Dish antenna equatorial mount.

The receiver uses a heterodyne architecture to downconvert the 10GHz signal from the antenna to a 5 MHz IF. The block diagram of the electronics is shown in Fig. 2.18. All of the components in Fig 2.18 are COTS parts. The low-noise amplifier (LNA) is placed directly at the dish antenna feed to minimize noise figure, and the rest of block diagram is housed in a custom-built enclosure shown in Fig. 2.19. The total gain of the receiver electronics is 68.5 dB with a noise figure of approximately 5 dB. The noise figure takes a 3 dB hit at the mixer because there is no image rejection, causing the noise at the image frequency to fold into the same IF at 5 MHz. There are four identical channels in Fig. 2.19, two of which are used for each polarization of the dish and the other two for each polarization of a redundant antenna.

The IF output is digitized by a 20Msps Keysight M3100A PXIe. To maximize SNR, the received signal is match filtered using a Fast Fourier Transform (FFT). However, the coherent processing interval cannot be too large as the Doppler shift from the spacecraft trajectory will cause frequency smearing between adjacent FFT bins. A processing interval of 26 ms, equivalent to a frequency bin bandwidth of approximately 40 Hz, maximizes the SNR while avoiding frequency smearing.

Using the transmitter, channel, and ground station design parameters discussed above and summarized in 2.1, the SNR of the link can be calculated using the equation below.




Figure 2.18: Ground station receiver block diagram. There is an additional 3dB of loss from the cabling between the LNA output and the mixer input.




Figure 2.19: Receiver electronics enclosure.

$$SNR = EIRP - PL + G_{rx} - 10\log_{10}(kT_eB), \quad (2.1)$$

PL is the path loss, G_{rx} the dish antenna gain, k the Boltzmann constant, T_e the system noise temperature, and B the bandwidth. The system noise temperature is simply $T_s = T_a + T_e$, where T_a is the antenna noise temperature and T_e the effective noise temperature of the electronic blocks. The expected values for these parameters are summarized in the table below. Plugging the values from table 2.1 into equation 2.1, the resulting SNR is 52 dB. It should be emphasized that this is the upper bound on the SNR, which only occurs when the alignment between TX/RX is perfect and the spacecraft directly overhead of the ground station. Still, there is significant margin in the system to account for non-ideal degradations.

Given the narrow beamwidth of the dish antenna, the accuracy of the pointing is

$EIRP$	27 dBm
PL	167 dB
G_{rx}	33.5 dBi
k	1.38×10^{-23} J/K
T_a	25 K
T_e	627 K (NF = 5dB)
B	40 Hz

Table 2.1: Link budget parameters.

critical to closing the link during the BTE pass. Before each pass, the dish and corresponding equatorial mount are calibrated with the following procedure.

1. The dish antenna is placed on the equatorial mount on the roof of Moore Laboratory, as shown in Fig. 2.20.
2. A broad beamwidth transmitter is placed on the top floor of Caltech Hall, approximately 0.35 km from the dish.
3. The dish is connected to a spectrum analyzer and then swept in azimuth and elevation. The power at each position is recorded with the spectrum analyzer. Then, elevation and azimuth position corresponding to peak received power is recorded. This position corresponds to the boresight of the dish antenna aligning with the transmitter.
4. While the dish antenna is aligned with the transmitter, an optical red-dot sight fixed to it (Fig. 2.21) is adjusted such that the red-dot aligns with the known physical location of the transmitter at Caltech Hall. This ensures the optical sight is aligned with the antenna boresight.
5. Before the BTE pass, during the evening, the optical sight is aligned with a known celestial object (moon, planet, star etc.) in the sky at a known time. Combined with the physical location of the equatorial mount provided by GPS, this provides enough information to determine where the antenna is pointing in the sky. This is a standard procedure for telescope mounts. Antenna calibration is ensured by the alignment between its boresight and the optical sight.

With the pointing calibrated, this data is fed into the equatorial-mount to provide the appropriate offsets. Then, tracking is provided by the SkyTrack program [43], which




Figure 2.20: Receiver calibration setup. The receiver is located on the roof of Caltech's Moore Laboratory. The target is located on the top floor of Caltech Hall (formerly Millikan Library).



Figure 2.21: Dish antenna with optical sight fixed at the top.

uses publicly available two-line element sets from [40] to point the dish antenna at the projected trajectory of Vigoride-5.

Spacecraft Attitude

As mentioned previously, MAPLE's beam must be pointed at the ground station dynamically as it passes overhead. This is accomplished using the ACS of Vigoride-5, which rotates the entire Vigoride-5 spacecraft. This was provided by the Momentus Inc. operations team, and the details of that system are beyond the purview of this work. Importantly, the attitude of the spacecraft can be verified after the pass by

downlinking the quaternions calculated using Vigoride-5's on-board star tracker, allowing us to verify the accuracy of the pointing.

Experiment #1 - 5/23/2023

A MATLAB script using the Satellite Communications Toolbox in conjunction with publicly available TLE's was used to project access times with Vigoride-5. Close fly-by's with near overhead passes were prioritized, as those provide the best SNR with longest access window.

Using this script, the first BTE attempt was made on 5/23/2023 UTC at approximately 5:00 A.M. Notably, Vigoride-5's ACS was not operational during this attempt. Vigoride-5 was tumbling in its orbit, and the exact attitude was unknown before the attempt. Thus, the attempt was purely opportunistic, in the hope that Vigoride-5's tumbling resulted in an attitude that was sufficient for detection. Thankfully (and extremely lucky), this worked out in our favor. A plot of the results from this experiment are shown in Fig. 2.22.

The orange dot in Fig. 2.22 is the ground station location in Pasadena, and the orange ellipse is MAPLE's 3dB beam projection on the ground. Clearly, there is significant pointing inaccuracy in the spacecraft, which is expected given the lack of any ACS. Still, it is close enough to provide detection above the noise floor, as shown in the IF signal power plot in Fig. 2.22. Even more compelling is the Doppler shift, which does not depend on the pointing accuracy, very closely matching a simulation from the orbital trajectory. Given that this occurs at precisely the time predicted from the TLE projection, this confirms the signal detected is from MAPLE.

Experiment #2 - 6/30/2023

The next BTE attempt was made on 6/30/2023 at approximately 6:00 P.M. UTC. For this attempt, Vigoride-5's ACS system was functional and maintained a nadir-pointing attitude. The nadir is the point on the ground directly below the spacecraft. The results of that experiment are shown in fig. 2.23. The attitude was well-controlled by Vigoride-5, tracking the nadir as it passes over Pasadena.

More data was acquired in this experiment vs. that on 5/23/2023. However, there is a gap that starts at approximately 17:55:30 UTC where the Doppler data drops out. This is due to an unfortunate glitch in the ground station where a power-cord was unplugged after tripping on it. Tracking was lost during this time. This was fixed and tracking acquired by approximately 17:57:00 UTC and was maintained for the




Figure 2.22: BTE #1 (5/23/2023) results. IF signal power in the upper right, Doppler shift on the lower right. Snapshots of MAPLE's beam projected on the ground (orange ellipse) shown on the left with corresponding time.

rest of the pass. There is again excellent agreement in the available Doppler data, which confirms that the signal is from MAPLE

Experiment #3 - 7/28/2023

The final BTE attempt was made on 7/28/2023 at approximately 5:00 A.M. UTC. A picture of the ground station setup on the roof of Moore Laboratory is shown in Fig. 2.24. The ground station was able to track Vigoride-5 for the entire duration of the pass. The power amplifier current, measured using flight computer telemetry, indicates the interval of time during which MAPLE was actively transmitting power: starting at 4:55 UTC and ending after MAPLE crosses the horizon. MAPLE's turn-on time corresponded to its crossing the horizon and matches the time at which IF




Figure 2.23: BTE #2 (6/30/2023) results. IF signal power in the upper right, Doppler shift on the lower right. Snapshots of MAPLE's beam projected on the ground (orange ellipse) shown on the left with corresponding time.

power begins to exceed the system noise floor. MAPLE was closest to the ground station at 4:58:15 UTC; this is approximately the time of peak IF signal power and zero Doppler shift, as predicted.

Snapshots of MAPLE's location relative to Caltech and the -3dB isopower contour of its main beam are shown at the bottom of Fig. 2.25 for three different times. Unfortunately, the ground station was never within MAPLE's -3dB contour, due to Vigoride-5's inability to maintain alignment, thus reducing peak IF power relative to the power budget prediction. Curves showing the range of predicted IF signal power in Fig. 2.25 account for pointing mismatch using a beam pattern measured on the ground with the engineering model. The exact pattern in orbit was likely slightly different due to the difference in environment and phase settings. This likely contributed to the discrepancy between predicted and measured IF signal power.

The signal was lost as MAPLE descended below the horizon at approximately 5:02

UTC. The measured Doppler shift over time closely matches a prediction made using the known payload orbital dynamics. Close alignment between IF signal power and power prediction, alignment between Doppler shift measurement and frequency prediction, and spacecraft telemetry confirm to a high degree of certainty that the signal detected was MAPLE.



Figure 2.24: BTE #3 setup on Moore Laboratory rooftop.

2.5 Anomalous Results

“Amomalous” used herein describes any observable and measurable phenomenon in the data collected from the MAPLE Flight Model (FM) during in-orbit testing which was not present in the data collected from the MAPLE Engineering Model (EM) on Earth during normal operation. In most cases, these phenomenon were adverse. A number of these phenomenon are relevant to the discussion for the purpose of completeness, understanding the system better, and identifying potential areas for improvement. These phenomenon, called “symptoms” below, are presented, followed by a discussion on the testing performed to uncover the possible cause(s).

Rectified Power

Primary among the anomalous symptoms observed during in-orbit experiments are properties of the rectified power waveforms produced at the rectenna. During optimizations, rectified power should increase more or less monotonically, save for expected and occasional dips in power as the algorithm stumbles on local maxima in




Figure 2.25: BTE #3 (7/28/2023) results. PA current draw (top), IF signal power (middle), and Doppler shift plotted vs. time. Snapshots of MAPLE's beam projected on the ground (orange ellipse) shown at the bottom with corresponding time.

the search space. After 10-30 seconds, the algorithm should stabilize at a maximum power (often close to the global maximum), which it is able to hold for the remainder of the optimization period. Inability to remain at this maximum, demonstrated by a slow drift downward ("drooping") or sporadic changes to the output power, indicates the system is varying in time and undermining the algorithm's computations.

Some of the results measured in-orbit are anomalous with respect to both the absolute power level achieved by normal optimization runs and the shape of the rectified power curve. With respect to the absolute power level: after the first optimization, the system is unable to reproduce peak power at the receiver. Optimizations stabilize, instead, at lower power values.




Figure 2.26: (a) On-Earth optimization sequence performed directly prior to launch, and (b) In-orbit sequence performed on March 3rd, 2023.

With respect to the shape of the power curves: optimizations often failed to produce the typical increasing asymptotic curve that is indicative of a healthy, working system. Instead, curves showed sporadic behavior including dramatic but continuous decreases in rectified power and disjointed decreases in rectified power. This is evidenced by the data presented from the first in-orbit test performed on March 3rd, 2023 in Fig. 2.26b.

Presented in Fig. 2.26a is data collected during integration testing wherein the expected power curves are produced by optimization. During a typical optimization, power levels repeatedly and consistently asymptotically increase and then stabilize at some maximum. When an optimization begins at the stable point of the previous optimization, as is done in the second half of the data in Fig. 2.26a, there should be marginal increases in rectified power, but generally a flat curve.

However, in Fig. 2.26b, a different pattern emerges. The first optimization, to




Figure 2.27: Supply current drop anomaly on DB1. Left: Supply current for RFIC1 (lines A and B) measured over time. Right: Supply current plotted against board temperature. t_{knee} , T_{knee} , and I_{equl} labeled.

rectenna 2, initially rises to 250mW but, instead of settling, begins to drop continuously, finishing at 175mW. The second optimization, also to rectenna 2, does not approach the previous peak of 250mW, but momentarily stabilizes at 140mW before dropping to 100mW. The third optimization, to rectenna 1, has the appropriate shape, but peaks at 130mW, lower than the integration peak of 175mW. The fourth optimization peaks still lower, but demonstrates the expected flat curves. The last two optimizations are similarly low power but flat.

Optimizations performed while PAs drew less power demonstrated lower performance degradation. Power reduction was accomplished by either lowering output PA gain directly or by only using a subset of the array elements instead of all 32. In the latter case, optimizations had nominal shapes and peak power levels given the number of active elements.

PA Supply Current

PA supply currents measured during in-orbit tests also appeared anomalous in some cases. PA supply lines are routed from the flight computer (FC) to the AP where they are voltage-regulated and then passed to the FA for routing to the two RFICs. Four PA supply lines (1A, 1B, 2A, 2B) route to the two RFICs (1/2); for RFIC1, quads 1 and 2 share supply line 1A and quads 3 and 4 share supply line 1B. On the AP, these lines are voltage- and current-monitored during regular “health checks.” These health checks provide the telemetry data referenced herein which are indicative of PA operation during testing.

Health checks measured supply lines 1A and 2A slowly rising ($\approx 5\%$) and then rapidly dropping ($\approx 24\%$) after PAs were on for a sufficiently long time, as shown in Fig. 2.27. This point in time, called t_{knee} , varied between RFICs (t_{knee} , RFIC1 $\approx$

100s, t_{knee} , RFIC2 ≈ 20 s) and between tests. For reasons discussed below, it's likely the drop in current is triggered by increased temperature. A similar phenomenon was seen during radiation testing, where the chip temperature rose rapidly during operation due to lack of a heatsink. The temperature of the AP at t_{knee} is called T_{knee} . Both quantities are labelled in Fig. 2.27.

PA supply lines 1B and 2B never dropped - they were observed rising about 10% and then stabilizing, even for very long tests.

Chip-to-Chip Interference Pattern

MAPLE is equipped with algorithms that create interference patterns between the two RFICs as a way to characterize chip-to-chip coherence. To accomplish this, the array of phases for one RFIC is held constant while the array of phases for the other RFIC is progressively and discretely swept through a list of phases from 0° to 360° . To add, for example, 10° to one chip, each core's previously measured CMU phase function is inverted to solve for the necessary unique phase code that will accomplish this phase shift. On Earth, these algorithms worked, successfully producing interference patterns with periods commensurate with the programmed phase step.

In-orbit, the same algorithm did not produce clean cyclic interference patterns, thus undermining at least one of the embedded assumptions. When run, the algorithm produces less predictable interference patterns, as shown in Fig. 2.28. Data from both tests was taken ≈ 43 s after turning the PAs on: a moderate "warm-up" period but not long enough to let the system "settle" completely. This points to significant changes to the CMU phase function(s).

Degradation Metrics Over Time

These anomalies, and system performance in general, are captured by parameters like t_{knee} , the time between system turn-on and the PA supply current drop, T_{knee} , the temperature of the AP when PA supply current drops⁴, and I_{eq} , the supply current after a sufficiently long time (seven minutes).

Low values for t_{knee} and T_{knee} indicate the RFIC's absolute temperature rise or sensitivity to temperature rise; both could be indications of system degradation. Likewise, low I_{eq} could serve as an indication of system degradation.

⁴As demonstrated in Fig. 2.27, plotting PA current against temperature, it appears that PA current drops in response to the crossing of some threshold. This threshold is called T_{knee} .




Figure 2.28: (a) On-Earth chip-chip interference test performed directly prior to launch, and (b) in-orbit test performed on March 21, 2023.

These metrics were tracked during simple, long-duration tests performed over the mission lifetime to characterize long-term system degradation. In Fig. 2.29, it's clear RFIC1 T_{knee} and I_{equil} drop over time, with cumulative PA usage. RFIC2 appears much more sensitive to abnormalities, with the supply current dropping immediately upon being turned on. Thus, T_{knee} is the nominal board temperature when the system is turned on; for this reason, we see no long-term degradation on RFIC2.

Thermal metrics and results presented in Fig. 2.15 demonstrate different levels of degradation:

1. Short-term, reversible output power degradation is the primary degradation, observed in the mission's first tests and responsible for significant output power decreases during testing.

2. Long-term degradation, observed in the decline of thermal metrics over time (Fig. 2.29), exacerbates output power decreases and appears to worsen over time. Additionally, an $\approx 23\%$ decrease in peak achievable rectified power over eight months (Fig. 2.15) quantifies long-term power output degradation.
3. Individual, core-level degradation was observed in lower output amplitudes in six cores in Fig. 2.15.

Symptom #3 (the six degraded cores) may be the driving factor in Symptom #2 (long-term degradation.) Additionally, Symptoms #2 (long-term degradation in thermal metrics and output power) and #3 (individual core damage) may be the result of normal flight-related ageing, or the prolonged effects of the issue(s) causing Symptom #1.




Figure 2.29: Long-term system degradation as indicated by two metrics, T_{knee} and I_{eq}

Anomalies: Likely Cause

Anomalies likely stem from a single issue: poor thermal contact between the RFIC and the heatsink. Though designed to operate at slightly above chassis temperature, the RFICs, due to damage to the thermal bond with the aluminum heatsink pedestal, progressively heat up when PAs are turned on. The rate and asymptotic chip temperature are functions of thermal bond degradation—RFIC2 might perform generally worse because its bond is degraded further. This means that RFIC2 crosses thermal thresholds faster and stabilizes at a higher temperature overall.

As the temperature of the RFICs rise, PA output amplitudes and CMU phase functions vary, thus introducing instability to the optimization. Though only select elements' phases are digitally varied at once, the other elements' output phasors vary as the chip overheats. This produces erratic and non-monotonic optimization curves. Final rectified power values are lower even after the RFICs stabilize as a result of temperature-induced, degraded PA performance.

Nominal performance when only a subset of the quads is active, or an overall lower PA output power is programmed, is the result of lower power dissipation and, thus, lower steady-state temperature. Transitory temperature-induced performance degradation is nonlinear with temperature and, thus, degradation is only observed under full or mostly full-power operation. Performance restores when the chips cool down; turning the PAs off, pausing briefly, and then restarting restores full functionality until the chips reheat.

Inability to generate a sinusoidal interference pattern between chips is indeed the result of temperature-induced variations on the CMU phase functions. The functions at higher temperatures are different, and the underlying assumptions about what codes generate what phase offsets is broken. Time variations in core output also introduce aperiodicity.

A temperature threshold, T_{knee} , crossed also induces PA current supply drop, but only on supply line A. Supply line B experiences an increase in supply current draw as a result of temperature-dependent current draw to the underlying PA transistor.

This hypothesis is the result of extensive testing and verification, in space and in TVAC chambers on Earth using the EM. When thermistors were mounted close to the chip⁵, in an EM model of MAPLE, during on-Earth TVAC testing, temperatures

⁵The RFICs are thermally heat sunk on one side and flip-chipped to an interposer installed on the backside of the flexible PCB on the other. The closest location, thermally, to the chip is on the

in excess of 330K and 360K were measured for RFIC1 and RFIC2, respectively. RFIC1, at this temperature, behaved nominally while RFIC2 presented all of the aforementioned symptoms. All testing indicates that the fact that anomalous behavior presented on RFIC2 in both space and TVAC testing, on two different MAPLE models, is probably coincidental. RFICs 1 and 2 are constructed and wired almost identically and are arranged symmetrically vis-a-vis thermal pathways.

The discrepancy between chips is likely explained by discrepancies in thermal bond damage. When the EM was placed in a thermal chamber under ambient pressure, symptoms were replicated on both RFICs only when air temperature was in excess of 415K. This temperature is close to the limit of device models available from the IC foundry and likely high enough to induce the symptoms observed. An experiment on the ground where the RFIC was operated under ambient conditions with its heatsink intentionally removed showed similar discrepancies compared with a chip whose heatsink was intact.

A difference in the temperature necessary to replicate symptoms in TVAC and in thermal chamber (ambient pressure) can be explained by poor RFIC-epoxy bond quality. When the intended thermal pathway is compromised, heat flows through the flex array, which has higher thermal resistance and increases the temperature of the RFIC. This effect is more pronounced in vacuum due to the absence of convective heat transfer.

The RFIC-epoxy bond can be degraded both by vibration-induced failure and by workmanship flaws not revealed during environmental testing. Additionally, ground testing indicates contact pressure can affect the integrity of the bond, as expected based on the discussion in [44]. MAPLE's design does not apply pressure to the RFIC-epoxy bond, potentially increasing sensitivity to workmanship flaws and resulting in degraded performance.

2.6 Conclusion

MAPLE was designed to test wireless power transfer in space and stress the core technology to its extreme. Results of that 8 month mission were presented, and MAPLE achieved every mission goal. To the best of my knowledge, MAPLE is the first demonstration of end-to-end wireless power transfer in space. Demonstrations of focusing and WPT in space, rapid refocusing between different targets, and

front side of the flexible PCB, opposite the chip. This is where a thermistor was installed. During proper heatsink operation, the temperature difference between the RFIC and the thermistor is only 10K.

individual radiator core performance, and long-term system performance studies were presented as well.

Moreover, presented above are results from the on-Earth detection of wireless power transferred from a payload in orbit. Despite its small size and output power, the flexible aperture could focus power successfully to deliver power up to 20dB above the ground station noise floor and with the expected Doppler shift signature.

The results from this mission were highly informative and identified a number of key areas for improvement: namely, the system's thermal design. Likewise, mission successes serve as a proof-of-concept and a step forward for the practical implementation of SBSP while system weaknesses offer areas to build upon in future projects. This will be discussed in the next chapter.

Chapter 3

NEXT GENERATION FLEXIBLE WIRELESS POWER TRANSFER ARRAY USING A FULLY INTEGRATED 22-NM CMOS FDSOI RFIC

3.1 Introduction

The RFIC is the core-component of any wireless power transfer system, being responsible for the key functionalities of the RF phased array. These include:

- Synthesizing a coherent RF tone with low phase-jitter.
- Independent phase-shifting for each radiator element.
- Amplifying the coherent, phase-shifted tone at each element with high DC-RF efficiency.
- Maintaining low temperature for optimal operation.
- Long-term reliability over 10+ years of continuous operation.
- Provide RFIC telemetry (temperature, currents, power etc.) to higher-level system.

In addition to these requirements, for SBSP the system must be lightweight and low-cost. This motivates the use of CMOS, where all of these components, and the digital control/interface, can be integrated on the same die.

Beyond simply the performance of the RFIC, the main challenge of wireless power transfer is scale. Efficient wireless power transfer at a distance requires collecting most of the power radiated by an array, necessitating an increase in the array size as the transfer distance increases to counteract the diffraction limit. The array must also have a near 0.5λ radiator spacing and maintain its coherence over long distances to avoid power loss in the grating lobes (The so-called "sparse array curse"). The RFIC must be designed with scalability in mind, it is not enough to simply have one or even several operating coherently. A practical SBSP in GEO requires billions of chips operating in concert [23], an extremely challenging proposition.

There are two primary challenges with respect to the RFIC and scalability. The first is the reference distribution, where a master frequency reference must be distributed centrally to all elements in the array to maintain coherence across the whole system. This most-likely involves a combination of wireless, optical, and electrical distribution of the signal [23], which is beyond the scope of this discussion but is an active area of research. With respect to scalability at the RFIC-level, large arrays consisting of multiple tileable subarrays (tiles) centered around custom RFICs that locally synthesize high frequency signals from a distributed low frequency reference are attractive as the low frequency (100 MHz) reference is significantly easier to distribute across the flexible PCB. Each RFIC drives multiple radiators for practical component counts. The tiling can be virtual and the actual flexible substrate can remain contiguous.

The second challenge is calibration, in particular the phase-shifter calibration. The fundamental operation of a phased array, forming a coherent beam in a known direction, relies on precise control of the phase via a phase-shifter. This phase-shifter, due to random variations in the RFIC fabrication, must be calibrated to provide precise and linear phase steps. Although this can be done using an external calibration scheme, such as a far-field range, this is not feasible for the array scale needed for wireless power transfer. This will be addressed with the RFIC using a self-calibrating phase-shifter design, where no external calibration of the phase-shifter is required.

In this work, we demonstrate a lightweight and flexible array tile based on a custom, fully integrated beamformer RFIC designed in a 22 nm CMOS FDSOI process. The RFIC improves upon nearly every aspect of the 65-nm chip used in MAPLE, including DC-RF efficiency, frequency stability, and thermal dissipation. Using 8 identical RFICs, a 64-element flexible array for wireless power transfer at a distance is assembled and tested. The rest of the chapter is organized as follows. Section 3.2 covers the specifications used in planning and designing the RFIC. Section 3.3 discusses the design and measurements of the RFIC. Section 3.4 presents the design and measurement of the fully assembled flexible array, and finally section 3.5 concludes.

3.2 Specifications

Specifications are used to condense the system-level requirements to the metrics the RFIC must satisfy. To that end, the important specifications, their target value, and

the reasoning behind them are listed in Table 3.1. A discussion of each of these specifications is provided in the following subsections.

Table 3.1: Specifications for WPT RFIC in 22-nm CMOS FDSOI.

Specification	Target	Justification
Available Power	80-100 mW/channel	Energy Conservation
Overhead Power	10 mW/channel	~10% Efficiency Hit
Phase Noise	$< 2^\circ$ Integrated Jitter	Coherent Gain $>99\%$
Phase-Shifter Resolution	>5 Bits	Quantization Efficiency $>99\%$
Channels Per Chip	8	Heat Dissipation
Chip Area	$<4\text{mm}^2$	Cost

Available Power

These RFICs are intended for SBSP, as discussed in Chapter 2 and in [23]. The sandwiched flexible PV and phased array sheets have equal area (no sunlight concentration), so there is a fundamental limit to the power available to each RF element. This is also a function the RF frequency, as the element spacing must be 0.6λ to avoid grating lobes. In other words, the density of radiators is higher as the RF frequency increases, reducing the available power per element. This is plotted in Fig. 3.1 with the following assumptions.

- Area matched PV and phased array.
- The sunlight power density is 1362 W/m^2 .
- 20% PV efficiency.
- 0.6λ element spacing.

At very low frequencies, the available power explodes to extremely high values. This, in addition to large apertures required for focusing, limit the viability of low-frequency ($<1\text{GHz}$) arrays. At the other extreme at very high frequencies (100GHz), the available power is small ($<1\text{mW}$). 10GHz is a nice sweet-spot (for many reasons [23]), with approximately 87mW of power available per element. A range of 80-100mW is provided in Table 3.1 to account for any variation in the realized PV efficiency.




Figure 3.1: Available DC power per element vs. frequency.

Overhead Power

Related to the available power is the overhead power. This is the power consumed by everything that is not the PA, which include the PLL, phase-shifter, LO distribution, biasing etc. These should be minimized as much as possible as it doesn't directly contribute to the output power, which is mostly provided by the PA. Still, there is a minimum DC power consumption to satisfy certain requirements. For example, the phase noise of a VCO is proportional to the power consumption and cannot be made arbitrarily low. An overhead power of 10mW/channel was chosen to correspond to an approximately 10% hit to the DC-RF efficiency. This value is aggressive, but achievable given the other specifications the RFIC must meet.

Phase Noise

A phased array relies on coherence across elements to focus its power to a point in space. If that coherence degrades, then the focusing capability will degrade. In particular, if there is uncorrelated phase noise at each element then power in the main beam will be lost to sidelobes. So, the phase noise specification in the RFIC directly translates to the "coherent gain", or the gain of the noisy array as a percentage of that in an ideal array. For an infinite array with uncorrelated phase noise drawn from a Gaussian random variable, an rms phase jitter of 1.81° will result in a coherent gain of 99.9%. This is derived in Appendix A. A specification of 2° was chosen to provide 99% coherent gain with margin.

Phase-Shifter Resolution

Similar to the phase noise specification, quantization noise due to finite phase-shifter resolution will lead to a degradation in coherent gain. This can be analyzed like the phase noise specification and is provided in Appendix B. A phase-shifter resolution of 5-bits provides coherent gain >99% and is the target specification for this design. Higher resolution phase-shifters are possible, however they entail higher power consumption.

Channels Per Chip and Chip Area

The 65-nm chip in Chapter 2 integrates 16 channels onto a $2.8\text{mm} \times 2.8\text{mm}$ chip. Integrating more channels onto a given die is beneficial in terms of cost (fewer chips needed on the flex PCB), however is detrimental for heat dissipation. More energy concentrated in the small chip area makes it harder to remove the heat generated while keeping the chip cool. As a trade-off, the specification was revised to 8-channels per chip. Accompanying that, the area allotted to the chip should be no greater than 4mm^2 to keep the area per channel comparable with the 65-nm chip.

3.3 22-nm CMOS FDSOI RFIC Design

Process Node

The semiconductor technology chosen for the RFIC is a trade-off between cost and performance. Our system requires good RF performance, advanced digital capabilities, and low-cost. CMOS satisfies those requirements, however many different process nodes in CMOS are available. The previous chip utilized 65-nm bulk CMOS, which has a nice balance between all of these requirements. Its worth keeping in mind the 65-nm CMOS is, at this point, is 20+ years old and significant advancements in CMOS technologies have been made. This motivated the evaluation of another technology for the next generation WPT RFIC.

A excellent technology node for RF design is 22-nm CMOS FDSOI from Global-Foundries. The transistors are faster, with an f_{max} (maximum frequency of oscillation) of 350GHz [45] compared to 220GHz in 65-nm [46]. This increase in f_{max} , and correspondingly higher gain, improves the PA performance in particular. With respect to the analog/mixed-signal blocks, the device scaling and reduction in V_{dd} from 1V to 0.8V results in minimum-sized 65-nm inverters consuming 2.5x more power than those in 22-nm CMOS FDSOI. For all of these advantages, in addition to prior experience (see Chapter 6), 22-nm CMOS FDSOI was chosen for the RFIC. A discussion of the system design and different circuit blocks and their designs is

presented in the following subsections.

System Design




Figure 3.2: Flexible wireless power transfer array driven by fully integrated CMOS RFICs.

A block diagram and selected schematics of the 8-channel transmitter RFIC are shown in Fig. 3.2. A 100 MHz reference is distributed on the flexible array tile to each RFIC using microstrip transmission lines. A key difference in the frequency synthesizer compared to the 65-nm chip is a centralized generation scheme. That is, there is only one synthesizer, at which point all of the distribution, phase-shifting, and amplification are done at RF (10GHz). The 65-nm chip, in contrast, utilizes 16 10GHz PLLs (one per channel) to distribute the LO at a lower frequency (2.5GHz) on-chip. The downside of this approach is frequency stability, particularly pulling between all of the adjacent 10GHz VCOs in the PLLs. A centralized architecture minimizes the number of PLLs interacting with each other, greatly reducing this effect. Still, self-pulling from the PA to the centralized VCO must be mitigated, however this is simpler in a centralized architecture where the VCO and PAs can be placed farther apart from one another.

Next, the 10 GHz tone is distributed across the chip into 8 channels using on-chip

differential transmission lines buffered by CMOS inverters. Each channel has an independent voltage-controlled delay line (VCDL) based phase shifter implemented with a current-starved inverter topology (Fig. 3.7) for low power consumption. The output of the VCDL is then distributed to a transformer-tuned common source driving stage, which provides the input power to the power amplifier.

PLL

The PLL is responsible for synthesizing the 10GHz RF tone from a low-frequency (100MHz) crystal reference. This is done in two-stages, first with a 2.5GHz PLL which cascades into a 10GHz PLL. The block diagrams of each of those PLLs is shown in Fig. 3.3. A two-stage synthesizer architecture is used to mitigate the reference-spur induced deterministic jitter's impact on the coherence of the focused beam.




Figure 3.3: PLL1 block diagram (a) and (b) PLL2 block diagram.

The architecture of PLL1 is a standard charge-pump PLL with a programmable division ratio from 2 to 255. The loop has a sampling switch for further reducing the reference spur level. The voltage-controlled oscillator (VCO) is a cross-coupled NMOS/PMOS pair with an LC tank. Nearly every aspect of the PLL is digitally programmable, including the loop filter, VCO biasing, division ratio, PFD biasing, and a capacitor bank for coarse tuning of the VCO center frequency. The architecture of PLL2 is similar to PLL1 except it does not have a sampling loop filter and has only a fixed division ratio of 4.

The synthesizer is designed to consume approximately 6.7mW and have an integrated RMS phase jitter of approximately 1.8° . It should be noted that since the PLL is shared between all 8 channels on the RFIC its power consumption is amortized to approximately 0.8mW/channel.

VCDL Phase Shifter

Due to the speed of the transistors in 22-nm FDSOI, it is possible to use traditional CMOS mixed-signal circuits at the 10GHz RF frequency. The advantage of mixed-signal circuits is their reduced area on the chip, their improved isolation from the elimination of inductive tuning, and lower power consumption. The circuit topology chosen for the phase-shifter is a current-starved VCDL.

A block-diagram of the delay line is shown in Fig. 3.4 and the schematic of each delay cell in Fig. 3.5. The delay-line works by cascading 16 copies of individual delay cells, each of which is tuned using a control voltage V_{cont} . This control voltage adjusts the delay of the LO signal passing through each delay-cell, which is equivalent to producing a phase shift at the output. Dummy delay cells, which are exact copies of the tunable ones but fixed at their minimum-delay, are placed at the beginning and end of the cascade to provide uniform drive-signals and loading across the VCDL. This improves the matching of each delay-cell.

Finally, there is a cell that can be digitally controlled to provide an additional 180° of phase-shift. This is done by simply swapping the differential outputs, which can be easily done using a tri-state buffer network. This means the voltage-controlled delay cells only need to provide 180° of phase-shift to reach a full 360° , relaxing their design and reducing power consumption.

The advantage of a VCDL over a phase rotator, the traditional RF phase-shifter solution, is this design does not require I/Q signals which reduces the overall power consumption of the phase-shifter. This is especially true in the case of wireless

power transfer, where an I/Q LO is not inherently required like it is in communication transceivers.




Figure 3.4: VCDL Block diagram.




Figure 3.5: VCDL delay-cell circuit schematic.

The schematic of each unit cell is shown in Fig. 3.5. The delay-cell core is a pair of differentially driven inverters. The inverters are loaded by cross-coupled inverters to enforce a differential signal at the output. The sizing of the inverters are 16x of a minimum-sized, balanced inverter in the 22-nm technology. The larger sizing consumes more power, however it has better matching across PVT. The cross coupled inverters are sized at 4x a minimum inverter as balance between enforcing the differential signal at the output and not loading it too much.

The tunable delay is provided by the current sources in series the Vdd and ground of the 16x inverter in Fig. 3.5. By adjusting the gate bias of the NMOS and PMOS transistors, the inverter supply current is throttled, limiting the slew-rate of the output

waveform and thus introducing a delay. A keeper transistor that is a tenth of the size of the tunable transistor is placed in parallel and is always fixed to its maximum bias. This ensures that even if the control voltage turns the tunable transistor off, the delay-cell will not completely turn off and prevent the waveform from passing through to the output. The NMOS control voltage in Fig. 3.5 is generated by mirroring the PMOS control-voltage through a diode-connected NMOS transistor. The current-source transistors have a length of 60nm as opposed to the minimum length of 20nm supported by the technology. This is done to improve the matching between devices. This does not degrade the speed of the inverter as the current-sources are not connected to high-speed nodes.




Figure 3.6: VCDL output phase vs. control voltage over temperature (a), and (b) across process skews.

Table 3.2: VCDL phase range across PVT corners. Phase range is the phase-shift from tuning $V_{cont} = 0V$ to $V_{cont} = 0.8V$

Process Skew	Phase Range (0C)	Phase Range (50C)	Phase Range (100C)
tt	254°	261°	266°
ff	192°	202°	210°
ss	273°	277°	280°
sf	251°	259°	265°
fs	264°	268°	270°

SPICE simulations of the VCDL are shown in Fig. 3.6. The most important parameter that the VCDL must satisfy is it must provide at least 180° of phase shift so it can provide a full 360° when combined with the 180° cell. In all process and temperature corners, the VCDL achieves at least a 180° phase shift across the control-voltage tuning range (Table 3.2). The VCDL block consumes 2.8mW of power.

Although the VCDL satisfies the main specification, the VCDL requires calibration to provide precise and linear phase steps. This is due to both the PVT variation (how do we know what V_{cont} corresponds to what phase-shift a priori?) and the non-linear transfer characteristic of the VCDL. In principle this can be accomplished by generating a look-up table (LUT), however this would require an external measurement performed for each RFIC. This is time-consuming and is not scalable to the array sizes required for SBSP. We desire a phase-shifter that is self-calibrating to compensate for PVT without an external calibration scheme. This is accomplished using a delay-locked loop, which is discussed in the next subsection.

DLL Self-Calibration

One of the major challenges in maintaining coherence as the array is scaled is the phase shifter accuracy. PVT-induced phase error will degrade the focusing capability of the array, resulting in power loss. Although this could be compensated for by calibrating each phase shifter, complex calibration schemes become impractical as the array size increases. Here we introduce a self-calibrating VCDL-based phase shifter that uses a DLL to provide linear, uniformly-spaced phase steps across PVT, as shown in Fig. 3.7.




Figure 3.7: VCDL phase shifter reference voltage generation with DLL (left) distributed to VCDL channel phase shifters (right.)

Consider the left block of Fig. 3.7, whose product will be a bias voltage for the VCDL phase-shifter at each channel on the RFIC. Two identical, differential VCDLs are driven in-phase with respect to each other, with VCDL1's control voltage fixed at ground (minimum delay) and VCDL2's control voltage set by the DLL loop. The

DLL's phase and frequency detector (PFD) is fed by the positive output of VCDL1 and the negative output of VCDL2, which results in a built-in 180° phase shift in the DLL forward path. The loop will correct this built-in phase shift at VCDL2. Once the DLL is locked, the control voltage corresponds to a $180/N^\circ$ relative phase shift between VCDL1 and VCDL2 at each stage, where N is the number of unit-cells in the VCDL. Importantly, this assumes uniform matching of the delay cells.

The DLL control voltage is then routed to the phase shifter VCDL operating within each channel's signal path. By switching the control voltage of each delay-element in the phase shifter VCDL between the DLL control voltage and ground, the VCDL acts as a thermometer-encoded digital phase shifter with a temperature-robust LSB of $180/N^\circ$ that ideally only depends on transistor matching. The implemented 5-bit VCDL-based phase shifter consists of $N=16$ stages, with an additional 180° MSB provided by a MUX which swaps the differential outputs. By referencing the variable delay line to another at minimum delay, the phase shifter is able to provide a full 360° of phase shift with the smallest excess delay, minimizing the additive jitter within the signal path.

To mitigate the systematic error in the phase shifter stemming from the DLL static phase offset, the current-steering charge-pump was designed with op-amp-based current mismatch and charge-sharing compensation loops. To provide the option of a one-time factory calibration, a programmable IDAC was also included to compensate for the static phase offset by injecting a DC current into the loop filter, but this is not used in the full array operation to preserve the self-calibrating performance. The resulting 5-bit phase shifter provides sufficient resolution to minimize sidelobes in power beaming due to phase shifter quantization error (see Appendix B).

An exemplary transfer curve is shown in Fig. 3.8a. The DLL will calibrate (ideally) the phase-shift steps over PVT. That is, each code step will ideally introduce a $180^\circ/N$ phase shift regardless of PVT variation. It does not, however, calibrate the phase-offset due to variations in the minimum delay of VCDL1. The phase-shift curve in Fig 3.8a will drift up and down according to that variation. A simulation of the minimum delay phase vs. temperature is shown in Fig. 3.8b. The total phase deviation is less than 5° from 0°C to 100°C , which is significantly less than an LSB and should minimally affect the array's performance.




Figure 3.8: Self-calibrated DLL phase-shift curve vs. code (a) and minimum delay deviation over temperature (b).

LO distribution

After passing through the VCDL phase-shifter, the signal must be routed across the RFIC to each channel's PA. This is done using CMOS inverters, shown in Fig. 3.9. The inverters are sized up with a fanout of 2 (FO2) to drive the large transmission line capacitance before again buffering up the to PA's driver with a fanout of 2. A fanout of 2 was chosen to keep the edges of the waveform sharp over PVT, which limits the additive jitter of the buffers.




Figure 3.9: Inverter-based LO distribution. Only the half-circuit is displayed in this figure, in actuality the distribution is differential.

Power Amplifier

The power amplifier (PA) is the most critical circuit block in terms of the total DC-RF performance of the transmitter, consuming most of the DC power on the chip. The PA efficiency should be maximized while providing output power consistent with the available power in Table 3.1. A schematic of the PA is shown in Fig. 3.10. The PA utilizes a cascode topology to operate at an elevated, 1.4V supply rail which boosts the output power. The cascode transistor is a thick-oxide device to improve

the long-term reliability of the amplifier from the large voltage swing at the drain.




Figure 3.10: PA and driver schematic. Driver is the common-source amplifier on the left and PA the cascode amplifier on the right.

Load-pull simulations of the cascode transistor core (including parasitics) are plotted in Fig. 3.11. The target impedance (shown as the red dot in Fig. 3.11) results in an output power of 17.5dBm and an efficiency of $\sim 64\%$. Note that this does not include the loss from the matching network.




Figure 3.11: Cascode core load-pull power contours (a) and efficiency contours (b). The red dot corresponds to the nominal impedance from the matching network at the drain.

The goal of the matching network is the transform a standard impedance, in this instance 100Ω differential, to that shown in the load-pull contours. This must be done with as minimal insertion loss as possible to maintain performance. A shunt inductor was chosen to minimize the passive loss of the output matching network.

This is due to its high quality factor (~ 20) and a 1:1 impedance transformation ratio, which maximizes efficiency compared to high transformation ratios [47]. The network has a simulated passive efficiency of 93% and is compatible with the differential drive required for the pop-up dipole antenna that will be used for the flexible array PCB, similar to those in [33]. Opportunistic tuning at of the second-harmonic is done using the center-tap of the shunt inductor matching network to place the PA in class-J operation [48][49].

The PA is stabilized using capacitive neutralization of the common-source device and the common-gate device (Fig. 3.10.) The purely reactive stabilization scheme for the differential mode reduces the efficiency loss. Additionally, an RC biasing network is used at the gate of the cascode transistor for common-mode stability. The PA is unconditionally stable over PVT.

Simulated performance of the PA in saturation (0.8V peak-to-peak swing at each gate of the driver) is summarized in Table 3.3.

Table 3.3: 22-nm PA simulations results.

corner	P_{sat}	PAE @ P_{sat}	P_{dc} @ P_{sat}
tt @ 50C	17.1dBm	56.2%	90.3mW
ff @ 0C	17.5dBm	59.9%	94.1mW
ss @ 100C	16.5dBm	51.3%	86.7mW

3.4 RFIC Test Structure Measurements

To test individual circuit blocks, a separate chip with probed test structures was fabricated. A labeled die photo is displayed in Fig. 3.12. Included in the chip are the PLL, DLL, and PA + driver test structures. Additionally, there are copies of the digital shift-registers for controlling the chip and a copy of the PTAT circuit used for biasing and temperature detection. The latter in particular is shown in Fig. 3.13, and we see good match between measurement and simulation of the PTAT output voltage vs. temperature when the test chip was placed in a temperature chamber. This circuit is also present on the full RFIC and can be read as an analog voltage to determine the operating temperature of the chip.

A phase-noise plot of the PLL test-structure is provided in Fig. 3.14. The match between measurement and simulation is good, and measured integrated phase jitter is 1.87° . DLL locking was verified, however measurements of the output waveforms were not due to lower than expected signal swing from the buffers at the probe pad.




Figure 3.12: Test structure chip die photo. Chip dimensions are 1.62mm x 1.1mm.

Finally, probed measurements of the PA + driver test structure were done. With a calibrated input power 0.8V peak-to-peak sinusoidal swing at each gate of the driver, the measured output power is 16.4dBm and power consumption 100mW. Power consumption is higher than expected, and this is due to sub-optimal common-mode impedance being presented at the second harmonic. Specifically, with a GSSG probe the common mode impedance is well-maintained at 25Ω . In practice, the common-mode second-harmonic impedance will be nearly a short-circuit, which is optimal for PA performance. As we will see in the next subsection, this additional power consumption disappears when the PAs are used in a connectorized setup. Output power is minimally affected by this change in the common-mode second-harmonic impedance in simulation. Output power is also lower than expected. This is most likely due to the driver under-driving the PA input.

3.5 Connectorized Testboard Measurements

The full-chip die photo is shown in fig. 3.15. Circuit blocks and their locations are labeled, and the grid of dark circles in the photo are $80\mu\text{m}$ diameter, $180\mu\text{m}$ pitch solder bumps for flip-chip bonding. The chip is flipped onto a 11mm x 11mm interposer (Fig. 3.16), which fans out all of the RF, digital, and power connections from the chip to castellated vias on the edge of the interposer. For the RF traces, the interposer has a two-section microstrip matching network to compensate for the via discontinuity, and then routes a matched 100Ω line to the RFIC bumps. The




Figure 3.13: PTAT schematic (left) and PTAT voltage vs. temperature (right).




Figure 3.14: PLL phase-noise measurement vs. simulation. Loop bandwidth setting refers to the digital code for adjusting loop-filter for optimal phase-noise.

interposer is connected to a motherboard by soldering the castellated vias. The motherboard fans out all of the RF traces to SMA connectors for characterization on a VNA and spectrum analyzer. A picture of that motherboard is shown in Fig. 3.16.

The VCDL-based phase shifter is tested on the connectorized board that houses the complete RFIC. The relative phase between two channels is measured using a VNA. The phase shifter accuracy is compared with the DLL and without (fixed voltage bias). The DLL control voltage was optimized for phase accuracy by adjusting 1




Figure 3.15: Full chip die photo.




Figure 3.16: Interposer (left) and connectorized motherboard (right) photos. Interposer dimensions are 11mm x 11mm.

LSB of the IDAC injected into the loop filter and similarly the constant voltage bias was adjusted for minimum phase error at room temperature.

Figure 3.17 shows the phase-shift curves of channels 4-8, demonstrating good matching across the chip with an average RMS phase step error of less than 1 degree. The chip temperature was then swept in a thermal chamber from 5° C

to 100° C and the phase shifter performance was measured in increments of 5° C. The results are shown in Fig. 3.18. The DLL-based control voltage results in significantly lower phase error over temperature, with a maximal rms phase error < 2° compared to 6° without the DLL. The RMS phase jitter of the entire channel was measured and shown in Fig. 3.18. The VCDL adds 26 fs of jitter when tuned to its max delay.




Figure 3.17: DLL phase shifter curve for channels 4-8.

Using the connectorized testboard, the total DC-RF efficiency of the chip was measured. The PAs consume approximately 721 mW (90mW per channel, in line with simulation) from a 1.4V supply and the rest of the electronics on the chip consume 95 mW from a 0.8V supply for a total power consumption of 816 mW. The output power of each channel is +16.4 dBm from the probed test structure measurement, resulting in a measured total chip DC-RF efficiency of 42%. The PA efficiency (output power divided by power consumed by just the PAs) is 48%. These are compared with the widely available efficiencies of PA literature in Fig. 3.19 [50]. The PAs achieve the highest reported efficiency in CMOS at X-band (8-12GHz). Even using the total DC-RF efficiency, performance is competitive with prior art.

3.6 8x8 Flexible Array Prototype

The assembled 8x8 element flexible array is shown in Fig. 3.20. Eight RFICs are used to drive the 64 elements. The flex PCB is an 8.4 mil thick, 4-layer polyimide board. The antennas are differential, ground-plane backed popup dipoles. Performance of the array is initially validated on a far-field range, where beam steering from -45 to 45 degrees in azimuth and elevation is demonstrated (Fig. 3.21). The measured EIRP at broadside is 53.5dBm.




Figure 3.18: (a) Phase noise at PA output. (b) Channel #4 VCDL RMS phase error with and without DLL calibration over temperature.




Figure 3.19: Plot of PA efficiency vs frequency. Prior art shown in blue (data from [50]), PA shown in red, and full RFIC in green.

Scanning asymmetry in Fig. 3.21 and lower EIRP (4dB) than expected is mostly due to the active reflection coefficient (2.5 to 3dB) causing impedance mismatch which gets reflected to the PA. This can be fixed by retuning the antennas. The remaining discrepancy is likely due to additional loss in the PCB / interposer transmission lines and matching networks.

To validate the power-focusing capability of the array while demonstrating its flexibility, the array tile was conformed to convex, flat, and concave frames. For each shape, a focusing algorithm [34] was run to optimize the power at a near-field scanner probe. Then, near-field scans of the power pattern were taken and are shown below the corresponding shape in Fig. 3.22. Power recovery is verified with a




Figure 3.20: Assembled 64-element flexible array front (left) and back (right). Active array size is 14.4cm x 14.4cm.




Figure 3.21: Azimuth and elevation far-field scans of 64-element array.

rectenna board tuned to 10 GHz. The 64-element flexible array radiates 850 mW of total power, and the 56.25 cm² rectenna delivers 385 mW of DC power to an optimal load, with a total end-to-end efficiency of approximately 7.2%. For visualization purposes, the optimal rectenna load is replaced by a multi-LED load, as seen in Fig. 3.22.

To verify the formation of a focal point, scans of the radiative near field (Fresnel region) were taken at different distances and are shown in Fig. 3.23. The location of the beam waist near the rectenna position clearly shows the array focusing in the radiative near field.

Once the focused states have been saved, the array can be refocused to an arbitrary point in space using a computationally efficient refocusing algorithm [34]. However, the efficacy of the algorithm is predicated on an accurate (and linear) phase shifter. Thus, the DLL-based phase shifter is evaluated in terms of its refocusing efficiency,




Figure 3.22: Power beaming in concave, flat, and convex shape configurations.

with efficiency defined here as the received power from on-the-fly refocusing normalized by the received power after running a closed loop optimization.

A heat map of the refocused beams from broadside steered in a square of side length 20 cm around broadside is shown in Fig. 3.24. The plots also display the refocusing efficiency for the VCDL-based phase shifter with DLL-based control voltage (first percentage) versus the refocusing efficiency for the VCDL-based phase shifter with a fixed external control voltage (second percentage). It should be emphasized that the DLL is uncalibrated, whereas the fixed control voltage uses the calibrated voltage found in Fig. 3.18, although applied to chips with different process variations. On average, the VCDL-based phase shifter utilizing the DLL has a refocusing efficiency of 84.2% vs. 76.7% for the fixed external bias.

3.7 Conclusion

A proof-of-concept scalable and flexible phased array tile at X-band driven by custom CMOS RFICs was designed and tested. The array was validated in a wireless power transfer system and achieves state-of-the-art efficiency. This technology provides a path towards similar technologies in communications and sensing.




Figure 3.23: Near-field scanner power plot cuts at varying distances from array. The 3dB contour (in y-direction) is shown in the bottom plot.




Figure 3.24: Near-field scanner power plots of refocused beams. The first percentage is the refocusing efficiency with the DLL and the second without.

MUTUAL POWER OPTIMIZATION

4.1 Introduction

For SBSP to be economically viable, the total DC-to-RF conversion efficiency must be maximized while minimizing the weight of the array [23]. These simultaneous specifications require new solutions to existing challenges that traditional techniques cannot satisfy. One example is Maximum Power Point Tracking (MPPT). Consider the situation in Fig. 4.1, which shows the IV characteristics of a photovoltaic cell under arbitrary illumination.




Figure 4.1: Typical IV curve of a photovoltaic cell. The point (V_{MPP}, I_{MPP}) is the maximum power point of the cell.

When a short-circuit is applied at its terminals, the maximum current I_{sc} is drawn from the cell. If an open-circuit is applied then the maximum voltage V_{oc} is seen. In both of these cases, zero power is drawn from the cell. There exists an intermediate impedance $R_{MPP} = V_{MPP}/I_{MPP}$ at which maximum power is drawn from the cell. This is shown in the orange curve of Fig. 4.1. Any deviation from this operating point will result in performance degradation. Importantly, the maximum power point is a dynamic parameter that changes with varying illumination from shadowing or sun-angle. Therefore, the RFICs in a SBSP system must dynamically adjust their DC power draw to operate at the maximum power point of the cell. In the next section, the traditional solution to this problem, a DC-DC converter, is investigated.

4.2 Maximum Power Point Tracking Using a DC-DC Converter

Consider the schematic in Fig. 4.2. The equivalent impedance looking into the RFIC supply (in contrast with the output impedance at RF) is R_{out} . In general, R_{out} does not equal R_{MPP} , the optimal impedance for PV output power. Traditionally, a buck-converter is inserted in-between the RFIC and the PV cell to transform R_{out} to R_{MPP} at the input.




Figure 4.2: Maximum power point tracking with a buck converter.

The ideal buck-converter has the well-known voltage transfer function given by the following equation [51]

$$V_{out} = D * V_{in}, \quad (4.1)$$

where D is the duty-cycle of the buck converter. Assuming the buck-converter has low loss ($P_{out} \approx P_{in}$), then

$$\frac{V_{in}^2}{R_{in}} = \frac{V_{out}^2}{R_{out}} \quad (4.2)$$

and thus using equations 4.1 and 4.2, the input impedance looking into the buck-converter is:

$$R_{in} = \frac{R_{out}}{D^2}. \quad (4.3)$$

Looking at equation 4.3, we see that the buck converter acts as an impedance transformation network, analogous to a transformer with a $D:1$ turns-ratio, which can be adjusted with the duty-cycle. So, the DC-DC converter can be used to "match" (a familiar concept in RF engineering) the output load to the MPP of the

photovoltaic cell, shown in Fig. 4.3. The load-line of R_{out} is plotted with the IV-curve of the PV cell. The achievable region is shown in gray. Note that R_{out} must be greater than R_{MPP} for the MPP to be within the achievable region of a buck converter. The opposite would be true for a boost converter.




Figure 4.3: PV IV curve with output resistor load-line. The gray shaded region of the plot represents the achievable impedances using a buck-converter.

The PA dominates the RFIC power consumption (Chapter 4), so it will dominate the supply impedance seen by the DC-DC converter. To evaluate this, the 22-nm CMOS FDSOI PA from Chapter 3 is simulated with a constant input drive while sweeping the Vdd. Those results are shown in Fig. 4.4 and 4.5. The output power scales quadratically with the Vdd supply, which aligns with the reduction in the voltage swing at the output.

The efficiency is fairly constant, but does reduce as the Vdd is scaled to lower voltages. Part of this is due to suboptimal biasing which does not scale with the supply, which could be corrected for with a dynamic biasing scheme. Finally, the equivalent impedance looking into the PA supply is shown in Fig. 4.5. The impedance is relatively constant over the simulated supply range. These results confirm that the PA power draw can be scaled by simply changing the supply voltage without significantly degrading the efficiency or modifying the equivalent supply impedance.

DC-DC converters are the de facto solution for MPPT [52] and are compatible with the RFICs in a SBSP system. However, there is a major downside to DC-DC




Figure 4.4: 22-nm CMOS FDSOI PA output power and efficiency vs. Vdd.




Figure 4.5: 22-nm CMOS FDSOI PA output power equivalent supply impedance vs Vdd.

converters with respect to SBSP. DC-DC converters are typically limited to MHz switching frequencies due to losses associated with higher frequencies [52]. As a result, the inductor in Fig. 4.2 must be large (several μH). A low-profile inductor (which will have higher loss vs. larger alternatives) has a weight of approximately 0.35 g [53]. Assuming that each RFIC, which drives 8-channels, requires its own DC-DC converter, this corresponds to an areal mass of 13.5 mg/cm^2 . Given that the areal mass of a cost competitive system is approximately 28.4 mg/cm^2 [23], this additional weight is not feasible for practical SBSP systems, and must be eliminated.

4.3 PA Duty Cycling

The inductor dominates the weight of a DC-DC converter, and must be removed for SBSP to be economically viable. Other components such as the switching electronics can be readily integrated onto the RFIC, or for the capacitors are inherently smaller and low-mass. Consider the schematic in Fig. 4.6a where inductor has been removed from the switching converter. Applying a test current to the input, the equivalent input impedance can be found by solving for the test voltage V_t . As C_d grows larger, particularly if the RC time-constant is much larger than the switching period, then it appears as a voltage source in the steady-state. This is a good approximation since we require C_d to stabilize the voltage supply rail. This gives us the circuit in 4.6b, which is used for the rest of the analysis.

When the switch is open, charge is dumped into the capacitor. When the switch is closed, the capacitor discharges some its stored charge onto the resistor. For the circuit to be in steady-state, the charge must be balanced between the open and closed switching states.

$$Q_{open} = -Q_{closed} \quad (4.4)$$




Figure 4.6: (a) Inductorless switching converter, and (b) equivalent circuit with very large C_d .

The charge in the open and closed switching states is given by the following set of equations

$$Q_{open} = DTI_t \quad (4.5)$$

$$Q_{open} = \left(I_t - \frac{V_t}{R_{out}} \right) (1 - D)T, \quad (4.6)$$

where D is the duty cycle and T the switching period. Plugging equations 4.5 and 4.6 into 4.4 and solving for V_t/I_t yields

$$\frac{V_t}{I_t} = R_{in} = \frac{R_{out}}{D}. \quad (4.7)$$

Note that in this scheme, no voltage transformation between the input and output is achieved, as is the case with the DC-DC converter. However, this is acceptable for our system as long as the PV is co-designed with the RFIC such that V_{MPP} aligns with the nominal operating voltage. This can be done by exploiting series-parallel combinations of typical PV junctions.

Additionally, the output voltage is pulsed rather than continuous wave (CW). This is also fine, and simply means the RF voltage is modulated as shown in Fig. 4.7. The PAs will operate efficiently during the on-time, and then turn off and consume no power, so the efficiency is the approximately the same as CW operation. The rectenna must have enough bandwidth to capture the broader transmit-signal, but given the switching frequency will be a few MHz compared to a 10GHz center frequency, this is readily achievable.

There are two distinct advantages of duty-cycling vs a DC-DC converter. The first is supply voltage of the PA's while tracking the MPP. The PA operates at V_{MPP} , which is roughly equivalent to the built-in potential of a PN junction and has a logarithmic dependence on I_{MPP} . So, even when the available power from the PV is significantly lower than nominal ($\sim 10\%$), the supply voltage of the PA's stays roughly the same. This is in contrast to a DC-DC converter, which must reduce the PA supply voltage in proportion to the reduction in available power from the PV. Eventually, this collapses PA supply and reduces efficiency, shown in Fig. 4.4. We refer the simultaneous maximization of both the PV efficiency via MPPT and the PA efficiency afforded by duty-cycling as mutual power optimization (MPO).

Secondly, the elimination of the inductor with duty-cycling not only reduces the weight, but also improves the efficiency compared to a DC-DC converter. Analyzing Fig. 4.2, the power flows from the PV through a switching device and an inductor before terminating at the PA supply. In contrast, power only flows through a switching device 4.6 for duty-cycling, eliminating the ohmic loss of the inductor entirely.




Figure 4.7: (a) PA duty cycling equivalent supply impedance, and (b) incorporation in the PA.

4.4 Hardware Implementation

To verify PA duty-cycling for MPPT, two experiments were run. The first was done in a laboratory environment using an emulator model of a PV cell. The second was done outside using a silicon PERC PV cell powered with sunlight powering the array. In both experiments, the array has 16-elements at a spacing of 0.6λ at 10 GHz and is driven by a custom 65-nm CMOS RFIC [24]. A block diagram of that system is shown in fig 4.8.

An on-chip switch in series with the PAs (Fig. 4.8) is controlled by an off-chip microcontroller to adjust to duty cycle of the PAs. Due to the limited speed of the SPI interface, the switching speed is limited to 16.2 KHz. The duty cycle D has a range $D \in [0.16, 0.87]$ and the decoupling capacitance (C_L) is 300 μ F. The decoupling capacitance can be reduced by going to higher switching frequencies, which can be done by integrating the switching logic onto the RFIC.

Laboratory Experiment

The experimental setup is shown in Fig. 4.9. The 16-element array is first focused at a horn antenna terminated by an RF power meter. The PA supply of the RFIC is connected to a solar-cell emulator (Fig. 4.9a), which is equivalent to a current-limited DC supply in parallel with a discrete diode. Adjusting the current limit shifts the IV curves up and down, as shown in Fig 4.9b and corresponding to varying levels of sunlight. The solar cell emulator sufficiently approximates the PV behavior while allowing for controlled experiments in an indoor setting.

Next, the duty-cycle of the PAs is swept for different currents, shown in Fig. 4.10. There is a clear peak in the RU power vs duty-cycle. It should be emphasized that




Figure 4.8: Generating Unit (GU) and Receiving Unit (RU) block diagram.




Figure 4.9: (a) PV emulator schematic, (b) measured IV curves of PV emulator, and (c) experimental setup.

if the PAs were connected to a voltage supply, the maximum power would always occur at 100% duty-cycle. The reduction in power in Fig. 4.10a at higher duty-cycle is a result of a collapse in the voltage from the PV due to its inability to provide enough current.

The RU power at optimal duty-cycle is extracted and plotted against I_{sc} (photocurrent) in Fig. 4.10b. This curve is compared with a baseline measurement (100% duty-cycle) where no power point tracking was attempted, and with an extrapolated

DC-DC converter. The DC-DC curve was extrapolated assuming 90% efficiency, which is optimistic given the low-voltage operation of the system. The extrapolation was done by measuring the MPP of the PV emulator and mapping it to the measured IV curve of the RFIC's PA supply. For each point on the PA IV curve, the power at the RU was recorded.




Figure 4.10: MPO laboratory experiment results.

The RU power with duty-cycling is significantly better than the baseline and DC-DC converter, converging with the baseline at higher photocurrent due to the MPP exiting the achievable region with duty-cycling (Fig. 4.10). Notably, the DC-DC converter has much lower RU power at low photocurrents. This is because the DC-DC converter reduces the PA voltage during MPPT. Invariably, this leads to a reduction in the PA efficiency as the V_{dd} gets too low, as seen in Fig. 4.4. In contrast, duty-cycling only modifies the current draw while allowing the PAs to still operate at V_{MPP} , which remains steady over varying photocurrent due its logarithmic dependence on I_{sc} .

Outdoor Experiment

To validate duty-cycling in a more realistic environment, the solar emulator was replaced with an array of Passivated and Rear Contact (PERC) polycrystalline silicon cells. The whole system was placed outside and measured every 30 minutes between 8:30 A.M and 12:30 P.M. The setup is displayed in Fig. 4.11. The RU in this experiment is a rectenna rather than a calibrated power detector. Power varies over the course of the day, with the maximum occurring at approximately noon.

The experimental results are plotted in Fig. 4.12. Due to the nonlinear efficiency curve of the rectifier (Fig. 4.11), an additional efficiency benefit is seen at the RU with




Figure 4.11: MPO outdoor experiment setup.

duty cycling due to the higher peak power compared to a DC-DC converter. Thus, the duty-cycled output will operate higher on the rectifier efficiency curve [54]. The duty-cycled PA delivers 26% more energy over the four hour measurement window compared with no duty cycling.




Figure 4.12: MPO laboratory experiment results.

Conclusion

MPO, a technique which dynamically duty-cycles the PAs to track the maximum power point of a PV cell, was presented and a proof of concept demonstrated in hardware. Not only does MPO introduce minimal mass, in contrast with a DC-

DC converter, it also achieves higher efficiencies by eliminating the inductor loss and operating the RU at higher peak powers. By integrating all of the switching functionality onto the RFIC, the technique can be scaled to larger arrays and operate at higher switching frequencies.

TRANSMITARRAYS FOR WIRELESS POWER TRANSFER

5.1 Introduction

The most significant challenges when designing long distance wireless power transfer systems stem from their extreme size. The diffraction limit, a fundamental physical limit in electromagnetics, dictates that the smallest diameter within which power can be focused is proportional to the distance from the source and inversely proportional to the source's diameter. The Airy disk radius equation [55] illustrates this relationship for a uniformly illuminated transmitting circular aperture of radius a (radius is used here rather than diameter as is typical). At a distance R from a source transmitting at wavelength λ , the main beam's radius (measured at the first null) is given by

$$r_0 = \frac{0.61\lambda R}{a}. \quad (5.1)$$




Figure 5.1: Transmitarray wireless power transfer system. The transmitarray, acting as a lens, refocuses the divergent beam from the transmitter.

The total integrated power over radii smaller than r_0 is roughly 84% of the total power in the pattern. Accordingly, to capture 84% of the transmitted power, the transmitter or receiver must scale at least linearly with the distance [4, 34]. For example, a transmitter operating at 10GHz in geostationary orbit (GEO) ($\sim 36,000$ km above the Earth's surface) requires a 1.6 km diameter to focus its main beam onto a 1 km diameter receiving station on Earth [23, 34].

This scaling challenge is further exacerbated by the number of array elements in a phased array of this size. Phased arrays' efficient spatial power combining and

ability to create dynamic focal points using electronic steering makes them an ideal choice for wireless power transmitters [24, 26, 27, 34]. However, they require dense element spacing ($\sim \lambda/2$) to avoid grating lobes (the so-called “sparse array curse”), leading to extremely large element counts. Avoiding grating lobes is especially important for wireless power transfer since they lead to major efficiency reductions.

Furthermore, elements must be synchronized to the same frequency source and their relative phases must be maintained over all operating conditions. Complex distribution networks are required to provide a phase-stable, low phase-noise frequency reference to elements in such arrays. This challenge also exists for retroreflective arrays where an absolute phase reference is assumed (sometimes implicitly) in order to perform the necessary phase conjugation¹. Work has been done to address challenges associated with large-scale synchronization of phased arrays beyond typical RF distribution networks. These include optical distribution techniques for low loss distribution over long distances [56] and wireless techniques for distributed array synchronization and timing alignment [57]. Still, the prospect of synchronizing large numbers of elements, potentially billions of elements in the case of space solar power [23], is daunting.

One way to address this challenge is to refocus the power from the transmitter (hereafter called the active array) on its way to the receiver using a transmitarray, working as a passive relay, as shown in Fig. 5.1. The active array is focused to a transmitarray, which acts as a microwave lens [58], reshaping the wavefront and refocusing the beam to the receiver (or to yet another transmitarray to form a multi-stage relay chain). Since the transmitarray refocuses the active array’s beam, the active array can be made significantly smaller without increasing the rectenna array’s size. This is desirable for two reasons. First, unlike the active array, the transmitarray requires no synchronization since the incident wavefront is the transmitarray’s reference. Second, the transmitarray is significantly cheaper than the active array, lacking much of the latter’s active components and complexity. Note that the transmitarray is a passive relay; it does not introduce new energy to the beam nor does it modify the beam’s frequency. The only electronic components required for the transmitarray are electronically tunable phase shifters and basic digital control.

In this work, we propose a space solar power concept at 10 GHz consisting of a multi-

¹In fact, it may even be more challenging in retroreflective arrays where an absolute reference is needed instead of a stable yet arbitrary phase reference as used in [34].

orbit constellation of phased arrays in geostationary orbit (GEO) and transmitarrays in medium Earth orbit (MEO). 10 GHz is chosen following the trend of recent wireless power transfer research conducted at X-band [3, 7, 23, 26, 31, 34, 59]. Using physical optics and the technoeconomic analysis in [23], a cost model of the system is developed. It is shown that using transmitarrays for space solar power can reduce the system's levelized cost of electricity (LCOE) considerably for certain active array diameters.

The remainder of the paper is organized as follows: Section 6.2 introduces our cost model based on the results in [23] and the methods used to accurately simulate the proposed system, and discusses a key scenario in which transmitarrays significantly lower the system's LCOE; Section 6.3 presents a static transmitarray hardware prototype and measurements illustrating its ability to refocus power in the context of wireless power transfer; finally, Section 6.4 presents our conclusions.

5.2 Cost Analysis

Transmitarray Constellation

Consider the system shown in Fig. 5.2 composed of one active array, one transmitarray and one rectenna array. As the transmitarray orbits the Earth, it will periodically approach the line-of-sight path between the active and rectenna arrays, cross that path, and then move away, eventually being eclipsed by the earth. For the small orbital segment where the transmitarray is more or less aligned with the other two arrays, transmission through the transmitarray can occur as intended (see Fig. 5.2ab). Otherwise, both the arrival/departure angles at the transmitarray and the additional distance traveled between the arrays will lead to energy losses that make the transmitarray virtually useless (see Fig. 5.2c).

To avoid this energy loss, we propose a constellation of N_{ta} equally-spaced transmitarrays occupying the same orbit. As one transmitarray moves away from the active array, another transmitarray will approach so that when the first moves out of range, the active array will quickly refocus its beam to the next transmitarray (see Fig. 5.3). While this will increase the system's total cost, it will decrease its LCOE. For example, if a single transmitarray can only be used for 1/5 of its orbit, using a transmitarray constellation of 5 will increase the energy beamed to Earth by a factor of 5. Meanwhile, the system's total cost increases by a factor much less than 5.

We also propose increasing the number of active and rectenna arrays to match the number of transmitarrays (see Fig. 5.3). If only one active array and rectenna




Figure 5.2: Cartoon showing (a) power transfer when transmitarray is aligned with direction of power transfer, (b) power transfer when transmitarray is not well-aligned with the direction of power transfer, and (c) power transfer with a transmitarray constellation—the active array can now beam through a transmitarray that is both closer and better aligned with the direction of power transfer.




Figure 5.3: Cartoon showing a full constellation including active arrays, transmitarrays, and rectenna arrays. As the transmitarrays move through their orbits, the active arrays continuously adjust their beams to focus on the nearest transmitarray.

array are used, only one transmitarray will ever be in use at a given time, leading to significant overhead. This overhead can be eliminated by matching each unused transmitarray with its own active array and rectenna array (of course, the specific transmitarray used by a given active array will change as the transmitarrays orbit). This triple constellation will decrease the LCOE even further: a constellation of 10 active arrays and rectenna arrays used along with 10 transmitarrays increases the total energy output of the system by 10 relative to a system with 10 transmitarrays but only one active and rectenna array. However, the total cost of the system clearly increases by less than a factor of 10 meaning LCOE must decrease.

A full constellation is shown in Fig. 5.3 where there are 6 of each array. Each active array will always steer to follow the closest transmitarray. As the transmitarrays orbit the Earth, the active arrays will periodically refocus their beams to the next transmitarray.

Unit-Area Cost Model

Since we ultimately want to minimize some cost function, we relate the area of each array to its total cost using a per-unit-area cost. For the active and rectenna arrays, we rely on the technoeconomic analysis presented in [23] (see Table 2, “middle” column) to calculate unit area costs for the active array and the rectenna array (ground station). The unit area costs are calculated by dividing the total cost of each array ($\text{COST}_{\text{total}} = \text{COST}_{\text{initial}} + 15 \cdot \text{COST}_{\text{annual}}$) by the area of each array. For the active array we obtain $c_{\text{tx}} = 354 \text{ \$/m}^2$, and for rectenna array, $c_{\text{rx}} = 121 \text{ \$/m}^2$. These values naturally assume that all costs scale linearly with area which is approximately true as long as none of the arrays become too large or too small.

Because transmitarrays are not part of [23], we use the active array cost breakdown as the basis for our transmitarray cost model due to the many similarities between the transmitarray and active array hardware. In Table 5.1 we break down the transmitarray costs into initial, launch, and annual upkeep costs (transmitarray module replacement) and show how they are derived from the active array cost model. The reduction in WPT costs and mass is due to a lack of RF integrated circuits, pop-up antennas and heavy DC transmission lines. Assembly, testing costs and failure rate are also reduced since the transmitarray is significantly less complex than the active array. Finally, mass per launch increases because the transmitarrays are located in MEO instead of GEO.

One part of the transmitarray that will be more expensive than the active array will be radiation shielding due to the higher radiation effects in MEO versus in GEO. However, because the vast majority of the transmitarray is passive, the additional shielding that will be required for the electronics is small.

Results

As developed in [23], the system’s LCOE and the system’s total cost are two possible cost functions. When transmitarrays are used, the system’s total cost is always greater due to the need of a constellation as discussed in the preceding section. On the other hand, adding a transmitarray constellation *can* reduce the LCOE. Of course, if only LCOE is considered, the largest possible active array should be used. Clearly this is impractical since there will be some maximum active array size that will be practical, likely due to synchronization challenges. Since research on synchronization schemes for large-scale arrays is ongoing, we leave the active array size as the independent variable in our analysis. For each active array size, we

Table 5.1: Unit-area cost of the transmitarray (T.A.) derived from the active array (A.A.) costs provided in [23].

	Unit	A.A.	Multiplier	T.A.
Initial Costs	\$/m²	293		102
PV	\$/m ²	95	0	0
WPT	\$/m ²	95	0.25	24
Structures	\$/m ²	34	1	34
Spacecraft	\$/m ²	28	1	28
Assembly	\$/m ²	28	0.5	14
Testing	\$/m ²	14	0.2	3
Mass	kg/m²	0.28		0.14
PV	kg/m ²	0.090	0	0
WPT	kg/m ²	0.089	0.4	0.035
Structures	kg/m ²	0.032	1	0.032
Spacecraft	kg/m ²	0.073	1	0.073
Launch Costs	\$/m²	36.4		12.0
Cost/launch	\$M	7.0	1	7.0
Mass/launch	kg	54500	1.5	81750
Upkeep Costs	\$/m²	24.7		4.2
Lifetime	yr	15	1	15
Failure rate	%/yr	0.5	0.5	0.25
Total Cost	\$/m²	354		118

compute the system's cost with and without transmitarrays to determine whether, for that particular active array, transmitarrays are advantageous.

Before discussing the optimization procedure, however, a model of the electromagnetic between the different arrays is required. Due to the large (in wavelengths) electrical size of the arrays, a physical optics model can be used to compute the electromagnetic fields at the transmitarray and on the ground throughout the orbit, as shown in Fig. 5.4. This technique is discussed more in [60].

The program used to compute the LCOE takes the following parameters:

- Size of the active array: independent variable
- Desired capture efficiency of the transmitarray: $\eta_{TA,cap} = 0.7$ (The capture efficiency specifies the percentage of the active array's radiated power that is




Figure 5.4: Diffraction patterns of the active array at the transmitarray plane (middle) and of the transmitarray at the rectenna plane (bottom) for five points as the transmitarray moves through its orbit (top). The part of the active array's pattern that falls outside the transmitarray is greyed out; The power in the circle is reradiated by the transmitarray towards the rectenna array. In this case, $N = 9$ so that each transmitarray has an effective orbital span of 40° .

incident upon and captured by the transmitarray.)

- The transmitarray's internal efficiency (includes losses in phase shifters and antennas): $\eta_{TA} = 0.9$.
- # of active arrays/transmitarrays/rectenna arrays in the constellation:
 $N = 10$.
- Operating frequency: 10 GHz.
- Array Orbits/Locations: Active array is in GEO, rectenna array is directly below the active array, transmitarrays are in MEO with an altitude half that of GEO.
- Array Orientations: Active array and transmitarrays are pointed at nadir.
- The number of points in the transmitarray's discretized orbit.

- The efficiency of the system due to factors other than non-unity capture efficiencies. This includes photovoltaic losses, DC to RF conversion loss, rectenna losses, etc.: 6% (see [23])².

We neglect the effects of polarization mismatch and additional losses at large steering angles. The latter can be due to antenna element patterns and impedance mismatch resulting from mutual coupling changes. These factors have the potential to increase losses beyond what is considered here but can be mitigated. Polarization mismatch is typically mitigated by using circular polarization. Issues at large scan angles can be reduced by restricting the maximum scan angle; in this system, that is accomplished by using a larger constellation.

Given the parameters above, the program computes the minimum achievable LCOE. The program chooses the minimum LCOE because the rectenna array size is not fixed; for every rectenna size there is a different LCOE. The transmitarray size, however, is indirectly fixed by specifying η_{TA} . A summary of the algorithm is below.

1. The effective transmitarray orbital span is calculated as $360^\circ/N$. The orbital span is the part of the transmitarray's orbit actually used for computation since the diffraction patterns are periodic.
2. The active array's focused and steered diffraction pattern is computed at the transmitarray's location for each point in the transmitarray's orbital span (see Fig. 5.4 middle row).
3. The transmitarray's captured power as a function of its diameter is computed from the diffraction pattern.
4. The size of the transmitarray is determined to meet the specified capture efficiency parameter.
5. The transmitarray's focused and steered diffraction pattern is computed at the rectenna array's location for each point in the transmitarray's orbital span (see Fig. 5.4 bottom row).
6. The rectenna array's captured power as a function of its diameter is computed from the diffraction patterns and, from that, the LCOE as a function of the diameter.

²6% is based on the middle column in [23] Table 1 with η_{ground} and η_{geom} excluded.

7. Finally, the rectenna array size is chosen to minimize LCOE.

This process is repeated for each active array diameter.

A critical part of this process is computing the diffraction pattern of the active array on the transmitarray, and of the transmitarray on the rectenna array. Exemplary patterns are shown in Fig. 5.4. We assume that the active arrays, transmitarrays and rectenna arrays are periodically spaced around Earth. While the rectennas (ground stations) will not necessarily be periodic, making this assumption significantly reduces computation time since the periodicity of each array type means the patterns themselves are periodic as the transmitarray moves through its orbit. The patterns need only be computed for one active array/transmitarray/rectenna array set and only for the part of the transmitarray's orbit where it is closest to the active array, i.e., the effective orbital span $360^\circ/N$.

In Fig. 5.5a, we show the minimum LCOE, $\text{LCOE}_{\min}$, that can be achieved by the 3-tier constellation. We also plot the minimum LCOE that can be achieved in a system of just one active array and rectenna array but no transmitarray (dashed black line). The percent increase/decrease in the LCOE compared to the system without transmitarrays is shown in Fig. 5.5b. This is a critical comparison because it tells us which systems (in terms of active array diameters) benefit from using a transmitarray. We conclude that introducing transmitarrays to the system lowers the LCOE for active array diameters below roughly ~ 1000 m for this particular set of parameters.

Intuitively, this is motivated by the inverse scaling of a diffraction pattern relative to the size of its source. Since the active array illuminates the transmitarray, the transmitarray diameter is inversely proportional to the active array diameter to the first order (see Fig. 5.6a). Similarly, the rectenna array scales inversely with the transmitarray and so it must scale directly with the active array (see Fig. 5.6b). If the active array is made very large, then the transmitarray will become very small requiring the use of a very large rectenna array. At some point, the transmitarray will be small enough that it would be better to just beam directly to earth.

An important side-effect of this analysis is that systems where a transmitarray is advantageous will have very small spot-sizes on the ground, as shown in Fig. 5.6b. This can be extremely valuable for locations where land is available at a premium. A small rectenna array would also be useful for military applications and natural




Figure 5.5: Cost analysis results. A.A. = active array. T.A. = transmitarray. For each A.A. diameter we plot (a) the minimum achievable LCOE both with and without T.A.'s and (b) LCOE percent increase/decrease relative to the LCOE without T.A.'s.

disaster relief. Small, temporary rectenna arrays could be constructed very quickly and torn down equally fast when no longer needed.

The downsides of using transmitarray systems are twofold. First, the transmitarray's size effectively substitutes for the active array's size. Thus, the transmitarray sizes are extremely large, on the order of 1-3 km in diameter as shown in Fig. 5.6a. However, we note that constructing a transmitarray of this size is much more feasible than constructing a similarly sized active array since a synchronization scheme is not required. Second, to achieve the LCOE's presented in Fig. 5.5a, the system requires a constellation of 10 arrays at each level. As the constellation size grows beyond 10, the gap between the two curves in 5.5a grows, but only to a point. If the constellation size is reduced sufficiently, the gap will close and there will not be any point where transmitarrays are beneficial. This is indicative of a trade-off between




Figure 5.6: The transmitarray (a) and rectenna array (b) diameters at minimum LCOE in Fig. 5.5a.

total system cost and complexity, and the system's LCOE.

5.3 Transmitarray Prototype

An important assumption underlying the cost analysis presented above is that transmitarrays can be made low-cost and lightweight. To demonstrate the feasibility of a lightweight and affordable transmitarray, we fabricated a static transmitarray using a low-cost, standard, commercial flexible PCB process and low-cost off-the-shelf materials.

Each layer consists of a flexible single-layer PCB each fabricated on a $25\ \mu\text{m}$ thick polyimide substrate with 1/2 oz. copper. A gap of 2 mm is maintained between each of the layers using low-cost EVA foam with a measured dielectric constant of 1.12 and negligible loss. Electromagnetically, the transmitarray is a $16 \times 16\ \lambda/2$ unit cell grid where each cell consists of two identical patch antennas (the outer two layers)




Figure 5.7: Transmitarray prototype stack-up. The outer layers are patch antennas with a slot in the middle for coupling.

coupled through a slot (Fig. 5.7) [61, 62]. The phase gradient was introduced by varying the slot length between 4 mm and 9 mm across the transmitarray. A photo of the fabricated prototype is shown in Fig. 5.8.

To measure the transmitarray's focusing capability, we illuminated it with a 256-element phased array located 2.5 m directly behind it. At a distance of 1.3 m on the other side of the transmitarray, a near field scanner was used to measure the power distribution across a roughly $1\text{m} \times 1\text{m}$ square both with and without the transmitarray in place (Fig. 5.10). The measurement setup is shown in Fig. 5.9. The transmitarray increases the peak power density by 3.5 dB compared to the phased array by itself, and reduces the -3dB beam area by approximately 75%. Integrating the power in Fig. 5.10ab, the transmitarray reduces the total power incident on the $1\text{m} \times 1\text{m}$ area by 10%, equivalent to 90% efficiency.

To demonstrate the transmitarray's ability to boost power delivered to a rectenna array receiver, a 5×5 , 0.5λ pitch ($7.5\text{ cm} \times 7.5\text{ cm}$) rectenna array [59] was placed at the center of the main-beam using the probe scanner. The phased array was focused to the rectenna using a closed-loop feedback algorithm to maximize power transmission [31]; a plot of the voltage at the rectenna while focusing is shown in Fig. 5.11. The peak rectified voltage, V_r , which feeds into a $54\ \Omega$ load, increases by a factor of 1.54, corresponding to a 2.37x increase in the rectified power.

5.4 Conclusion

At the system level, the cost analysis shows that transmitarrays can increase the feasibility of space solar power by lowering the system's LCOE, thus lowering the cost of power delivered to consumers. In terms of flexibility, transmitarrays also reduce the size of the rectenna array on the ground, increasing the system's




Figure 5.8: Fabricated transmitarray prototype: individual layers (left) and assembled array (right)




Figure 5.9: Prototype transmitarray measurement setup.

potential for emergency and military applications where quick access to electricity is essential. Finally, we conclude with a practical demonstration showing that low-cost, low-mass transmitarrays can significantly increase the amount of power transferred in a wireless power transfer system.

We believe that our results demonstrate that transmitarrays can play an important role in the implementation of a economically viable space solar power system. More importantly, they may prove key to enabling space solar power from a technological perspective by decreasing the synchronization requirements, one of the most significant challenges facing space solar power today. The work presented in this




Figure 5.10: Normalized power densities (a) without the transmitarray and (b) with the transmitarray. Intensity values are normalized to the maximum intensity measured *with* the transmitarray. Black ellipses correspond to the -3db contours of each plots respective peak.

paper can serve as a starting point for fruitful discussion, research, and development leading us closer to a vision of clean solar power delivered from space.




Figure 5.11: Rectenna output voltage while focusing. The rectenna DC output is terminated by a 54Ω load in both cases.

Chapter 6

SCALABLE D-BAND ACTIVE RADIATOR TILES IN 22-NM CMOS FDSOI

6.1 Introduction

The D-band frequency range (110-170GHz) can enable high-speed wireless communications and high-resolution sensing due to its wide spectrum and physically small apertures [63, 64, 65, 66, 67, 68, 69, 70, 71, 72]. This comes at the cost of high path loss and degraded device performance, especially in CMOS technologies. The phased array is a promising solution to these issues [63, 64, 65, 66, 67, 68, 69, 70, 71, 72]. Phased array transmitters utilize efficient spatial power-combining and aperture gain to achieve an N^2 improvement in EIRP, where N is the number of elements in the array. High EIRP systems can be constructed from relatively low-power, CMOS-based elements [65] [66]. The challenge common to D-band systems and RF WPT is that the arrays must be scaled to large element counts to provide sufficient SNR for communications/sensing links.

A key system-level design choice for D-band phased arrays is the level of on-chip integration [73]. Beamformer architectures have the lowest level of integration, where the RFIC provides only phase shifting and amplification at D-band [64]. All of the reference distribution, routing, and antennas are fabricated on a PCB or package. To completely suppress grating lobes in the array, the element pitch must be approximately $\lambda/2 = 1.07\text{mm}$ at 140GHz. At this scale, the element pitch becomes comparable to the size of the RFIC, making routing and fanout to the antennas difficult. The antennas themselves also scale down in dimensions, putting increased demand on the PCB tolerances. Combined with the need for impedance-controlled transmission lines and sensitive transitions from chip to PCB, the cost and complexity of D-band phased array PCBs and packages are high [74, 75, 76, 77, 78, 79, 80, 81].

Increasing levels of on-chip integration significantly relax the PCB requirements. Fully integrated frequency generation and integrated antennas [63] confines all the high-frequency operations to the chip, lowering the cost of the PCB. Integrated antennas, however, suffer from low efficiency due to the high permittivity and low resistivity of conventional doped silicon [82][83]. A full survey of on-chip radiators




Figure 6.1: Hemispherical Dielectric lens mounted to a silicon RFIC (left) and ray-tracing diagram (right) of slot antenna on the RFIC.




Figure 6.2: Aperture coupled integrated radiator.

at D-band is documented in [84].

Simple post-processing techniques that improve of on-chip radiator efficiency have been demonstrated, the most common being dielectric lenses [82][85][86]. Dielectric lenses and their effect on integrated antennas can be intuitively visualized in Fig. 6.1. A slot antenna located at the center of RFIC radiates mostly in the silicon, generating rays that propagate outwards over a wide range of angles. The lens is index matched to the silicon, so the rays propagate without refraction at the Si-lens interface. The lens is designed such that those rays are (mostly) collimated when they hit the lens-air interface, generating useful radiation. This is in contrast when no lens is present, where rays incident at the Si-air interface beyond a critical angle

θ_c will be fully reflected and end up as lossy surface modes [82]. The downside of dielectric lenses, however, is they limit the field-of-view of the antenna, reducing an array's scan range without the use of mechanical steering or mechanical adjustment of the lens [87]. The bulky lens also results in large element spacing, and is not amenable to densely spaced arrays.

Another popular post-processing technique is substrate-thinning combined with mode-suppression, which have demonstrated radiation efficiencies as high as 52% [63]. Substrate thinning is a common post-processing technique in semiconductor manufacturing and does not add significant cost. These techniques, however, do not scale when used in an array that requires dense element spacing. Mode-suppression techniques require multiple antennas on-chip with fixed phase relationships to limit the excitation of dominant surface-wave modes. Once the array is steered, these surface modes are no longer suppressed due to the excitation phase gradient, resulting in lower efficiency. The wirebond packaging in [63] also increases the inter-chip element pitch well beyond $\lambda/2$.

Techniques involving advanced post-processing steps such as localized back etching (LBE) selectively removes silicon around the antenna, improving efficiency [88, 89, 90]. This requires complicated, non-standard processing steps that increase cost and complexity, reducing their attractiveness.

Off-chip antennas fabricated on a high-performance substrate will typically achieve the highest efficiency, with radiation efficiencies of 85% demonstrated in the literature [91, 92]. The transition loss from the chip to the antenna, however, must be included in the overall efficiency when comparing with on-chip designs. Recent work has shown transition losses of 0.9dB [93] with LTCC and 1dB with a glass substrate [94]. When combined with radiation efficiencies of 85%, however, the overall efficiency can drop to 67.5-70%. This loss is significant despite using high-performance and high-tolerance packaging techniques that greatly increases fabrication and assembly costs.

Aperture coupling, long used for microwave PCB antennas [95], balances the trade-offs between on-chip and off-chip antennas. Depicted in Fig. 6.2, the aperture coupled integrated radiator utilizes a slot to excite the radiative mode of a patch antenna off-chip. The slot is directly driven by on-chip transistors, and when combined with integrated frequency generation, confines the high-frequency routing to the chip. As previously discussed, normally the electromagnetic radiation from the slot would be predominantly sucked into the silicon due its extremely high




Figure 6.3: Chip-tile and PCB assembly with flip-chip aperture coupled antennas. Four differential PAs excite four sets of differential slots on-chip. The slots drive off-chip patch antennas at $\lambda/2$ pitch fabricated on the PCB. Chips are tiled with narrow spacing to form large, dense arrays.

permittivity. By exploiting the resonance of the patch antenna, however, energy is efficiently coupled upward, avoiding the lossy silicon.

No physical routing to the antenna is required due to the electromagnetic field coupling, eliminating the need for costly chip-to-PCB transitions and transmission lines. The patch antenna is fabricated in a low-loss substrate and achieves good efficiency, broad bandwidth, and wide steering range.

Prior works utilizing aperture coupling at D-band [66][71] have achieved radiation efficiencies of 45% and 58%, respectively. These works, however, use non-standard assembly approaches for the patch superstrate and chip alignment. Standard and commercially available assembly is desirable due to its cost-savings and simplified manufacturing. Additionally, [66] uses wirebond packaging, which require chip-to-chip bonds between tiles to maintain a narrow element spacing. This leads to lossy on-chip routing for the reference, data, and power distribution.

This chapter introduces flip-chip aperture coupled antennas as a solution to relax package requirements using standard fabrication and assembly techniques. Combined with electronics fabricated in a 22-nm CMOS FDSOI process, the active radiator tile provides a scalable solution with narrow element spacing for large arrays at D-band.

6.2 Antenna Design

Flip Chip Aperture Coupled Antennas




Figure 6.4: Flip-chip aperture coupled antennas side-view.




Figure 6.5: Flip-chip aperture coupled antennas stackup.

The flip-chip aperture coupled antenna concept is shown in Fig. 6.3, 6.4, and 6.5. A chip with four sets of differential slots and $\lambda/2$ pitch is flipped and reflowed onto a commercial PCB. The PCB is fabricated on a standard $127\mu\text{m}$ (5 mil) Duroid substrate with $35\mu\text{m}$ (1 oz) copper. Duroid was chosen for its well-established




Figure 6.6: 2x2 radiator E-plane and H-plane gain simulations at 140GHz.

low-loss and uniform permittivity, however low-loss substrates such as RO3003, Megtron 7 or Astra MT77 are compatible as well.

The bumps are $150\mu\text{m}$ pitch, $80\mu\text{m}$ diameter lead-free solder offered by most foundries. The flip-chip interconnects provide the chip's microwave reference, power, and control. The D-band LO is generated by on-chip electronics from the off-chip microwave reference. Flip-chip bonding also provides alignment between the slots and patch antennas located on the top-side of the PCB. Thus, the patch antenna alignment and the chip's interconnects are accomplished using standard flip-chip bonding and commercial PCB fabrication, with no D-band routing or transitions required. Only factory thinning of the substrate to $300\mu\text{m}$ is required to suppress backside radiation.

The tile can be scaled, as represented in Fig. 6.3, by flipping additional chips onto a PCB with the appropriate pads and patches. The small lateral spacing enabled by flip-chip bonding, in contrast with wirebonding, allows for close spacing between tiles. An inter-chip element spacing of 0.6λ is the target for this work, although 0.5λ is possible with modifications to the PCB design, the challenge being the narrower gap between chips. The 2x2 tile size provides flexibility in the construction of different array sizes and pitches, in contrast with wafer-scaled systems where the number of channels on a given chip is large. Small die sizes also improves yield [96].

A full 3D electromagnetic model of the PCB, bumps, aperture coupling, and the silicon substrate (Fig. 6.5) is used to design the radiator. This holistic structure is required to precisely model the antenna performance due to the close proximity between the chip, interconnects, and PCB. The PCB itself is composed of two metal layers. The top layer includes the patches and a ground-plane. The bottom layer includes all the pads for flip chip bonding and the interconnects. Additionally, the cross structure on the bottom layer both shields the on-chip electronics from antenna coupling and distributes Vdd to the chip. Each arm of the cross has a width of 0.38mm and extends 1.58mm across the chip.

The simulated H-plane and E-plane cuts are shown in Fig. 6.6. The 2x2 array has a simulated gain of 7.8dBi, 45% radiation efficiency, and a front-to-back ratio of 9.5dB. The silicon substrate has a low resistivity of $10\Omega\cdot\text{cm}$ and accounts for most of the loss in the antenna. Detailed parametric studies of the antenna design are presented in the next subsection.

Parametric Studies

The patch's dimensions generally behave as expected based on the theoretical analysis from [97]. The patch's length determines the resonant frequency and the width is used to tune the impedance. The nominal dimensions for the patch's length and width are $610\mu\text{m}$ and $250\mu\text{m}$ as displayed in Fig. 6.3.

The slot, located on the top metal layer of the chip, does not have a shield on the lower metal layers for isolating it from the silicon substrate. A shield limits the efficiency and bandwidth of the antenna. As a result, the silicon substrate thickness requires tuning to optimize the power coupled into the patch antenna. A simulation of the antenna gain and the power radiated into the front hemisphere (front-side radiation) normalized to the total radiated power is shown in Fig. 6.7. Although gain and front-side power are maximized for very thin silicon, in practice this is difficult to achieve with conventional back grinding due to the mechanical instabilities it introduces [98]. The gain is minimized for a $\lambda/4$ ($\sim 150\mu\text{m}$) thickness and increases up to $\lambda/2$ ($\sim 300\mu\text{m}$). A $300\mu\text{m}$ thickness was chosen for its strong front-side radiation, suppressed backside radiation, good efficiency, and mechanical robustness.

The bump height after bonding, shown in Fig. 6.4, results in an air gap between the chip and the PCB substrate. Furthermore, the bump height is not well-controlled with lead-free solder and the antenna's sensitivity must be evaluated. A simulation




Figure 6.7: 2x2 radiator gain and front-side radiation vs. silicon substrate thickness at 140GHz.




Figure 6.8: 2x2 radiator gain and radiation efficiency vs. bump height at 140GHz.

of the antenna gain and efficiency is shown in Fig. 6.8. Both the gain and efficiency of the antenna improves with smaller bump heights. This is due to the stronger coupling from the slot to the patch antenna rather than the silicon substrate. Scanning electron microscopy indicates that in practice, the bump height reflows consistently to approximately $50\mu\text{m}$ with the current PCB design. More precise control over the bump height can be achieved using a copper-pillar bump if necessary [99]. Alternatively, the air-gap can be filled with a capillary underfill that acts as a matching layer between the chip and patch to improve performance. This is discussed in detail in the next subsection.

Similar to the bump height, a thinner PCB substrate generally improves the gain




Figure 6.9: 2x2 radiator gain and radiation efficiency vs. PCB substrate thickness at 140GHz.




Figure 6.10: 2x2 radiator gain and efficiency vs slot length (a) and Smith chart impedance (b).

and efficiency of the antenna (Fig. 6.9). For very thin substrates, the metal and dielectric losses result in lower radiation efficiency [97]. The optimal PCB substrate thickness is approximately $70\mu\text{m}$, but the final design was limited to $127\mu\text{m}$, the thinnest standard RT/Duroid 5880 laminate.

The length of the slot also influences the coupling into the patch. This is shown in Fig. 6.10a, where the gain and efficiency of the antenna improves as the slot length increases. A nominal value of $550\mu\text{m}$ was chosen to maximize the coupling while limiting area consumption on-chip. Additionally, the slot length adjusts the feed impedance, which is used to match the impedance of the on-chip PA without




Figure 6.11: 2x2 radiator gain degradation for varying x and y translations at 140GHz.

requiring a complex matching network. The feed impedance is plotted on the Smith chart in Fig 6.10b, and the nominal feed impedance is $45 + j68\Omega$. This is analyzed in more detail in section III along with the PA design.

Finally, an attractive property of aperture coupling is its resilience to translational misalignment. A simulation of the gain degradation vs. misalignment between the PCB and the chip is shown in Fig. 6.11. The y-axis is co-linear with the slot and x-axis orthogonal. Only a single quadrant of misalignment is shown in Fig. 6.11, but due to the symmetry of the structure the other three quadrants are equivalent. The results show that for a misalignment of $\pm 100\mu\text{m}$ in any direction, the antenna gain is still within 0.4dB of its peak value. Typical registration between the top and bottom layers is $25.4\mu\text{m}$ (1mil), more than sufficient based on this simulation.

Underfill Matching Layer

Underfill is an essential component of flip-chip interconnects, especially when the coefficient of thermal expansion (CTE) of the chip and PCB are mismatched. The underfill epoxy encapsulates the gap between the chip and substrate (Fig. 6.12), spreading the stress that accumulates at the bumps during thermal cycling. Underfill increases the mean fatigue lifetime by more than an order of magnitude [100]. Thus, for future practical implementations of the tile in this work, underfill will be necessary.

The underfill's electrical characteristics affect the antenna performance, as it alters the patch antenna's effective substrate. Using a commercial underfill from NAMICS Corporation ($D_k = 3.4$, $D_f = 0.016$), the gain vs. patch length was simulated in Fig.




Figure 6.12: Flip-chip aperture coupled antenna with underfill.




Figure 6.13: 2x2 radiator gain vs. patch length at 140GHz (a) and (b) gain vs. frequency (patch length = 610 μm).

6.13a. The peak occurs at a smaller patch length due to the higher permittivity of the underfill. The peak gain, however, is higher with underfill than without. The underfill's permittivity results in more coupling from the slot to patch, enhancing the gain and efficiency despite its higher loss tangent. Figure 6.13b shows the gain vs. frequency of the antenna with the nominal patch length of 610 μm . Although the antenna is detuned from 140GHz, it should still be within the electronics' bandwidth.

Stacked Patch Antennas and Backside Reflector

There are two promising modifications to the antenna architecture that improves the electromagnetic and system-level performance. These modifications are stacked patch antennas and the addition of a metal reflector on the backside of the silicon

substrate, shown in the diagram of Fig. 6.14.




Figure 6.14: Flip-chip aperture coupled antenna with underfill, stacked patch antennas, and backside reflector. Slot dimensions and patch width are the same as Fig. 6.3.




Figure 6.15: 2x2 stacked patch radiator pattern (a) and gain vs. frequency (b).

On-board and in-package stacked patch antennas have been demonstrated at mm-wave frequencies [101, 102], and offer wide bandwidths. These antennas can be excited via chip-to-PCB aperture coupling in the same manner as the architecture presented in Fig. 6.3. The antennas are not subject to the same relationship between PCB substrate thickness and efficiency shown in Fig. 6.9, permitting thicker multi-layer PCB designs without the same performance degradation. Thus, the stacked patch topology provides the additional benefit of another metal layer on the PCB, which can be used for routing reference distribution and data lines while maintaining tight element spacing for larger arrays.

The second modification is a backside metal reflector, in this case copper, placed on the silicon substrate and shown in Fig. 6.14. The metal reflects the waves in the substrate back into the patch's radiative mode, improving the efficiency. This requires the substrate be thinned to $200\mu\text{m}$ to maximize its performance. The reflector can also serve as a thermal heatsink for the chip, which is critical as the array size and heat dissipation increases.

Simulations of this antenna architecture are shown in Fig 6.15. The peak simulated gain is 10.3dBi and the radiation efficiency 63%. The front to back ratio is greater than 20dB and the gain bandwidth is approximately 40GHz. These simulations show significant performance improvements compared to the single patch, reflector-less design previously presented while still providing all the advantages of aperture coupling from the chip. Although measurements of this topology were not done for this work, the stacked-patch architecture provides a path towards further improvements in future designs.

6.3 Electronics Design

The chip integrates the beamforming, mm-wave LO synthesis, LO distribution, power amplification and antenna coupling onto a chip fabricated in a 22-nm CMOS FDSOI technology. The chip block diagram and associated circuit schematics are displayed in Fig. 6.16. All the circuits use a 0.8V supply. A 15.56GHz microwave reference frequency is synthesized off-chip and distributed on the PCB. The reference frequency is converted from single-ended to differential by an on-chip balun, which is also used for matching. The differential signal is amplified and distributed to four frequency multiplier chains.

The frequency multiplier chain (Fig. 6.17) is composed of two self-mixing triplers [63] with a driver in between for saturating the output of tripler #2. Schematics of the tripler circuits are shown in Fig. 6.16. Inter-stage matching and biasing are provided by transformers stacked on the top two metal layers of the chip. Inductive source degeneration is tuned to the second harmonic, which mixes with the fundamental at the gate to enhance the harmonic generation.

The tripler chain can also generate a phase shift at the output. This is done by adjusting the DC bias, V_ϕ , of tripler #1 in Fig. 6.16. The phase shift is multiplied by the subsequent tripler for full 360° control [103]. The tripler-based phase shifter requires no additional area or I/Q generation. A detailed discussion of the analysis and design of this tripler-based phase shifter is provided in section 6.4.




Figure 6.16: Chip block diagram and circuit schematics.




Figure 6.17: Channel block diagram and power breakdown.

The D-band LO frequency is amplified and distributed to the on-chip PA. The D-band buffers are pseudo-differential pairs with capacitive neutralization. The buffers are biased and matched using transformers. The D-band signal is distributed to the PA located directly at the center of the two slots. The slots are placed on the top aluminum metal layer of the chip. The metal slot-plane also distributes Vdd across the chip. The transistors' drains are connected to a short microstrip line that passes directly below the center of the slots and is shorted to the slot-plane, which provides Vdd to the PA.

Due to the tightly coupled electromagnetics and electronics, the PA, antenna, PCB, and chip interconnects are co-designed. The antenna dimensions (slot length, slot pitch, and patch width) are tuned to the optimal impedance of the PA. A small,




Figure 6.18: PA input distribution and antenna feed.




Figure 6.19: PA output power contours (in dBm) load-pull simulation.

single-turn shunt inductor is placed across drains of the transistor core to provide additional tuning. This inductor is incorporated into the 3D HFSS model of Fig. 6.5, whose broadband S-parameters are simulated and exported to a Touchstone format for SPICE simulations with the transistors. A diagram of the PA/antenna feed structure is shown in Fig. 6.18. The differential impedance of the antenna feed is approximately $45 + j68\Omega$. This is plotted on the Smith chart along with the power contours of the PA transistor core from load-pull simulations in Fig. 6.19. The maximum output power is 5.1dBm, with the antenna impedance resulting in an output power of 4.5dBm.




Figure 6.20: Tripler #1 and tripler #2 output power vs. tripler #1 gate bias. Simulation done using full RF chain in Fig. 6.16.




Figure 6.21: Full channel simulated PA output power and phase shift vs. tripler #1 gate bias (a), and output phase corner simulations (b).

The performance of a single channel is simulated at 140GHz. A complete electro-

magnetic model of the aperture coupled antenna is provided at the PA's output, and individual transformers/inductors are extracted separately. The simulated DC power of the channel varies between 70.3mW to 88.7mW depending on the phase setting (Fig. 6.17). The phase-gradient needed for beamforming results in the per-channel power consumption averaging to approximately 80mW for large arrays.

Tuning the gate-bias of tripler #1 for phase-shifting results in output ripple from varying conversion gain, which is shown in Fig. 6.20. The output ripple is 6.5dB for tripler #1, however this is reduced to 2.2dB for the D-band output from tripler #2. This is due to the gain compression of the intermediate driver (buffer #4) in between the two triplers in Fig. 6.17.

The PA output power and phase versus the control voltage is plotted in Fig. 6.21a. The peak output power is 4.5dBm per channel. The tripler-based phase shifter achieves a simulated frequency shift of $>360^\circ$ using the DC gate bias, with a small power variation of approximately 0.65dBm over all phase settings. The output phase shift for the extreme PVT corners are also shown in 6.21b. The phase shift of the corners is referenced to the nominal curve.

Finally, the simulated inter-element isolation is plotted in Fig. 6.22. The simulation includes the full electromagnetic simulation of the chip and the loading of the PAs at the antenna feed. The isolation is >20 dB for all element combinations over a bandwidth of 130GHz to 150GHz.




Figure 6.22: Element isolation simulation.

6.4 Tripler-Based Phase Shifter

Tripler Model

A critical component of phased arrays is 360° of independent phase-control at each element, providing the most flexibility when beamforming. However, it can be challenging to fit the phase shifters into a narrow $\lambda/2$ pitch. Here we outline a technique that generates a full 360° phase shift at each element without the need for I/Q generation or significant area overhead using the frequency triplers. It does this by exploiting two sources of third harmonic current content generated by the transistors. Interfering these two third-harmonic phasors results in a vector at the output of the tripler with a phase-shift. The relative strength of these two phasors is controlled by the DC bias of the tripler, which provides the tunable phase-shift. The advantage of this approach is the tripler, which is already required for synthesizing the D-band signal, serves double duty as a phase-shifter. This saves area and reduces power consumption.

The accompanying analysis expands upon the work in [104] by including the effect of a second harmonic at the transistor's source terminal and investigating the phase of the generated harmonic. Consider the situation in Fig. 6.23. This schematic represents the simplified operation of the frequency triplers in Fig. 6.16. The transistor is driven at its gate terminal by a sinusoidal source at the fundamental frequency with amplitude V_{AC} and a DC voltage V_g . The transistor source terminal is driven at the second harmonic with amplitude V_s and a phase offset θ with respect to the fundamental. It should be noted that, in practice, the second harmonic at the source is generated from the fundamental by the transistor non-linearity and the tuned tail inductor (Fig. 6.16). To simplify analysis, the transistor's I-V curve in Fig. 6.23 is approximated with a piecewise linear function. The drain current is given by the equation

$$\begin{cases} I_{ds} = g_m [v_{gs}(t) - V_t] & v_{gs}(t) \geq V_t \\ 0 & v_{gs}(t) < V_t, \end{cases} \quad (6.1)$$

where g_m is the transconductance of the transistor, $v_{gs}(t)$ the gate-source voltage, and V_t the threshold voltage.

Equation (6.1) is mathematically equivalent to multiplying $g_m v_{ov}(t)$ (where $v_{ov}(t) = v_{gs}(t) - V_t$ is the overdrive voltage) by a pulse function, $m(t)$, with duty-cycle d and time-delay τ . This function is referred to as the masking function, and is periodic




Figure 6.23: Transistor schematic for drain current harmonic analysis.




Figure 6.24: Masking function modulates $V_{ov}(t)$ to produce the rectified drain current.

with $T = 2\pi/\omega_0$. A visual representation of the masking function is plotted in Fig. 6.24. Assuming $V_{AC} \gg V_s$, then a closed form expression for the masking function duty-cycle and time-delay can be found

$$d \approx \frac{1}{\pi} \cos^{-1} \left(\frac{V_t - V_g}{V_{AC}} \right), \quad (6.2)$$

$$\tau \approx 0. \quad (6.3)$$

Thus, the time domain drain current is given by

$$g_m v_{ov}(t) \cdot m(t), \quad (6.4)$$

where the masking function, $m(t)$, has a duty-cycle and time-delay given by (6.2) and (6.3). Multiplication in the time-domain is equivalent to convolution in the frequency-domain. Thus, convolving the Fourier series of $g_m v_{ov}(t)$ and $m(t)$ and solving at the third harmonic results in,

$$I_{ds,3rd} = \frac{-a_1}{2}(g_m V_s)e^{j\theta} + \frac{a_2}{2}(g_m V_{AC}) + a_3 g_m (V_g - V_t) + \frac{a_4}{2}(g_m V_{AC}) - \frac{a_5}{2}(g_m V_s)e^{-j\theta}, \quad (6.5)$$

$$a_n = \frac{2}{n\pi} \sin(n\pi d). \quad (6.6)$$

Equation (6.6) is the well-known single-sided Fourier series coefficients of a pulse function [105]. Equation (6.5) can more conveniently be written as

$$I_{ds,3rd} = I_{int} + I_{up}, \quad (6.7)$$

where,

$$I_{int} = \frac{a_2}{2}(g_m V_{AC}) + a_3 g_m (V_g - V_t) + \frac{a_4}{2}(g_m V_{AC}), \quad (6.8)$$

$$I_{up} = \frac{-a_1}{2}(g_m V_s)e^{j\theta} - \frac{a_5}{2}(g_m V_s)e^{-j\theta}. \quad (6.9)$$

Intuitively, (6.7) can be interpreted as the summation of two nonlinear mechanisms for generating third-harmonic content. The first, I_{int} , corresponds to the intrinsic non-linearity of the transistor's IV characteristic. Under the condition $V_{AC} \gg V_s$, I_{int} has no dependency on the source terminal voltage. The second, I_{up} , corresponds to the fundamental at the gate terminal and second-harmonic at the source terminal mixing and up-converting to the third-harmonic at the drain. I_{up} is strongly dependent on θ , and this can be tuned to provide a monotonic phase shift at the output third harmonic of the drain current.

Plots of I_{up} for $\theta = 0^\circ$, 45° , and 90° and I_{int} are shown in Fig. 6.25. With the assumption $V_{AC} \gg V_s$, I_{int} remains unchanged with respect to the second-harmonic at the source terminal. I_{int} experiences a phase-inversion and a null at the threshold voltage of approximately 300mV, and is at its maximum at 100mV or 500mV. I_{up} , however, has no null, its phase is strongly dependent on θ , and is relatively insensitive to the gate terminal DC bias V_G .

SPICE simulations of the third harmonic current using the full transistor model compared with the analytical model from (6.7), equivalent to $I_{int} + I_{up}$ in Fig. 6.25, are plotted in Fig. 6.26. For ease of comparison in terms of the total achievable phase shift, the phase-maximum of the plots in Fig. 6.26 are normalized to zero. For an in-phase second-harmonic ($\theta = 0^\circ$), I_{int} and I_{up} combine coherently for $V_g > V_t$ and produce larger harmonic content than I_{int} alone. The third-harmonic content is maximized when the fundamental and second-harmonic are in-phase. Notably, in the presence of an in-phase second harmonic at the source terminal the phase still experiences an inversion at V_g slightly less than V_t due to the interference from I_{up} . The phase-inversion occurs over too small of a voltage range to be usable as a phase shifter when $\theta = 0^\circ$, however.




Figure 6.25: I_{int} and I_{up} magnitude (a) and phase (b) for $\theta = 0^\circ, 45^\circ$, and 90° . $V_{AC} = 400\text{mV}$, $V_s = 50\text{mV}$, and θ displayed in graph.

As θ increases to 90° , the interference is no longer perfectly in phase above V_t . This causes a reduction in the peak magnitude of the third-harmonic and a softening of the phase transition. Thus, tuning θ closer to 90° results in a more gradual phase transition, useful for a phase shifter, but comes at the cost of a lower phase shift range and slightly lower harmonic generation. This trade-off is summarized in Fig. 6.27. Figure 6.27a plots the maximum phase shift over the tuning voltage range of $[0, 500]\text{mV}$ vs. θ . Figure 6.27b plots the tuning voltage required to adjust the phase from 10% to 90% of its maximum phase shift, a measure of the sensitivity. It is clear from Fig. 6.27 that a more gradual phase-transition results in lower phase range.

Figure 6.28 plots the magnitude variation, or ripple, of the drain current third




Figure 6.26: Third harmonic magnitude (left) and phase (right) of the schematic in Fig. 6.23. $V_{AC} = 400\text{mV}$, $V_s = 50\text{mV}$, and θ displayed in graph.

harmonic over $V_g = [0, 500]\text{mV}$. The ripple reduces as θ increases towards 90° . This is related to the phase-inversion, and thus phase range, because the inversion coincides with a null in the magnitude of the drain current. For $\theta = 0^\circ$, the ripple ideally approaches infinity due to perfect cancellation between I_{int} and I_{up} . Ripple is undesirable for the phase shifter, as some phase settings will have conversion gain that is too low to saturate the subsequent tripler. These trade-offs between phase range, phase sensitivity, and ripple must be carefully considered in the design of the tripler-based phase shifter.




Figure 6.27: Maximum phase shift over $V_g = [0, 500]$ mV (a) and tuning voltage required to achieve 10% to 90% of the maximum phase shift (b).




Figure 6.28: Drain current third harmonic ripple over $V_g = [0, 500]$ mV.

Circuit Simulations

There are two primary design parameters for tuning the phase shift properties of the tripler. First is the tail inductance, which modifies both V_s and θ . Second is splitting the ratio of differential pair widths into two portions, one pair with inductive degeneration and one pair without. This is shown in Fig. 6.29b. This is equivalent to multiplying I_{int} by a weighting coefficient in (6.7), adjusting the strength of the intrinsic non-linearity relative to the up-conversion non-linearity. The width ratio in Fig. 6.29b is 1:1.

Simulations of the phase range, tuning range, maximum conversion gain, and output voltage ripple are plotted in Fig. 6.30a-d. The phase range decreases as the tail inductance increases, and the opposite is the case for the tuning range. This is expected based on prior discussion of the trade-off between the phase and tuning

A tail inductance of 97pH was chosen for the final design, and offers $>120^\circ$ of total phase shift at the third harmonic of the output drain current ($>360^\circ$ at D-band), a tuning range of 340mV, a maximum conversion gain of -2.9dB, and a ripple of 5dB. The ripple increases to 6.5dB in Fig. 6.20 when placed in the full RF chain due to additional variations, mainly from the reference distribution network.




Figure 6.29: Frequency tripler schematic with full degeneration (a) and partial degeneration (b).

The chips were fabricated and assembled with the PCB, photos of which are shown in Fig. 6.31. Due to the highly integrated structure of the radiator, probed measurements are challenging. Thus, a far-field measurement of the fully-assembled radiator is used to test its performance. Close match between simulation and measurement of the EIRP, pattern, and bandwidth verify the tile functions as expected.

The far-field measurement setup is shown in Fig. 6.32. The chip and PCB are placed on a 2D turntable for pattern measurements. The chip's 15.56GHz reference is synthesized and distributed off-chip by a signal generator. A 21dBi standard gain conical horn antenna is placed in the far-field (110cm) to measure the radiated D-band signal. This is down-converted by a VDI WR6.5 SAX-M with calibrated




Figure 6.30: Frequency tripler circuit phase range (a), 10% to 90% phase shift tuning voltage (b), maximum conversion gain (c), and output voltage ripple (d).

conversion gain. The power is detected at an IF frequency of 322.5MHz using a spectrum analyzer. The EIRP of the radiator is measured by calibrating the known gains and losses of the system in (6.10)

$$EIRP = P_{IF} - G_{horn} - G_{VDI} + L_{cable} + 20\log_{10} \frac{4\pi d}{\lambda}. \quad (6.10)$$

Single-Chip 2x2 Array

First, the phase-shifter was characterized by turning on two adjacent elements and sweeping the phase control voltage of one element while keeping the other at its nominal value of 0.4V. Measuring the received power on the spectrum analyzer, shown in Fig. 6.33, a pseudo-sinusoidal interference pattern corresponding to the phase shifter curve can be measured. This measurement, along with the expected simulated behavior, is shown in Fig. 6.34a.

The interference pattern cycles between strong destructive and constructive interference consistent with $>360^\circ$ of phase shift. The same measurement was repeated




Figure 6.31: Chip die photograph (a) and PCB (b).




Figure 6.32: Far-field measurement setup.

at 135GHz and 145GHz and is plotted in Fig. 6.34b. The interference pattern is stable over a broad range, with largest variation occurring at the extreme edge of the tuning range near 0.5V.




Figure 6.33: Spectrum analyzer output using VDI WR6.5 SAX-M frequency extender.




Figure 6.34: Two element interference pattern vs. tripler #1 gate bias voltage at 140 GHz (a), and at 135GHz, 140GHz, and 145GHz (b).

Next, the EIRP of a single-chip 2x2 tile was measured on the far-field setup. The phase-control voltages were optimized for boresight and the frequency swept by adjusting the reference signal generator. The phase-control voltages were kept constant during the frequency sweep. The results are shown in Fig. 6.35. The peak measured EIRP was +17.5 dBm at 138.2 GHz compared to a simulated peak of 18.3dBm. The measured 3 dB bandwidth was approximately 18.3 GHz (130.2 to 148.5 GHz). Good match between measurement and simulation also demonstrates

the stability of the tripler-based phased shifters over a wide frequency range. The total power consumption of the chip was 308mW, and the reference signal power was ~ 0 dBm.




Figure 6.35: Single chip-tile boresight EIRP vs. frequency.

The far-field pattern in the E-plane and H-plane were measured using the turntable (Fig. 6.36). The measurement and simulation results are in good alignment. The close match between measurement and simulation in these experiments (Fig. 6.35 and Fig. 6.36) confirms the performance of the system and validity of the modeling techniques used to design the tile. Beamsteering was also demonstrated in both the E-plane and H-plane, and those are plotted in Fig. 6.37.

Finally, the harmonic distortion of the tile can be measured using the same far-field setup¹. This was done by measuring the received power over the WR19 (40-60GHz), WR15 (50-75GHz), WR10 (75-110GHz), WR6.5 (110-170GHz), WR4.3 (170-260GHz) and WR2.8 (260-400GHz) frequency bands using the corresponding VDI SAX-M module. A carrier of 139GHz was used for convenience in fitting the harmonics within the waveguide bands. The resulting measurement is shown in Fig. 6.38.

Using the received power and (6.10), the harmonic rejection with respect to the EIRP of the 139GHz carrier can be computed, and those results are shown in the table of

¹This experiment measures EIRP, which includes the antenna's directivity. Additionally, the phases of the harmonics are not precisely controlled, and may result in directivity variation due to destructive/constructive interference between the radiating elements. With the current chip design, elements can only be turned on/off in pairs. To maximize the SNR and average out the phase error across elements, all four transmitting elements focused to boresight were used to measure the harmonic distortion.




Figure 6.36: Single chip-tile H-plane (a) and E-plane (b) patterns.

Fig 6.38. Harmonics are generated at integer multiples of the reference frequency, however only the 6th through 11th harmonics were detectable, with all others being below the sensitivity of the measurement setup. The minimum harmonic rejection is 36.7 dBc at the 6th harmonic (92.67GHz).

Single-Chip 2x2 Array With Commercial Underfill

To evaluate the performance of commercial underfill with the tiles, underfill ($D_k=3.4$, $D_f=0.016$) from NAMICS Corporation was placed into the air gap. This detuned the antenna to a lower frequency, however the behavior of the underfill can still be tested. The radiator can be retuned to 140GHz by adjusting the length of the patch antenna in a future board revision.

The measured and simulated EIRP vs. frequency is plotted in Fig. 6.39. The peak measured EIRP is 17.5dBm at a center frequency of 134GHz and a bandwidth of 14.6GHz (126.6GHz to 141.2GHz). The measured EIRP is nearly identical to




Figure 6.37: Single chip-tile H-plane steering (a) and E-plane steering (b) patterns.

the air-gap tile, with the expected detuning and bandwidth degradation due to the sub-optimal patch dimensions. Still, this measurement confirms the viability of commercial underfill for use in mm-wave antenna structures.

Two-Chip 2x4 Array

To verify the scalability and beamforming capabilities of the tile, a 2x4 array composed of two chips was assembled. The inter-chip element spacing is 0.6λ at 140GHz. For these measurements, underfill was not applied to the gap. The EIRP vs. frequency was measured on the far-field range and is plotted in Fig. 6.40. The peak measured EIRP is 23.5dBm, which aligns with the expected 6dB increase in EIRP compared to the 2x2 array. The measured 3dB bandwidth is 14.7GHz (128.8 to 143.5 GHz). The bandwidth is slightly lower than the 2x2 array, and the center frequency is detuned to a lower frequency. The reduction is most likely due to a $25\mu\text{m}$ layer of soldermask placed on the PCB below the patch antennas to improve flip-chip bonding reliability during reflow.

Measurements of the array's H-plane and E-plane patterns are shown in Fig. 6.41. As expected, the H-plane pattern is narrowed due the larger number of elements,




Figure 6.38: Far-field received power spectrum and harmonic rejection ratios. Harmonic rejection is referenced to the EIRP of the carrier at 139GHz.




Figure 6.39: Single chip-tile with underfill boresight EIRP vs. frequency.

whereas the E-plane is similar to Fig. 6.36b due to the same effective element count in that axis. The beam was steered in the H-plane from -40° to 40° in increments of 10° (Fig. 6.42). A scan loss of $<3\text{dB}$ is maintained over the -40° to 40° steering range using the tripler-based phase shifters. A comparison of these results with




Figure 6.40: Two-chip 2x4 array boresight EIRP vs. frequency.




Figure 6.41: Two-chip 2x4 array H-plane (a) and E-plane (b) boresight patterns.

other published works is provided in Table 6.1.




Figure 6.42: Two-chip 2x4 array H-plane steering patterns.

6.6 Conclusion

A proof of concept 2x2 D-band active radiator tile was designed in 22-nm CMOS FDSOI. The tile uses standard PCB and flip-chip assembly techniques to minimize the cost and complexity of the PCB. Additionally, D-band routing and transitions are eliminated using flip-chip aperture coupled antennas.

Integration of the frequency synthesis, beamforming, LO distribution, and aperture coupling combined with narrow pitches offered by flip-chip technology enables dense element spacing. Scaling was demonstrated with a 2x4 array using two of the tiles, providing a path towards large phased arrays at D-band.

Table 6.1: Comparison with other works at D-band.

Process	This Work	[64]	[65]	[66]	[68]	[69]	[70]	[71]	[106]
	22nm CMOS FDSOI	130nm SiGe BiCMOS	45nm CMOS SOI	65nm CMOS	55nm BiCMOS	45nm CMOS SOI	22nm CMOS FDSOI + InP HBT PA	45nm CMOS SOI	45nm CMOS SOI
Elements	8	4-1024 ^b	63	16	16	4	8 ^g	8	1
Chip Tiles	YES	YES ^b	NO ^e	YES	YES	NO	YES ^h	NO ^e	NO
No D-band Routing	YES	NO	NO	YES	NO	NO	NO	YES	YES
Independent Phase Control	YES	YES	YES	YES	YES	NO	NO ^g	YES	NO
Element Pitch	0.5λ 0.6λ^a	0.7λ	0.5λx0.57λ	0.5λx0.5λ	0.5λ	0.67λ	0.65λ	0.5λ	N/A
Area/element (mm ²)	1	0.9	1.29	1	0.69	1.2	1.44 ⁱ	4.1	5 ^j
Frequency (GHz)	128.8-143.5	130-164 ^c	137.5-145	134-141	140-160	139-157	130-136	136-147	131-141
EIRP (dBm)	23.5	58-66 ^d	37.5	26	26	8.3	38.5	32	15
P _{dc} /element (mW)	77	265	170	255	275 ^f	150	760	231	814 ^f
Antenna eff. (%)	45*	not specified	81*	45*	not specified	not specified	not specified	58*	68 ^{*,k} (tx) 59 ^{*,k} (rx)
Scan range (deg)	±40 (H-plane)	±20*	±60, ±38	±40	±40	N/A	±5	±30	N/A
Data Capable	NO	YES ^{b,c}	YES	YES	YES	YES	YES	YES	NO

* = simulated, ^a = Intra-chip | Inter-chip, ^b = full system not tested, ^c = does not include antenna, ^d = extrapolated for 256 elements, ^e = wafer-scaling,

^f = antenna driver, ^g = series-fed patch, ^h = 1-d scaling, ⁱ = does not include InP

PA, ^j = includes Rx, ^k = Si resistivity not reported

BIBLIOGRAPHY

- [1] A Fikes et al. “The Caltech space solar power project: Design progress and future direction”. In: *Proc. IEEE WiSEE Space Sol. Power Workshop*. Mar. 2022, pp. 1–5.
- [2] Isaac Asimov. *Reason*. Street & Smith, Apr. 1941.
- [3] Behrooz Abiri et al. *A Lightweight Space-based Solar Power Generation and Transmission Satellite*. 2022. eprint: 2206.08373 (astro-ph.IM).
- [4] Paul Jaffe and James McSpadden. “Energy Conversion and Transmission Modules for Space Solar Power”. In: *Proceedings of the IEEE* 101.6 (2013), pp. 1424–1437. DOI: 10.1109/JPROC.2013.2252591.
- [5] B. Strassner and Kai Chang. “Highly efficient C-band circularly polarized rectifying antenna array for wireless microwave power transmission”. In: *IEEE Transactions on Antennas and Propagation* 51.6 (2003), pp. 1347–1356. DOI: 10.1109/TAP.2003.812252.
- [6] Michael Aaron Marshall. “Dynamics of Ultralight Flexible Spacecraft during Slew Maneuvers”. PhD thesis. California Institute of Technology, 2022.
- [7] Christopher T. Rodenbeck et al. “Microwave and Millimeter Wave Power Beaming”. In: *IEEE Journal of Microwaves* 1.1 (2021), pp. 229–259. DOI: 10.1109/JMW.2020.3033992.
- [8] William C Brown. “Experimental airborne microwave supported platform”. In: *Raytheon Co Burlington, MA Microwave and Power Tube Div* (1965).
- [9] Peter E. Glaser. “Power from the Sun: Its Future”. In: *Science* 162.3856 (1968), pp. 857–861. DOI: 10.1126/science.162.3856.857. eprint: <https://www.science.org/doi/pdf/10.1126/science.162.3856.857>. URL: <https://www.science.org/doi/abs/10.1126/science.162.3856.857>.
- [10] William C. Brown. “Satellite power stations: a new source of energy?” In: *IEEE Spectrum* 10.3 (1973), pp. 38–47. DOI: 10.1109/MSPEC.1973.5216798.
- [11] W.C. Brown. “The technology and application of free-space power transmission by microwave beam”. In: *Proceedings of the IEEE* 62.1 (1974), pp. 11–25. DOI: 10.1109/PROC.1974.9380.
- [12] RM Dickinson and William C Brown. *Radiated microwave power transmission system efficiency measurements*. Tech. rep. Jet Propulsion Laboratories, 1975.

- [13] RMDickinson. *Evaluation of a microwave high-power reception-conversion array for wireless power transmission*. Tech. rep. Jet Propulsion Laboratories, 1975.
- [14] DOE/NASA. *DOE/NASA Satellite Power System Concept Development and Evaluation Program, Satellite Power System (SPS) FY79 Program Summary*. DOE/ER-0037. Washington, DC, USA, 1980.
- [15] James McSpadden et al. “A receiving rectifying antenna for the International Space Year-Microwave Energy Transmission in Space (ISY-METS) rocket experiment”. In: *Intersociety Energy Conversion Engineering Conference*. 1994, p. 4134.
- [16] John C Mankins. “A Technical Overview of the “Suntower” Solar Power Satellite Concept”. In: *Acta Astronautica* 50.6 (2002), pp. 369–377. ISSN: 0094-5765. DOI: [https://doi.org/10.1016/S0094-5765\(01\)00167-9](https://doi.org/10.1016/S0094-5765(01)00167-9). URL: <https://www.sciencedirect.com/science/article/pii/S0094576501001679>.
- [17] W. Seboldt et al. “European Sail Tower SPS concept”. In: *Acta Astronautica* 48.5 (2001), pp. 785–792. ISSN: 0094-5765. DOI: [https://doi.org/10.1016/S0094-5765\(01\)00046-7](https://doi.org/10.1016/S0094-5765(01)00046-7). URL: <https://www.sciencedirect.com/science/article/pii/S0094576501000467>.
- [18] Masahiro Mori, Hideshi Kagawa, and Yuka Saito. “Summary of studies on space solar power systems of Japan Aerospace Exploration Agency (JAXA)”. In: *Acta Astronautica* 59.1 (2006). Space for Inspiration of Humankind, Selected Proceedings of the 56th International Astronautical Federation Congress, Fukuoka, Japan, 17-21 October 2005, pp. 132–138. ISSN: 0094-5765. DOI: <https://doi.org/10.1016/j.actaastro.2006.02.033>. URL: <https://www.sciencedirect.com/science/article/pii/S0094576506001032>.
- [19] Susumu Sasaki et al. “A new concept of solar power satellite: Tethered-SPS”. In: *Acta Astronautica* 60.3 (2007), pp. 153–165. ISSN: 0094-5765. DOI: <https://doi.org/10.1016/j.actaastro.2006.07.010>. URL: <https://www.sciencedirect.com/science/article/pii/S0094576506002815>.
- [20] Ian Cash. “CASSIOPeiA – A new paradigm for space solar power”. In: *Acta Astronautica* 159 (2019), pp. 170–178. ISSN: 0094-5765. DOI: <https://doi.org/10.1016/j.actaastro.2019.03.063>. URL: <https://www.sciencedirect.com/science/article/pii/S0094576518320708>.
- [21] Yang Yang et al. “A novel design project for space solar power station (SSPS-OMEGA)”. In: *Acta Astronautica* 121 (2016), pp. 51–58. ISSN: 0094-5765. DOI: <https://doi.org/10.1016/j.actaastro.2015.12.029>. URL: <https://www.sciencedirect.com/science/article/pii/S0094576515300680>.

- [22] Xun Li et al. “A new concept of space solar power satellite”. In: *Acta Astronautica* 136 (2017), pp. 182–189. ISSN: 0094-5765. DOI: <https://doi.org/10.1016/j.actaastro.2017.03.017>. URL: <https://www.sciencedirect.com/science/article/pii/S0094576516311092>.
- [23] Oren S. Mizrahi et al. “Space solar power generation: A viable system proposal and technoeconomic analysis”. In: *Joule* 9.6 (2025), p. 101928. ISSN: 2542-4351. DOI: <https://doi.org/10.1016/j.joule.2025.101928>. URL: <https://www.sciencedirect.com/science/article/pii/S2542435125001096>.
- [24] Mohammed Reza M Hashemi et al. “A flexible phased array system with low areal mass density”. In: *Nature electronics* 2.5 (2019), pp. 195–205.
- [25] Michael D. Kelzenberg et al. “Ultralight Energy Converter Tile for the Space Solar Power Initiative”. In: *2018 IEEE 7th World Conference on Photovoltaic Energy Conversion (WCPEC) (A Joint Conference of 45th IEEE PVSC, 28th PVSEC & 34th EU PVSEC)*. 2018, pp. 3357–3359. DOI: 10.1109/PVSC.2018.8547403.
- [26] M. Gal-Katziri et al. “Scalable, Deployable, Flexible Phased Array Sheets”. In: *2020 IEEE/MTT-S International Microwave Symposium (IMS)*. 2020, pp. 1085–1088. DOI: 10.1109/IMS30576.2020.9224066.
- [27] Matan Gal-Katziri, Austin Fikes, and Ali Hajimiri. “Flexible active antenna arrays”. In: *npj Flexible Electronics* 6.1 (2022), p. 85.
- [28] Austin C Fikes et al. “Frontiers in Flexible and Shape-Changing Arrays”. In: *IEEE Journal of Microwaves* 3.1 (2022), pp. 349–367.
- [29] Austin Fikes, Oren S. Mizrahi, and Ali Hajimiri. “A Framework for Array Shape Reconstruction Through Mutual Coupling”. In: *IEEE Transactions on Microwave Theory and Techniques* 69.10 (2021), pp. 4422–4436. DOI: 10.1109/TMTT.2021.3097729.
- [30] Oren S. Mizrahi, Austin Fikes, and Ali Hajimiri. “Flexible Phased Array Shape Reconstruction”. In: *2021 IEEE MTT-S International Microwave Symposium (IMS)*. 2021, pp. 31–33. DOI: 10.1109/IMS19712.2021.9574941.
- [31] Florian Bohn, Behrooz Abiri, and Ali Hajimiri. “Fully integrated CMOS X-band power amplifier quad with current reuse and dynamic digital feedback (DDF) capabilities”. In: *2017 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*. 2017, pp. 208–211. DOI: 10.1109/RFIC.2017.7969054.
- [32] Austin Fikes et al. “Fully Collapsible Lightweight Dipole Antennas”. In: *2021 IEEE International Symposium on Antennas and Propagation and USNC-URSI Radio Science Meeting (APS/URSI)*. 2021, pp. 545–546. DOI: 10.1109/APS/URSI47566.2021.9704302.

- [33] Oren S. Mizrahi et al. “Popup Arrays for Large Space-Borne Apertures”. In: *IEEE Transactions on Antennas and Propagation* 72.8 (2024), pp. 6387–6400. DOI: 10.1109/TAP.2024.3414593.
- [34] Ali Hajimiri et al. “Dynamic Focusing of Large Arrays for Wireless Power Transfer and Beyond”. In: *IEEE Journal of Solid-State Circuits* 56.7 (2021), pp. 2077–2101. DOI: 10.1109/JSSC.2020.3036895.
- [35] Youngsub Kim et al. “X-band printed rectenna design and experiment for wireless power transfer”. In: *2014 IEEE Wireless Power Transfer Conference*. 2014, pp. 292–295. DOI: 10.1109/WPT.2014.6839566.
- [36] Momentus Space. “Momentus’ Vigoride Spacecraft Integrated with @SpaceX Falcon 9 Launch Vehicle”. In: *twitter.com* (2023). URL: <https://twitter.com/momentusspace/status/1598344495904755713>.
- [37] A. Fikes et al. “The Caltech Space Solar Power Demonstration One Mission”. In: *2022 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. 2022, pp. 18–22. DOI: 10.1109/WiSEE49342.2022.9926883.
- [38] Gregory R. Allen et al. “2015 Compendium of Recent Test Results of Single Event Effects Conducted by the Jet Propulsion Laboratory’s Radiation Effects Group”. In: *2015 IEEE Radiation Effects Data Workshop (REDW)*. 2015, pp. 1–14. DOI: 10.1109/REDW.2015.7336703.
- [39] Gregory R. Allen et al. “2017 Compendium of Recent Test Results of Single Event Effects Conducted by the Jet Propulsion Laboratory’s Radiation Effects Group”. In: *2017 IEEE Radiation Effects Data Workshop (REDW)*. 2017, pp. 1–9. DOI: 10.1109/NSREC.2017.8115429.
- [40] USSF. *Space-Track.org*. United States Space Force. Arlington County, VA, USA, 2023.
- [41] Alex Ayling et al. “Wireless power transfer in space using flexible, lightweight, coherent arrays”. In: *Acta Astronautica* 224 (2024), pp. 226–243. ISSN: 0094-5765. DOI: <https://doi.org/10.1016/j.actaastro.2024.08.006>. URL: <https://www.sciencedirect.com/science/article/pii/S009457652400434X>.
- [42] Akihito Ebata, Takashi Shimizu, and Yoshinori Kogami. “Frequency dependence measurement of complex permittivity for C-, A- and R-plane sapphire substrates from 30 to 50GHz band”. In: *2014 Asia-Pacific Microwave Conference*. 2014, pp. 516–518.
- [43] Brent Boshart. *SkyTrak*. Heavenscape. Atwood, Ontario, Canada, 2023.
- [44] DF Gluck, V Baturkin, and D Gilmore. “Mountings and interfaces”. In: *Spacecraft thermal control handbook 1* (2002).

- [45] Wriddhi Chakraborty et al. “Characterization and modeling of 22 nm FD-SOI cryogenic RF CMOS”. In: *IEEE Journal on Exploratory Solid-State Computational Devices and Circuits* 7.2 (2021), pp. 184–192.
- [46] Dajiang Yang, Yuanli Ding, and Samny Huang. “A 65-nm High-Frequency Low-Noise CMOS-Based RF SoC Technology”. In: *IEEE Transactions on Electron Devices* 57.1 (2010), pp. 328–335. DOI: 10.1109/TED.2009.2034994.
- [47] I. Aoki et al. “A fully-integrated 1.8-V, 2.8-W, 1.9-GHz, CMOS power amplifier”. In: *IEEE Radio Frequency Integrated Circuits (RFIC) Symposium, 2003*. 2003, pp. 199–202. DOI: 10.1109/RFIC.2003.1213925.
- [48] Steve C Cripps et al. *RF power amplifiers for wireless communications*. Vol. 250. Artech house Norwood, MA, 2006.
- [49] Rouzbeh Kananizadeh and Omeed Momeni. “High-Power and High-Efficiency Millimeter-Wave Harmonic Oscillator Design, Exploiting Harmonic Positive Feedback in CMOS”. In: *IEEE Transactions on Microwave Theory and Techniques* 65.10 (2017), pp. 3922–3936. DOI: 10.1109/TMTT.2017.2690291.
- [50] Hua Wang et al. “Power Amplifiers Performance Survey 2000-Present”. In: (2025). URL: <https://ideas.ethz.ch/surveys/pa-survey.html>.
- [51] N. Mohan. *Power Electronics: A First Course: A First Course*. Wiley, 2011. ISBN: 9781118214343. URL: <https://books.google.com/books?id=rfsbAAAAQBAJ>.
- [52] Bidyadhar Subudhi and Raseswari Pradhan. “A Comparative Study on Maximum Power Point Tracking Techniques for Photovoltaic Power Systems”. In: *IEEE Transactions on Sustainable Energy* 4.1 (2013), pp. 89–98. DOI: 10.1109/TSTE.2012.2202294.
- [53] Coilcraft. “Shielded Power Inductors – XFL7015”. In: *XFL7015 Datasheet* (2021). URL: <https://www.coilcraft.com/getmedia/ccf09628-6e8c-462a-9dc9-fa4346e7cf0a/xfl7015.pdf>.
- [54] Yu-Lin Yang et al. “Using pulse width and waveform modulation to enhance power conversion efficiency under constraint of low input power”. In: *2012 Asia Pacific Microwave Conference Proceedings*. 2012, pp. 400–402. DOI: 10.1109/APMC.2012.6421611.
- [55] Joseph Goodman. *Introduction to Fourier Optics*. Englewood, CO, 2005.
- [56] Matan Gal-Katziri et al. “Optically Synchronized Phased Arrays in CMOS”. In: *IEEE Journal of Solid-State Circuits* 57.6 (2022), pp. 1578–1593. DOI: 10.1109/JSSC.2021.3136787.

- [57] Jason M. Merlo, Serge R. Mghabghab, and Jeffrey A. Nanzer. “Wireless Picosecond Time Synchronization for Distributed Antenna Arrays”. In: *IEEE Transactions on Microwave Theory and Techniques* 71.4 (2023), pp. 1720–1731. DOI: 10.1109/TMTT.2022.3227878.
- [58] Zhi Ning Chen et al. “Microwave Metalens Antennas”. In: *Proceedings of the IEEE* 111.8 (2023), pp. 978–1010. DOI: 10.1109/JPROC.2023.3287599.
- [59] Alex Ayling et al. “Wireless power transfer in space using flexible, lightweight, coherent arrays”. In: *Acta Astronautica* 224 (2024), pp. 226–243. ISSN: 0094-5765. DOI: <https://doi.org/10.1016/j.actaastro.2024.08.006>. URL: <https://www.sciencedirect.com/science/article/pii/S009457652400434X>.
- [60] Jesse Brunet, Alex Ayling, and Ali Hajimiri. “Transmitarrays for Wireless Power Transfer on Earth and in Space”. In: *IEEE Journal of Microwaves* 4.4 (2024), pp. 934–945. DOI: 10.1109/JMW.2024.3459859.
- [61] Jonathan Yun Lau and Sean Victor Hum. “Analysis and Characterization of a Multipole Reconfigurable Transmitarray Element”. In: *IEEE Transactions on Antennas and Propagation* 59.1 (2011), pp. 70–79. DOI: 10.1109/TAP.2010.2090456.
- [62] Jonathan Yun Lau and Sean Victor Hum. “A Planar Reconfigurable Aperture With Lens and Reflectarray Modes of Operation”. In: *IEEE Transactions on Microwave Theory and Techniques* 58.12 (2010), pp. 3547–3555. DOI: 10.1109/TMTT.2010.2086373.
- [63] Akshay Visweswaran et al. “A 28-nm-CMOS Based 145-GHz FMCW Radar: System, Circuits, and Characterization”. In: *IEEE Journal of Solid-State Circuits* 56.7 (July 2021), pp. 1975–1993.
- [64] Mohamed Elkhoully et al. “Fully Integrated 2D Scalable TX/RX Chipset for D-Band Phased-Array-on-Glass Modules”. In: *2022 IEEE International Solid-State Circuits Conference (ISSCC)*. San Francisco, CA, USA, Feb. 2022, pp. 76–78.
- [65] Amr Ahmed et al. “140-GHz 2-D Scalable On-Grid 8×8 -Element Transmit–Receive Phased Arrays With Up/Down Converters Demonstrating a 5.2-m Link at 16 Gbps”. In: *IEEE Transactions on Microwave Theory and Techniques* (2024), pp. 1–17. DOI: 10.1109/TMTT.2023.3346407.
- [66] Jicong Zhang et al. “24.2 A Scalable 134-to-141GHz 16-Element CMOS 2D $\lambda/2$ -Spaced Phased Array”. In: *2024 IEEE International Solid-State Circuits Conference (ISSCC)*. Vol. 67. 2024, pp. 414–416. DOI: 10.1109/ISSCC49657.2024.10454376.
- [67] Amr Ahmed et al. “A 140 GHz Scalable On-Grid 8×8 -Element Transmit–Receive Phased-Array with Up/Down Converters and 64QAM/24 Gbps Data Rates”. In: *2023 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*. 2023, pp. 93–96. DOI: 10.1109/RFIC54547.2023.10186155.

- [68] David del Rio et al. “A D-Band 16-Element Phased-Array Transceiver in 55-nm BiCMOS”. In: *IEEE Transactions on Microwave Theory and Techniques* 71.2 (Feb. 2023), pp. 854–869.
- [69] Abdelaziz Hamani et al. “A D-Band Multichannel TX System-in-Package Achieving 84.48 Gb/s With 64-QAM Based on 45-nm CMOS and Low-Cost PCB Technology”. In: *IEEE Transactions on Microwave Theory and Techniques* 70.7 (July 2022), pp. 3385–3395.
- [70] A. A. Faid et al. “A Fully Packaged 135-GHz Multiuser MIMO Transmitter Array Tile for Wireless Communications”. In: *IEEE Transactions on Microwave Theory and Techniques* 70.7 (July 2022), pp. 3396–3405.
- [71] Siwei Li, Zhe Zhang, and Gabriel M. Rebeiz. “An Eight-Element 136–147 GHz Wafer-Scale Phased-Array Transmitter With 32 dBm Peak EIRP and >16 Gbps 16QAM and 64QAM Operation”. In: *IEEE Journal of Solid-State Circuits* 57.6 (2022), pp. 1635–1648. DOI: 10.1109/JSSC.2022.3148385.
- [72] Siwei Li et al. “An Eight-Element 140-GHz Wafer-Scale IF Beamforming Phased-Array Receiver With 64-QAM Operation in CMOS RFSOI”. In: *IEEE Journal of Solid-State Circuits* 57.2 (2022), pp. 385–399. DOI: 10.1109/JSSC.2021.3102876.
- [73] Bodhisatwa Sadhu, Xiaoxiong Gu, and Alberto Valdes-Garcia. “The More (Antennas), the Merrier: A Survey of Silicon-Based mm-Wave Phased Arrays Using Multi-IC Scaling”. In: *IEEE Microwave Magazine* 20.12 (2019), pp. 32–50. DOI: 10.1109/MMM.2019.2941632.
- [74] Steffen Hansen et al. “A SiGe-Chip-Based D-Band FMCW-Radar Sensor With 53-GHz Tuning Range for High Resolution Measurements in Industrial Applications”. In: *IEEE Transactions on Microwave Theory and Techniques* 70.1 (2022), pp. 719–731. DOI: 10.1109/TMTT.2021.3121746.
- [75] Xiaofan Jia et al. “Antenna With Embedded Die in Glass Interposer for 6G Wireless Applications”. In: *IEEE Transactions on Components, Packaging and Manufacturing Technology* 13.2 (2023), pp. 219–229. DOI: 10.1109/TCPMT.2023.3251725.
- [76] Meng Wei et al. “A D-Band Packaged CMOS Integrated Transmitter for MU-MIMO Applications”. In: *2022 IEEE Asian Solid-State Circuits Conference (A-SSCC)*. 2022, pp. 6–8. DOI: 10.1109/A-SSCC56115.2022.9980538.
- [77] Sensen Li et al. “A fully packaged D-band MIMO transmitter using high-density flip-chip interconnects on LCP substrate”. In: *2016 IEEE MTT-S International Microwave Symposium (IMS)*. 2016, pp. 1–4. DOI: 10.1109/MWSYM.2016.7540177.
- [78] Wasif Tanveer Khan et al. “Characterization of a low-loss and wide-band (DC to 170 GHz) flip-chip interconnect on an organic substrate”. In: *2014 IEEE MTT-S International Microwave Symposium (IMS2014)*. 2014, pp. 1–4. DOI: 10.1109/MWSYM.2014.6848343.

- [79] Zhibo Cao et al. “D-Band Flip-Chip Packaging with Wafer-Level Cu-pillar Bumps”. In: *2023 IEEE 32nd Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*. 2023, pp. 1–4. doi: 10.1109/EPEPS58208.2023.10314877.
- [80] S. Beer et al. “Design and measurement of matched wire bond and flip chip interconnects for D-band system-in-package applications”. In: *2011 IEEE MTT-S International Microwave Symposium*. 2011, pp. 1–4. doi: 10.1109/MWSYM.2011.5972851.
- [81] S. Beer et al. “D-band grid-array antenna integrated in the lid of a surface-mountable chip-package”. In: *2013 7th European Conference on Antennas and Propagation (EuCAP)*. 2013, pp. 1318–1322.
- [82] D. Rutledge and M. Muha. “Imaging antenna arrays”. In: *IEEE Transactions on Antennas and Propagation* 30.4 (1982), pp. 535–540. doi: 10.1109/TAP.1982.1142856.
- [83] N.G. Alexopoulos, P.B. Katehi, and D.B. Rutledge. “Substrate Optimization for Integrated Circuit Antennas”. In: *IEEE Transactions on Microwave Theory and Techniques* 31.7 (1983), pp. 550–557. doi: 10.1109/TMTT.1983.1131544.
- [84] Martijn de Kok, A. Bart Smolders, and Ulf Johannsen. “A Review of Design and Integration Technologies for D-Band Antennas”. In: *IEEE Open Journal of Antennas and Propagation* 2 (2021), pp. 746–758. doi: 10.1109/OJAP.2021.3089052.
- [85] Aydin Babakhani et al. “A 77-GHz Phased-Array Transceiver With On-Chip Antennas in Silicon: Receiver and Antennas”. In: *IEEE Journal of Solid-State Circuits* 41.12 (2006), pp. 2795–2806. doi: 10.1109/JSSC.2006.884811.
- [86] Lili Chen, Andreia Cathelin, and Ehsan Afshari. “A High-Efficiency High-Power 170-176-GHz Frequency Stabilized Quadrature Radiator”. In: *IEEE Journal of Solid-State Circuits* 59.1 (2024), pp. 243–252. doi: 10.1109/JSSC.2023.3276388.
- [87] Maria Alonso-delPino et al. “Beam Scanning of Silicon Lens Antennas Using Integrated Piezomotors at Submillimeter Wavelengths”. In: *IEEE Transactions on Terahertz Science and Technology* 9.1 (2019), pp. 47–54. doi: 10.1109/TTHZ.2018.2881930.
- [88] Herman Jalli Ng et al. “Multi-Purpose Fully Differential 61- and 122-GHz Radar Transceivers for Scalable MIMO Sensor Platforms”. In: *IEEE Journal of Solid-State Circuits* 52.9 (2017), pp. 2242–2255. doi: 10.1109/JSSC.2017.2704602.

- [89] Herman Jalli Ng et al. “Miniaturized 122 GHz system-on-chip radar sensor with on-chip antennas utilizing a novel antenna design approach”. In: *2016 IEEE MTT-S International Microwave Symposium (IMS)*. 2016, pp. 1–4. DOI: 10.1109/MWSYM.2016.7540353.
- [90] Ruoyu Wang et al. “A micromachined double-dipole antenna for 122 – 140 GHz applications based on a SiGe BiCMOS technology”. In: *2012 IEEE/MTT-S International Microwave Symposium Digest*. 2012, pp. 1–3. DOI: 10.1109/MWSYM.2012.6258421.
- [91] Faisal Ahmed, Muhammad Furqan, and Andreas Stelzer. “120-GHz and 240-GHz Broadband Bow-Tie Antennas in eWLB Package for High Resolution Radar Applications”. In: *2018 48th European Microwave Conference (EuMC)*. 2018, pp. 1109–1112. DOI: 10.23919/EuMC.2018.8541732.
- [92] Stefan Beer et al. “Integrated 122-GHz Antenna on a Flexible Polyimide Substrate With Flip Chip Interconnect”. In: *IEEE Transactions on Antennas and Propagation* 61.4 (2013), pp. 1564–1572. DOI: 10.1109/TAP.2012.2232260.
- [93] Ali A. Farid et al. “A Packaged 135GHz 22nm FD-SOI Transmitter on an LTCC Carrier”. In: *2021 IEEE MTT-S International Microwave Symposium (IMS)*. 2021, pp. 713–716. DOI: 10.1109/IMS19712.2021.9575026.
- [94] Amit Singh et al. “A D-Band Radio-on-Glass Module for Spectrally-Efficient and Low-Cost Wireless Backhaul”. In: *2020 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*. 2020, pp. 99–102. DOI: 10.1109/RFIC49505.2020.9218437.
- [95] P. Sullivan and D. Schaubert. “Analysis of an aperture coupled microstrip antenna”. In: *IEEE Transactions on Antennas and Propagation* 34.8 (1986), pp. 977–984. DOI: 10.1109/TAP.1986.1143929.
- [96] Samuel Naffziger et al. “Pioneering Chiplet Technology and Design for the AMD EPYC™ and Ryzen™ Processor Families : Industrial Product”. In: *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. 2021, pp. 57–70. DOI: 10.1109/ISCA52012.2021.00014.
- [97] Constantine A Balanis. *Antenna theory: analysis and design*. John Wiley & sons, 2016.
- [98] Shoubhik Gupta et al. “Ultra-thin chips for high-performance flexible electronics”. In: *npj Flexible Electronics* 2.1 (2018), p. 8.
- [99] Bernd Ebersberger and Charles Lee. “Cu pillar bumps as a lead-free drop-in replacement for solder-bumped, flip-chip interconnects”. In: *2008 58th Electronic Components and Technology Conference*. 2008, pp. 59–66. DOI: 10.1109/ECTC.2008.4549951.

- [100] Liu Chen et al. “The effects of underfill and its material models on thermomechanical behaviors of a flip chip package”. In: *IEEE Transactions on Advanced Packaging* 24.1 (2001), pp. 17–24. DOI: 10.1109/6040.909620.
- [101] Md Hedayatullah Maktoomi et al. “A Sub-Terahertz Wideband Stacked-Patch Antenna on a Flexible Printed Circuit for 6G Applications”. In: *IEEE Transactions on Antennas and Propagation* 70.11 (2022), pp. 10047–10061. DOI: 10.1109/TAP.2022.3185497.
- [102] Akanksha Bhutani et al. “122 GHz aperture-coupled stacked patch microstrip antenna in LTCC technology”. In: *2016 10th European Conference on Antennas and Propagation (EuCAP)*. 2016, pp. 1–5. DOI: 10.1109/EuCAP.2016.7481147.
- [103] Ziad Iskandar et al. “A 270-to-300 GHz LO-Path Phase Shifting Architecture for Sub-mm-wave Phased Arrays”. In: *2018 48th European Microwave Conference (EuMC)*. 2018, pp. 804–806. DOI: 10.23919/EuMC.2018.8541650.
- [104] Nadav Mazor and Eran Socher. “Analysis and design of an X-band-to-W-band CMOS active multiplier with improved harmonic rejection”. In: *IEEE Transactions on Microwave Theory and Techniques* 61.5 (2013), pp. 1924–1933. DOI: 10.1109/TMTT.2013.2252914.
- [105] Alan V. Oppenheim, Alan S. Willsky, and S. Hamid Nawab. *Signals & systems (2nd ed.)* USA: Prentice-Hall, Inc., 1996. ISBN: 0138147574.
- [106] Vincent Lammert, Michael Leyrer, and Vadim Issakov. “A D-Band Transceiver with On-Chip Multi-Port Radiators and Leakage Cancellation in 45nm SOI”. In: *2022 IEEE BiCMOS and Compound Semiconductor Integrated Circuits and Technology Symposium (BCICTS)*. 2022, pp. 57–61. DOI: 10.1109/BCICTS53451.2022.10051760.
- [107] Jeffrey A. Nanzer et al. “Distributed Phased Arrays: Challenges and Recent Advances”. In: *IEEE Transactions on Microwave Theory and Techniques* 69.11 (2021), pp. 4893–4907. DOI: 10.1109/TMTT.2021.3092401.
- [108] Jian-Ming Jin. *Theory and computation of electromagnetic fields*. John Wiley & Sons, 2015.
- [109] Roy D Yates and David J Goodman. *Probability and Stochastic Processes: A Friendly Introduction for Electrical and Computer Engineers, International Adaptation*. John Wiley & Sons, 2025.
- [110] Jongsoo Lee et al. “A Tri-Band Dual-Concurrent Wi-Fi 802.11be Transceiver Achieving 46 dB TX/RX EVM Floor at 7.1 GHz for a 4 K-QAM 320 MHz Signal”. In: *IEEE Journal of Solid-State Circuits* 59.12 (2024), pp. 3966–3979. DOI: 10.1109/JSSC.2024.3441858.

Appendix A

TIMING REQUIREMENTS FOR LARGE ARRAYS

We wish to analyze the effect of uncorrelated phase noise at each radiating element on the transmitted power. To start, the array factor for an $N \times N$ array (fig. A.1) is written below




Figure A.1: 2D-planar $N \times N$ array

$$AF(\theta, \phi) = \sum_{n=1}^N e^{jn(kd_x \sin \theta \cos \phi + \alpha_x)} \times \sum_{m=1}^N e^{jm(kd_y \sin \theta \sin \phi + \alpha_y)}. \quad (\text{A.1})$$

In equation A.1, k is the propagation constant in free space, d_x and d_y are the element spacings in the x and y directions, α_x is the phase gradient in the x -axis, and α_y the phase gradient in the y -axis. Assume the array is steered to broadside, so $\alpha_x = \alpha_y = 0$. Now, assume that a random variable Δ_i corresponding to the phase noise is added at each element, then

$$AF(\theta, \phi) = \sum_{n=1}^N e^{jk d_x n (\sin \theta \cos \phi + \Delta_n)} \times \sum_{m=1}^N e^{jk d_y m (\sin \theta \sin \phi + \Delta_m)}. \quad (\text{A.2})$$

Or, more simply at broadside,

$$AF(0,0) = 1 + \sum_{n=1}^{N^2-1} e^{j\Delta_n}. \quad (\text{A.3})$$

Note that equation A.3 assumes the first element is a phase reference for the array without loss of generality, which evaluates to one and is pulled out of the summation. To evaluate the degradation in transmitted power, the coherent gain from [107] is used, defined as

$$G_c = \frac{G_{array}}{G_{max}}, \quad (\text{A.4})$$

where G_{array} is the gain using the array-factor defined in A.2 and A.3, and G_{max} is the ideal beamforming gain of the array with perfect coherence. The ideal beamforming gain for an NxN array is N^2 . The definition of antenna gain at broadside is

$$G_{array} = \frac{U_{array}(0,0)}{\frac{1}{4\pi} \int_0^{2\pi} \int_0^\pi U_{array}(\theta, \phi) \sin \theta d\theta d\phi}, \quad (\text{A.5})$$

where $U(\theta, \phi)$ is the radiant intensity with units of watt/steradian. The denominator of A.5 is proportional to the total radiated power, which will become useful later. The radiant intensity is related to the Poynting vector by

$$U(\theta, \phi) = \hat{r} \cdot r^2 \vec{S}(\theta, \phi, r) = \hat{r} \cdot r^2 \frac{1}{2} \vec{E} \times \vec{H}^*, \quad (\text{A.6})$$

where r^2 is the radial component in spherical coordinates, $\vec{S}$ the Poynting vector, $\vec{E}$ the electric field and $\vec{H}$ the magnetic field. To simplify the math, an isotropic radiation pattern is assumed for each element in the array. In that case, the $\vec{E}$ and $\vec{H}$ fields for an isolated element can be written as

$$\vec{E}(\theta, \phi, r) = \hat{\theta} \frac{jk\eta}{4\pi r} e^{-jkr}, \quad (\text{A.7})$$

$$\vec{H}(\theta, \phi, r) = \hat{\phi} \frac{jk}{4\pi r} e^{-jkr}, \quad (\text{A.8})$$

where $\eta = \sqrt{\frac{\mu_0}{\epsilon_0}}$ is the wave-impedance in free-space. Using equations A.7 and A.8, the power radiated by an isolated element in free space is

$$U_{el}(\theta, \phi) = r^2 \frac{\eta}{2} \left| \frac{k}{4\pi r} \right|^2 = \frac{\eta}{2} \left| \frac{k}{4\pi} \right|^2, \quad (\text{A.9})$$

$$P_{el} = \int_0^{2\pi} \int_0^\pi U_{el}(\theta, \phi) \sin \theta \, d\theta d\phi = \int_0^{2\pi} \int_0^\pi \frac{\eta}{2} \left| \frac{k}{4\pi} \right|^2 \sin \theta \, d\theta d\phi, \quad (\text{A.10})$$

which evaluates to

$$P_{el} = \frac{\eta}{2} \frac{k^2}{4\pi}. \quad (\text{A.11})$$

By energy conservation, the total power radiated by the $N \times N$ array should be the element power in A.11 multiplied by the number of elements,

$$P_{array} = P_{el} N^2 = \frac{\eta}{2} \frac{k^2}{4\pi} N^2. \quad (\text{A.12})$$

Finally, the array pattern is simply the multiplication of the array factor and the element pattern [108].

$$\vec{E}_{array} = \vec{E}(\theta, \phi) AF(\theta, \phi) \quad (\text{A.13})$$

$$\vec{H}_{array} = \vec{H}(\theta, \phi) AF(\theta, \phi). \quad (\text{A.14})$$

Thus, the radiant intensity of the array is

$$U_{array}(\theta, \phi) = \hat{r} \cdot r^2 \frac{1}{2} \vec{E}_{array} \times \vec{H}_{array}^* = \frac{\eta}{2} \left(\frac{k}{4\pi} \right)^2 \times |AF(\theta, \phi)|^2 \quad (\text{A.15})$$

Using equations A.12 and A.15 and plugging into A.5, the gain of the array is

$$G_{array} = \frac{\frac{\eta}{2} \left(\frac{k}{4\pi} \right)^2 |AF(0, 0)|^2}{\frac{\eta}{2} \left(\frac{k^2}{4\pi} \right)^2 N^2} = \frac{|AF(0, 0)|^2}{N^2}. \quad (\text{A.16})$$

An ideal, perfectly coherent array with isotropic elements steered to broadside has a maximum gain of N^2 , so

$$G_C = \frac{|AF(0,0)|^2}{N^4}. \quad (\text{A.17})$$

Thus, after a good deal of math, the coherent gain of the array can be analyzed by evaluating the relatively simple expression in A.17. We have the expression for $AF(0,0)$ in A.3. Let's expand this expression out using Euler's formula, then

$$\frac{|AF(0,0)|^2}{N^4} = \left[\frac{1 + \sum_{n=1}^{N^2-1} \cos(\Delta_n)}{N^2} \right]^2 + \left[\frac{\sum_{n=1}^{N^2-1} \sin(\Delta_n)}{N^2} \right]^2. \quad (\text{A.18})$$

Looking at equation A.18, we know that the random variables Δ_n are independent and identically distributed (i.i.d). Let $X_n = \cos(\Delta_n)$ and $Y_n = \sin(\Delta_n)$, then

$$\frac{|AF(0,0)|^2}{N^4} = \left[\frac{1 + \sum_{n=1}^{N^2-1} X_n}{N^2} \right]^2 + \left[\frac{\sum_{n=1}^{N^2-1} Y_n}{N^2} \right]^2. \quad (\text{A.19})$$

We know from probability and statistics that $\frac{X_1+X_2+\dots+X_n}{n} = E[X_n]$ by the Law of Large Numbers [109]. Equation A.19 can be rearranged to take advantage of this

$$\frac{|AF(0,0)|^2}{N^4} = \left[\frac{1}{N^2} + \left(1 - \frac{1}{N^2}\right) \sum_{n=1}^{N^2-1} \frac{X_n}{N^2-1} \right]^2 + \left[\left(1 - \frac{1}{N^2}\right) \sum_{n=1}^{N^2-1} \frac{Y_n}{N^2-1} \right]^2. \quad (\text{A.20})$$

Using the Law of Large Numbers, the limit as $N \rightarrow \infty$ in equation A.20 is simply

$$\lim_{N \rightarrow \infty} \frac{|AF(0,0)|^2}{N^4} = \left[E[X_n] \right]^2 + \left[E[Y_n] \right]^2. \quad (\text{A.21})$$

So far we have made no assumptions about the distribution of the random phase noise. To start, we will assume the distribution is uniform random. The results of a uniform random variable can be used to approximate the response when the distribution is Gaussian, which is a more accurate representation of the phase noise in RF systems. Let the probability density function of Δ_i be defined as

$$f_{\Delta}(\delta) = \begin{cases} \frac{1}{2B} & \text{if } |\delta| < B, \\ 0 & \text{if } |\delta| \geq B. \end{cases} \quad (\text{A.22})$$




Figure A.2: Uniform continuous probability density function for Δ_i

The PDF in equation A.22 is sampled from a range of $[-B, B]$ and is centered at zero. We know that $E[g(X)] = \int g(x)f_X(x)dx$ [109], so

$$E[X_n] = \int_{-B}^B \cos \delta \frac{1}{2B} d\delta \quad (\text{A.23})$$

$$E[Y_n] = \int_{-B}^B \sin \delta \frac{1}{2B} d\delta. \quad (\text{A.24})$$

Equation A.23 is easily solved, and A.24 evaluates to zero since $\sin(x)$ is an odd function. The final result is

$$G_c = \left[\frac{\sin B}{B} \right]^2 = \text{sinc}^2(B). \quad (\text{A.25})$$

Thus, the coherent gain is related to the phase noise variance by a sinc function. The assumptions used in deriving A.25 should be restated for clarity:

- The number of elements approaches ∞ .
- The phase error at each element is uncorrelated with respect to the others.
- The phase error takes a uniform random distribution, with expected value of zero and a range of $[-B, B]$.

Equation A.25 is plotted along with a statistical MATLAB simulation of a 10,000 element array with 10,000 trials in Fig. A.3. The match between simulation and theory is excellent.




Figure A.3: Coherent gain with uniform phase error distribution.

Coherent Gain with Gaussian Phase Error

As mentioned previously, a $\text{Gauss}(0, \sigma^2)$ distribution for Δ_n is a more accurate representation phase noise. Technically, this could be captured by solving A.21 with a Gaussian probability density function rather than a uniform one. In practice, the math is quite messy and most likely cannot be solved exactly. As an alternative, we note that the Central Limit Theorem states that, regardless of a random variable's distribution, as the summation of i.i.d. random variables grows large it will approach a Gaussian distribution. In particular [109],

$$\bar{X} = \frac{X_1 + \cdots + X_n}{n} \quad (\text{A.26})$$

$$(\bar{X} - \mu) \rightarrow \text{Gauss}\left(0, \frac{\sigma^2}{n}\right). \quad (\text{A.27})$$

Thus, the Central Limit Theorem tells us that if we have two random variables with different distributions but their mean and variance are the same, their sum will approach the same gaussian random variable. So, we can get the coherent gain with gaussian random variables by simply modifying A.25. Let σ^2 be the variance of the gaussian random variable. The variance of a uniform $[-B, B]$ random variable is $\frac{B^2}{3}$ and both have a mean of zero. So, by equating the variance of the two we get

$$G_c \approx \text{sinc}^2(\sqrt{3}\sigma). \quad (\text{A.28})$$

Plots of the coherent gain with a Gaussian distribution (MATLAB simulation and theory) are shown in Fig. A.4. The match between simulation and theory is close for small σ , but deviates for larger values. However, the important regime is $G_c < 90\%$ in practical systems, where the match is good. A table of important breakpoints in the coherent gain is given below.




Figure A.4: Coherent gain with Gaussian phase error.

Table A.1: G_c vs. RMS phase jitter with Gaussian random variables.

G_c	σ
99.9%	1.81°
99%	5.74°
90%	18.5°

To put these results in perspective, typical WiFi-7 transceivers, which must support very high-order modulation schemes such as 4K-QAM, typically target an RMS jitter of 50fs, or about 0.18° at 10GHz [110]. Thus, $G_c > 99.9\%$ is readily achievable with existing CMOS synthesizers with extra margin that can be used to reduce power consumption.

Appendix B

PHASE-SHIFTER QUANTIZATION LOSS FOR LARGE ARRAYS

A related problem to that analyzed in Appendix A is the phase-shifter accuracy necessary for beamforming. Specifically, we assume there is some pseudo-random and quasi-static phase offset at each element in the array. Physically, this would occur due to mismatches in the transmission line lengths, mismatches in the delays of the active components in the signal-path, temperature dependent phase shifts etc. These are random in terms of manufacturing variation, but the key point here is that once the system is assembled they are deterministic and slowly varying, meaning that they can be corrected. This is in contrast with something like phase noise analyzed in Appendix A, which is randomly varying over time and cannot be corrected.

What distribution do these random phase offsets take? As discussed above, these offsets are due to process variation in the manufacturing of things like the transmission-lines, RFICs, optical and wireless distribution systems etc. Thus, they are the sum of the phase offsets due to these individual components, which should be independent with respect to one another. So, we have the sum of many independent variables, which we know should approach a Gaussian distribution by the Central Limit Theorem [109].

However, something interesting happens when we consider that the phase is periodic with 2π radians. Consider a Gaussian distribution where $\sigma \gg 2\pi$. If we look at intervals of $[0, 2\pi]$, $[2\pi, 4\pi]$... $[2N\pi, 2(N+1)\pi]$, all of these will fold into the same $[0, 2\pi]$ interval when we wrap the phase. Since $\sigma^2 \gg 2\pi$, the PDF changes very little during this interval and looks effectively like a uniform distribution. All of these conditional uniform distributions can be added according to the law of conditional probability [109], which will simply yield an approximately uniform distribution over a $[0, 2\pi]$ interval. This can be seen a MATLAB simulation of the PDF in Fig. B.1. For $\sigma = 0.2\pi$, the Gaussian distribution is mostly preserved because little phase wrapping occurs. For $\sigma = 2\pi$, the phase wrapping forces the distribution to approach a uniform one.

Finally, is the assumption $\sigma \gg 2\pi$ reasonable? Considering the size of the array is extremely large in wavelengths (km scale array vs. cm scale wavelength), something

as simple as an optical fiber for reference distribution will be many wavelengths large. A small skew in the fiber index, leading to a change in the propagation velocity, will produce large phase deviations. This, among other effects, are compounded as the array is scaled, resulting in $\sigma \gg 2\pi$.




Figure B.1: Gaussian phase offset wrapping.

Let's return to the problem of phase shifter accuracy. We know the phase offsets will take a uniform $[-180^\circ, 180^\circ]$ distribution. The phase shifter can be controlled in steps of $\frac{360^\circ}{2^k}$, where k is the number of bits in the phase shifter. The phase shifter can correct the offset to within a $\left[-\frac{180^\circ}{2^k}, \frac{180^\circ}{2^k} \right]$ range taking a uniform random distribution. We have already solved this problem, however, in Appendix A. Noting that $B = \frac{180^\circ}{2^k}$ and using equation A.25, then the degradation in coherent gain from the phase-shifter quantization is simply

$$G_c = \text{sinc}^2\left(\frac{\pi}{2^k}\right). \quad (\text{B.1})$$

The table below shows the coherent gain degradation for different phase-shifter resolutions. The results are compared with a MATLAB simulation of a 64x64 element array and the two are in good agreement. With a phase-shifter resolution of 4-bits, the coherent gain is nearly 99%, getting above the threshold with 5-bits. Beyond that point there are diminishing returns. This is important as higher-resolution phase-shifters require more area and power consumption than lower resolution designs. A 5-bit phase-shifter was ultimately chosen for the 22-nm SSPP RFIC as it is the minimum phase-shifter resolution with $G_c > 99\%$.

Table B.1: G_c for different phase-shifter resolutions.

Phase-shifter bits	G_c (sim)	G_c (theory)
1	40.53%	40.53%
2	81.20%	81.06%
3	94.91%	94.96%
4	98.72%	98.72%
5	99.68%	99.68%
6	99.92%	99.92%

Appendix C

MPO TIMING REQUIREMENTS

It is worthwhile to investigate the timing requirements for the MPO switching signals given the previous analysis done for the phase jitter and phase shifter resolution requirements in Appendix A and B. To start, we assume an array whose RF signal is perfectly coherent, but each element is modulated by a pulse function of duty-cycle D with an arbitrary delay skew. In that case, the signal at the receiver will be

$$V_r = A_0 \cos(\omega_{rf}t) [A(t - \tau_1) + \cdots + A(t - \tau_n)], \quad (\text{C.1})$$

where A_0 is an arbitrary amplitude, $A(t)$ is a normalized pulse function of duty-cycle D representing the MPO switching signal, and τ_n is the delay skew at element n . We assume that $A(t)$ is periodic with period T_{sw} , and further more that $T_{sw} = MT_{rf}$. In other words, the rf frequency is an integer multiple of the switching frequency. This is reasonable given both will be derived from the same low-frequency reference on the RFIC. The total received-power is

$$P_r = \frac{1}{T_{sw}} \int_0^{T_{sw}} \left(A_0 \cos(\omega_{rf}t) [A(t - \tau_1) + \cdots + A(t - \tau_n)] \right)^2 dt. \quad (\text{C.2})$$

We can expand out $\cos(\omega_{rf}t)^2$ using the well-known trigonometric identity to yield

$$P_r = \frac{A_0^2}{2T_{sw}} \int_0^{T_{sw}} (1 + \cos(2\omega_{rf}t)) [A(t - \tau_1) + \cdots + A(t - \tau_n)]^2 dt. \quad (\text{C.3})$$

Equation C.3 can be greatly simplified when we note that $A(t)$ can be expanded out with a Fourier Series. In that case, the $\cos(2\omega_{rf}t)$ term multiplies out with all the Fourier components of each $A(t - \tau_i)$ function. We know that cosine functions are orthogonal in frequency when multiplied and integrated over a period, so the only terms that are non-zero occur when $\omega_{rf} = M\omega_{sw}$. Furthermore, we assume that $\omega_{rf} \gg \omega_{sw}$, which will be the case since the switching frequency will be several MHz and the rf frequency 10GHz. Thus, only the high-order Fourier Series coefficient will be non-zero, which can be safely neglected with respect to the DC term in C.3, so the $\cos(2\omega_{rf}t)$ can be eliminated such that

$$P_r \approx \frac{A_0^2}{2T_{sw}} \int_0^{T_{sw}} [A(t - \tau_1) + \dots + A(t - \tau_n)]^2 dt. \quad (C.4)$$

Equation C.4 has a maximum value when $\tau_1 = \tau_2 = \dots = \tau_n = 0$, with

$$P_{max} = \frac{A_0^2}{2} DN^2. \quad (C.5)$$

Now, let's expand $A(t - \tau_n)$ with its Fourier series [105]

$$A(t - \tau) = D + \sum_{n=1}^{\infty} a_n \cos(n\omega_{sw}(t - \tau)), \quad (C.6)$$

$$a_n = \frac{2}{n\pi} \sin(n\pi D). \quad (C.7)$$

Using equations C.5, C.6, and C.7, we can rewrite C.4 as

$$\frac{P_r}{P_{max}} = \frac{1}{T_{sw} DN^2} \int_0^{T_{sw}} \left[DN + \sum_{m=1}^{\infty} \sum_{n=1}^{\infty} a_m \cos(m\omega_{rf}(t - \tau_n)) \right]^2 dt. \quad (C.8)$$

The inner summation of C.8 are cosines with the same frequency, which will result in an equivalent cosine with the same frequency and some amplitude/phase. We can condense C.8 such that

$$\frac{P_r}{P_{max}} = \frac{1}{T_{sw} DN^2} \int_0^{T_{sw}} \left[DN + \sum_{m=1}^{\infty} B_m \cos(m\omega_{rf}t + \phi_m) \right]^2 dt. \quad (C.9)$$

We will elaborate on the definition of B_m and ϕ_m in a moment. Before that, however, we can again simplify C.9 using the orthogonality of cosines. When the square is expanded out, only the DC terms will produce a non-zero integral since all the harmonics will integrate to zero. Thus, only the DC and $[B_m \cos(m\omega_{rf}t + \phi_m)]^2$ terms (which has a DC term when expanded using a trigonometric identity) will be non-zero. Thus, C.9 simplifies to

$$\frac{P_r}{P_{max}} = \frac{1}{T_{sw} D} \int_0^{T_{sw}} \left[D^2 + \sum_{m=1}^{\infty} \frac{B_m^2}{2N^2} \right] dt. \quad (C.10)$$

Note in C.10 that ϕ_m cancels out for the non-zero terms. As for B_m , we can analyze it by rewriting using the phasor form of the cosine-sum

$$|B_m|^2 = a_m^2 \left| 1 + e^{-jm\omega_{sw}\tau_1} + \dots + e^{-jm\omega_{sw}\tau_{N-1}} \right|^2, \quad (\text{C.11})$$

where without loss of generality the delay of the first term in C.11 is set as the reference. Expanding out C.11 using Euler's identity,

$$\frac{|B_m|^2}{N^2} = a_m^2 \left[\frac{1 + \sum_{n=1}^{N-1} \cos(-m\omega_{sw}\tau_n)}{N} \right]^2 + a_m^2 \left[\frac{\sum_{n=1}^{N-1} \sin(-m\omega_{sw}\tau_n)}{N} \right]^2. \quad (\text{C.12})$$

Note that C.12 and A.18 are the same, and thus we already solved it in Appendix A. Assuming τ_m is a uniform random variable with mean zero and range $[-B, B]$, then

$$\lim_{N \rightarrow \infty} \frac{|B_m|^2}{N^2} = a_m^2 \text{sinc}^2(Bm\omega_{sw}). \quad (\text{C.13})$$

Finally, using the result in C.13 and applying it to C.10

$$\lim_{N \rightarrow \infty} \frac{P_r}{P_{max}} = D + \frac{1}{D} \sum_{m=1}^{\infty} \frac{a_m^2}{2} \text{sinc}^2(Bm\omega_{sw}). \quad (\text{C.14})$$

It is worth justifying the assumption of a uniform delay spread. The RMS phase-jitter associated with the switching-signal, which takes a Gaussian distribution, is extremely small given that it is derived from the same reference as the RF signal but with a much smaller frequency. That leaves the process variation induced delay spread discussed in Appendix B as the dominant value, which will take a uniform distribution.

Equation C.14 will be solved numerically at this point. A plot of $\frac{P_r}{P_{max}}$ vs. the delay spread is shown in Fig. C.1 and is compared with a numerical simulation in MATLAB using 10,000 elements, a switching frequency of 1MHz and an RF frequency of 10GHz. The match between simulation and theory is excellent. What's notable about Fig. C.1 is the floor of $\frac{P_r}{P_{max}}$ is simply D . This can be explained by the fact that as the switching signal decoheres, the RF carrier remains coherent, setting a floor and the power degradation.




Figure C.1: $\frac{P_r}{P_{max}}$ vs. delay spread for different duty-cycle values.

For purposes of system-level design, we plot the delay range corresponding to $\frac{P_r}{P_{max}} = 95\%$ vs. duty-cycle in Fig. C.2. We note the time requirements are stricter at smaller duty-cycles. For a duty-cycle of 10%, the delay spread (B) should be less than 1% of the switching period, which corresponds to a delay-line resolution of 8-bits. This is readily achievable in CMOS with low-power given the slow switching frequency.




Figure C.2: Delay range at $\frac{P_r}{P_{max}} = 95\%$ vs. duty-cycle.

