

Enabling Technologies and Algorithms for Antenna Arrays

Thesis by
Ailec Wu

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



CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2026
Defended January 13, 2026

© 2026

Ailec Wu

ORCID: 0009-0008-7726-7842

All rights reserved

ACKNOWLEDGEMENTS

The funny thing about a Ph.D. is that despite it representing what is often a small "bump" in humanity's circle of knowledge (see Matt Might's illustrated guide to a Ph.D.), it requires an enormous amount of resources. My reason for seeking a Ph.D. is straightforward: I wanted my program to train me to be able to independently work towards solving problems of weight and depth. I hope, that with the knowledge and skills that I have developed, that I can begin this pursuit.

In my perspective, a Ph.D. is definitely not primarily the product of an individual's efforts; rather it is the product of many different communities selflessly investing in that individual. When I reflect on the past, my journey was made possible by a number of people that I must express my utmost gratitude for. Foremost, I must thank my advisor. Professor Hajimiri has been there for me throughout the highest of highs and lowest of lows; he has guided me through dire straits that I would never be able to navigate alone, celebrated with me achievements that I did not think I would accomplish, and has done so with an unyielding patience and wisdom that I can only one day hope to establish.

Next, I would like to thank my thesis committee, Professors Yaser Abu-Mostafa, Baback Hassibi, and Sander Weinreb for taking time out of their very busy lives to participate in my candidacy and defense examinations and providing helpful feedback and suggestions.

I would like to thank Austin Fikes, who served as an excellent mentor that guided me through my earlier years working on the SSPP project. I would like to thank Elliott Williamstyer for his friendship and guidance as a senior student. I would like to thank Alex Ayling, who has been both a good friend and great colleague: there are few, if any, of my successes that I would not partially attribute to him. I would like to thank Oren Mizrahi, who is ever animated, ever entertaining, and one of the most genuine people I have met. There are few people that I have ever learned more from disagreeing with, and I am sure that most of the time Oren was simply playing devil's advocate. I would like to thank Samir Nooshabadi for his perspectives and advice. He has provided many suggestions that were immensely helpful to me, especially when I have been stumped on issues in my projects. I would like to thank David Baum, who has done the exact same thing, but often in the context of problems more loosely related to process engineering. I would like to

thank Debjit Sarkar, who has been a good friend. I would like to thank Jatin Mathur and Imaad Syed: in the process of working with them, their youth and bright minds give me high hopes for CHIC's next generation of students. I would like to thank all the other CHIC lab members that I have had the privilege of interacting with and learning from: Reza Fatemi, M. Reza, Noyan, Aroutin, Parham, Craig, Armina, Volkan, Zach, Suraj, Jesse, Oliver, Raha, Mohith, and Cole.

Finally, I would like to thank my girlfriend, Kelly. It truly does not feel like an exaggeration to say that she has shown me color in life when I was previously seeing in black and white.

ABSTRACT

Phased arrays provide high angular resolution, efficient spatial power combining, and dynamic electronic beam-steering; although there are many applications of phased arrays, there are considerable challenges in their realization. This thesis summarizes the author's work in developing a novel phased array calibration method to address timing and synchronization challenges in phased arrays, as well as work exploring the usage of phased arrays in wireless power transfer in the context of the Caltech Space Solar Power Project. First, the calibration method is described, which involves usage of amplitude-only sensor placed in the proximal field of the array's radiators: to obtain array phase information from these sensors, machine learning is used to solve the phase retrieval problem generated from sensor raw data. Then the satellite payload producing the first in-orbit wireless power transfer demonstration is described. Finally, a method to address areal power matching issues between photovoltaic and wireless power transfer subsystems is described.

PUBLISHED CONTENT AND CONTRIBUTIONS

- [1] A. Wu, I. A. Syed, A. Ayling, and A. Hajimiri, “Amplitude-only proximal-field sensors and machine-learning-based phase retrieval for array calibration and monitoring,” To be submitted to *IEEE Transactions on Microwave Theory and Techniques*, 2026.
- [2] O. S. Mizrahi, P. Jahelka, E. Gdoutos, *et al.*, “Space solar power generation: A viable system proposal and technoeconomic analysis,” *Joule*, vol. 9, no. 6, Jun. 2025, ISSN: 2542-4785. DOI: 10.1016/j.joule.2025.101928. [Online]. Available: <https://doi.org/10.1016/j.joule.2025.101928>,
A. Wu participated in the conception of the project, analyzed the data, and helped write the manuscript.
- [3] A. Wu, I. A. Syed, A. Ayling, and A. Hajimiri, “On-board array self-calibration using amplitude-only proximal-field sensors and machine-learning-based phase retrieval,” in *2025 IEEE/MTT-S International Microwave Symposium - IMS 2025*, 2025, pp. 930–933. DOI: 10.1109/IMS40360.2025.11103983,
A. Wu participated in the conception of the project, developed the theory, designed the hardware, wrote control software and training software, performed the measurements, analyzed the data, and wrote the manuscript.
- [4] A. Ayling, A. Fikes, O. S. Mizrahi, *et al.*, “Wireless power transfer in space using flexible, lightweight, coherent arrays,” *Acta Astronautica*, vol. 224, pp. 226–243, 2024, ISSN: 0094-5765. DOI: <https://doi.org/10.1016/j.actaastro.2024.08.006>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S009457652400434X>,
A. Wu participated in the conception of the project, developed the theory, designed the hardware (including PCBs, cameras, and rectennas), wrote software (for payload control and ground power detection) and performed the measurements, analyzed the data, and wrote the manuscript.
- [5] A. Ayling, A. Wu, and A. 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,
A. Wu participated in the conception of the project, developed the theory, designed the hardware, wrote control software, performed the measurements, analyzed the data, and wrote the manuscript.
- [6] A. Fikes, E. Gdoutos, M. Klezenberg, *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,
A. Wu participated in the conception, construction, and data analysis for the MAPLE sub payload.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	v
Published Content and Contributions	vi
Table of Contents	vii
List of Illustrations	viii
List of Tables	xiii
Chapter I: Introduction	1
Chapter II: Phased Array Fundamentals	3
2.1 Phased Arrays	3
2.2 Calibration	5
Chapter III: Proximal Field Phase Sensors	8
3.1 Introduction	8
3.2 Prior Art	11
3.3 The Linear Case	12
3.4 Amplitude Only Sensing and Phase Retrieval	15
3.5 Sensor Design and Placement	16
3.6 Machine Learning	23
3.7 Simulated Results	25
3.8 Angular Synchronization	36
3.9 First Generation Hardware	39
3.10 Second Generation Hardware	47
3.11 Conclusion and Future Work	54
Chapter IV: Microwave Array for Power-transfer Low-orbit Experiment	55
4.1 Introduction	55
4.2 Payload Hardware	58
4.3 Software	65
4.4 Mechanical and Thermal Design	67
4.5 Qualification Testing	68
4.6 In-Orbit Experiments	71
4.7 In-Orbit Wireless Power Transfer Experimental Results	72
4.8 Beam to Earth Ground Power Detection	80
4.9 Anomalous Results	85
Chapter V: Mutual Power Optimization	94
5.1 The Area Matching Problem	94
5.2 Principles of Power Amplifier Duty Cycling	99
5.3 Hardware and Measurement	100
5.4 Rectifiers in Wireless Power Transfer	103
5.5 Conclusion	105
Bibliography	106

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
3.1 Gain degradation of (infinite approximation) phased array in presence of phase errors	9
3.2 An example of a far-field range measurement in an anechoic chamber	11
3.3 two element patch array with proximal field sensors	13
3.4 2x2 dipole array augmented by dipole proximal field radiation sensors	14
3.5 Simplified block diagram of on-chip proximal field system from [7] .	15
3.6 Simulated surface current density profiles for an on-chip array	18
3.7 Process to obtain coupling matrix from a design	20
3.8 Comparison of reconstruction error of various phase retrieval algorithms from simulated data assuming knowledge of coupling matrix .	21
3.9 Entropy values averaged across all tile positions for a four by four array of elements from 3.3.	22
3.10 Comparison of training performance for 20 width hidden layer network for various dataset sizes	27
3.11 Model performance for optimistic noise values	28
3.12 Model performance for pessimistic noise values	29
3.13 Model performance for optimistic noise values supplying temperature	30
3.14 Model performance for pessimistic noise values supplying temperature	31
3.15 Model performance for optimistic noise values at SS corner	32
3.16 Model performance for pessimistic noise values at SS corner	33
3.17 Model performance for optimistic noise values at FF corner	34
3.18 Model performance for pessimistic noise values at FF corner	35
3.19 Model performance for moderate $100\mu\text{V}$ noise value supplying temperature at TT corner and 20°C operation	36
3.20 Comparison of least-squares and eigenvector angular synchronization methods	39
3.21 Single tile assembly	40
3.22 Single tile electronics board with labeled components	41
3.23 Single tile antenna and proximal field sensor board, front and back . .	42
3.24 Measurement setup block diagram	43
3.25 Phase offsets due to PLL and reference buffering	44

3.26	Four element array measurement setup	45
3.27	Four element array training results	46
3.28	Four element array beam patterns	46
3.29	Single tile electronics board with labeled components	48
3.30	Single tile antenna and proximal field sensor board	49
3.31	2 by 4 phased array with circular loop proximal field sensors	49
3.32	2 by 4 phased array element labeling (top left), RMS phase error heatmap for both methods (bottom left), and neural network block diagram (right)	50
3.33	Array far-field steering pattern measurements using far-field opti- mization (solid) and phase sensor feedback (dotted)	51
3.34	Neural network architecture initially used to train 4 by 4 array	52
3.35	Initial 10 ResNet-10 model used to train 4 by 4 array	52
3.36	Training loss for 8 different ResNet10 models	53
4.1	The Caltech SSP concept [38]: A photovoltaic-wireless power trans- fer array converts solar energy to DC power and DC power to RF power which is then collected on Earth by a ground station.	57
4.2	MAPLE system block diagram (flight computer is external to MAPLE, but has redundancy in parallel connections to both motherboards of the aggregation panel	59
4.3	Labeled CAD render of MAPLE	60
4.4	MAPLE flight model (FM)	60
4.5	Simplified RFIC block diagram and die micrograph	61
4.6	Top: Front side of flexible phased array (FA) attached to aluminum frame Bottom: Back of FA with vertical bars to heat sink RFIC (translucency for visualization) Labels: (a) popup antennas (b) alu- minum frame (c) RFIC 1 (obscured by frame) (d) microcontroller (e) RFIC 2 (soldered on interposer)	63
4.7	5 × 5 rectenna array of dual-polarized patch antenna receivers.	64
4.8	Rectifier circuit schematic with labeled harmonic terminations	64
4.9	Rectenna efficiency as a function of absolute output power.	65
4.10	SSPD-1 integrated on Vigoride-5 spacecraft (a) MAPLE (wireless power transfer experiment) (b) flight computer (c) ALBA (photo- voltaics experiment) (d) DOLCE (dynamic, deployable structures experiment)	67
4.11	RFIC thermal simulation with heat sink stackup	69

4.12	TAMU proton beam with mounted MAPLE component test panel . . .	70
4.13	Left: MAPLE mounted on vibration table on its X axis Right: Sine vibration excitation and response	71
4.14	MAPLE EM mounted for TVAC testing at JPL post-mission	72
4.15	Rectenna #1 (sidewall). Data from 2023/03/14, 03:19 UTC.	73
4.16	Rectenna #2 (broadside). Data from 2023/03/03, 23:51 UTC.	73
4.17	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.	74
4.18	Rectenna #2 focusing with only 16 elements controlled by RFIC 1. Data from 2023/03/21, 23:16 UTC.	75
4.19	Rectenna #2 focusing with only 16 elements controlled by RFIC 2. Data from 2023/03/22, 00:47 UTC.	75
4.20	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.	76
4.21	In black: unwrapped CMU phase functions, averaged over $N = 6$ interference tests performed during (on-Earth) integration prior to launch, for each core. In red: unwrapped CMU phase functions from the first in-orbit test. The label on the upper left indicates the “coordinates” of the core: (RFIC, quad, stack, core).	78
4.22	Comparing on-Earth to in-orbit PA amplitudes for 32 cores across two RFICs. Error bars indicate spread (standard deviation) among tests performed on Earth (x-axis) and in orbit (y-axis).	79
4.23	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 ($i \in [1, 32]$) for the first core in the stack. Degraded cores circled in red.	80

4.24	Block diagram of RF receiver chain for ground station. There is an additional 3dB of loss due to cabling from the low-noise amplifier to the mixer.	81
4.25	MAPLE antenna pattern through the sapphire window for radiation from Tile 1, Tile 2, and both tiles. Tiles are labeled in figure.	83
4.26	Full ground station assembled before the July 28th, 2023 UTC nighttime pass. a) Narrow beamwidth dish antenna, b) wide beamwidth horn antenna, c) RF receiver chain, power supplies, frequency source, digitizer, etc.	83
4.27	Beam-to-Earth experiment results from 7/28/2023 UTC. Top to bottom: power amplifier current draw for RFIC 1, IF power received at the digitizer, Doppler shift from IF center frequency (5MHz), and orbital tracking maps showing MAPLE during the pass. Orange ellipses in the bottom row depict the -3dB isopower contour for MAPLE's main beam. Three vertical, dotted lines (labeled "1", "2", "3") show the times at which the three bottom photos (1,2,3, respectively) were captured.	84
4.28	The first part of an "Initial Operation Sequence": Two optimizations to Rectenna 2, phase recall to Rectenna 2, one optimization to Rectenna 1, phase recall to Rectenna 1, and then three successive optimizations to Rectenna 1, each optimization starting its search at the last setting of the previous optimization.	86
4.29	Supply current drop anomaly on RFIC 1. Left: Supply current for RFIC 1 (Line A, Line B) measured over time. Right: Supply current for Line A plotted against board temperature. t_{knee} , T_{knee} , and I_{equil} labeled.	88
4.30	Rectified power during linear sweep of phase offset added to RFIC 2.	89
4.31	Long-term system degradation as indicated by two metrics, T_{knee} and I_{eq} , with respect to cumulative power amplifier (PA) usage.	91
5.1	Sankey Diagram [83] tracking losses using medium-term efficiency projections	95
5.2	Typical photovoltaic cell IV curve and MPP	95
5.3	Left: load impedance versus power extracted for typical photovoltaic cell; Right: the same cell's IV curve with superimposed resistive load lines	96

5.4	Ultralight integrated SSPP tile from [68] with areal mismatch: note that the RF aperture is only partially filled	96
5.5	Normalized solar irradiance as a function of PV panel's angle	97
5.6	Usage of a DC-DC converter for MPP by impedance conversion . . .	97
5.7	Projected asymptotic mass breakdown for the wireless power transfer system	98
5.8	Idealized boost converter circuit and switching waveforms	99
5.9	Modulation of a load with capacitive bank	100
5.10	Effect of duty cycle lowering on quiescent operating conditions . . .	101
5.11	Experimental setup for MPO validation	101
5.12	a) Solar cell emulator b) solar cell emulator IV curve c) lab experiment setup	102
5.13	Power received as a function of photocurrent, with assumed DC-DC converter efficiency of 90%	102
5.14	PERC cell setup and PV MPP. Experiment conducted at Caltech on Oct. 16th, 2021.	103
5.15	Rectenna power received as over the course of the day for PERC cells.	104

LIST OF TABLES

<i>Number</i>	<i>Page</i>
3.1 Comparison of various array calibration methods	10
3.2 Summary of simulated chip-level results	35
3.3 Summary of training results for various ResNet block counts	53

Chapter 1

INTRODUCTION

Engineering has served society since at least the dawn of civilization itself. Large scale construction projects relied on civil engineers, while military engineers helped aid in war. With the advent of the Industrial Revolution, the sub-disciplines of chemical, mechanical, and electrical engineering were born. Electrical engineering has had a pervasive role in shaping key aspects of our society involving information and energy transfer. In the 19th century, transatlantic telegraph cables reduced communications times between Europe and the Americas from weeks to minutes. Spark gap transmitters invented by Heinrich Hertz and developed by Guglielmo Marconi enabled maritime communications; later, the development of the vacuum tube and continuous wave systems greatly improved reliability of wireless communications. Alongside these developments in the transfer of information, the invention of the electrical generator and usage of transformers and alternating currents allowed practical long distance power transmission with centralized generators enabling mass production.

In the 20th century, electrical engineering advancements accelerated in pace, many of which were achieved thanks to Moore's Law: a self-fulfilled prophecy of semiconductor development for decades. The integrated circuit led to the microprocessor, and ushered in a new era: the Information Age. More recently, advances in computational power have enabled the practical application of machine learning in optimization, image recognition, and natural language processing—addressing problems previously considered unthinkably challenging. Initially developed for military applications due to their electronic beam-steering capabilities, phased arrays are now a crucial component in wireless communications: both in MIMO cellular networks, and in low-earth orbit high speed satellite internet networks that allow for global connectivity in very remote areas. Their ability to direct and focus power suggests that beyond transfer of information, phased arrays could additionally facilitate the transfer of energy—indeed, most of my time at Caltech has been spent exploring this topic.

In this thesis, I focus on technologies and algorithms that enable antenna arrays. There are considerable challenges with creating arrays, many of which are due to

their strict timing and synchronization requirements. In Chapter 2, I briefly discuss fundamental concepts in arrays and common calibration methods. In Chapter 3, I introduce a new array calibration technique that is meant to aid in the array calibration and monitoring. It utilizes cheap, minimally invasive hardware, and relies on recent technologies in machine learning for data interpretation. In Chapter 4, I discuss MAPLE, a low-earth orbit satellite electronics payload that I worked on that is the first known demonstration of in-orbit wireless power transfer. MAPLE is also a technology demonstrator for the larger Caltech Space Solar Power Project (SSPP), which is an ambitious vision to provide space solar power via a kilometer scale flexible spacecraft equipped with photovoltaics and radio-frequency integrated-circuit phased array wireless power transmission. Finally, in Chapter 5, I discuss Mutual Power Optimization, which is a duty-cycling technique to alleviate challenges encountered in SSPP when connecting photovoltaics and wireless power transfer systems with different optimal areal power densities.

Chapter 2

PHASED ARRAY FUNDAMENTALS

2.1 Phased Arrays

As phased arrays are thematic to the work described in this thesis, this chapter provides a brief summary containing fundamentals of arrays, as well as a description of some array calibration methods.

An antenna array consists of a multitude of antenna elements (which convert between radiated electromagnetic waves and voltages and currents or vice versa). In a phased array, the phases (and frequently also the amplitudes) of each element can be controlled. Adjusting the phases allows for beam-steering: the phased array can dynamically form lobes and nulls in different locations through constructive and destructive wave interference.

There are a few properties in antennas that are helpful for describing arrays.

Reciprocity

Antennas have the same properties whether viewed as transmitters or receivers: they possess the same radiation/receiving patterns. For consistency, subsequent properties will be worded for transmitting antennas though due to reciprocity, the analogous description for a receiving antenna applies.

Radiation Pattern

The radiation pattern of an antenna describes the power it transmits in each direction. In spherical coordinates, this can be described by an azimuth (θ) and elevation (ϕ) angle: $P(\theta, \phi)$.

Directivity

An array's ability to focus its beam over an omnidirectional isotropic radiator of equal total power can be described by its directivity:

$$D(\theta, \phi) = \frac{P(\theta, \phi)}{P_{isotropic}(\theta, \phi)} \quad (2.1)$$

Radiation Efficiency

In practice, an antenna will radiate only a fraction of the power that is delivered to its port. This can be due to both conductor and dielectric loss. The radiation efficiency of antenna is defined to be:

$$\eta = \frac{P_{rad}}{P_{port}} \quad (2.2)$$

Gain

The gain of an antenna describes its ability to focus power relative the the power sent to its port. Thus the gain includes both the aforementioned directivity and radiation efficiency. In other words, the gain of an antenna describes its radiated power in a certain direction divided by an isotropic radiator with the same port power:

$$G(\theta, \phi) = \frac{P(\theta, \phi)}{P_{isotropic,port}(\theta, \phi)} = \eta D(\theta, \phi) \quad (2.3)$$

Array Factor

When placed in an array, superposition can be used to determine the far-field pattern as a summation of the individual element contributions. Frequently, the element patterns are assumed to be the same (this approximation is most accurate for larger arrays with fewer edge elements). In this case, the element pattern becomes a common term in the summation and the array pattern can be viewed as the product of the element pattern and an array factor. For an array of N elements each driven with phase α and amplitude a , the array factor is

$$AF = \sum_{m=0}^{N-1} a e^{j\alpha_m} \quad (2.4)$$

Frequently, elements are linear and equally spaced with uniform amplitude excitation. For an element spacing of d , The array factor can be rewritten:

$$AF = \sum_{m=0}^{N-1} a e^{j\alpha_m} = \sum_{m=0}^{N-1} e^{jmkd \cos \theta} \quad (2.5)$$

2.2 Calibration

There are many challenges in creating physical arrays. The need to calibrate an array is one such challenge: while the design of the array might intend for elements to operate identically, due to differences in the active RF hardware behind each antenna element, significant differences in amplitude and phase can occur in practice.

Array calibration seeks to compensate for these effects; while the numerous methods of array calibration described in the literature follow only loosely standardized terminology, they can be roughly classified into different methods. Following the terminology adopted by [1], the most prominent calibration methods involve near-field scanning probes, peripheral fixed probes, calibration lines, mutual coupling, and phase toggling. As calibration methods rely on different equipments and require different environments, some are primarily suitable for in-factory calibration, while others allow usage in-field. The following discussion of each method seeks to compare their benefits and drawbacks, whereupon a new calibration method will be introduced that is discussed in detail in Ch. 3.

Near-field Scanning Probes

Near-field scanning probes are an in-factory calibration method that utilize a near-field probe to characterize the array. This method is closely related near-field antenna measurement itself: an anechoic chamber and near-field probe scanning system are used to calibrate the array. With a transmitting array, the signal measured by the receiving probe can be used to determine both the amplitude and phase effects of each element relative to the reference signal.

There are a few distinct benefits for this calibration method. This calibration method can be done without prior characterization of the array, and is able to compensate for mutual coupling. Also, unlike some subsequent methods, this calibration method does not require modification of the original array in any way. Finally, if the array is intended to be characterized in the near-field range, this method yields the array's far-field pattern at little additional cost.

Unfortunately, while this near-field probe calibration is quite comprehensive, the requirement for a full near-field range testing environment is very expensive and complex. Besides obviously being restricted to in-factory calibration, this method also scales poorly with a larger arrays as the probe must measure each element one at a time.

Peripheral Fixed Probes

Peripheral fixed probes require the placement of probes at fixed locations around the array. After an initial in-factory calibration, by measuring the coupling between the probes and each array element, amplitude and phase offsets can be determined for in-field calibration. This is possible through using prior knowledge (such as with a near-field scanning probe calibration) to compare what the coupling between the antennas and probes should be (in a good calibration state) against what they currently are in-field.

With this relatively small additional hardware, periodic in-field calibration is possible without needing the advanced motion system of the near-field scanning probe method. However, this method is possible only with a different prior calibration method. Furthermore, used for in-field calibration, space must be allocated for probe placement: in a large array, a calibration system with permanently mounted probes would need to accommodate the large dynamic range due to differences in element to probe distances.

As an example, the rotating-element electric-field vector (REV) method sweeps the phases of individual elements one at a time with a receiver located broadside in the far-field. Once the phase setting corresponding to the maximum of each element is determined, the set of phase codes yields the broadside steering settings for the array (and effectively yields the phase offsets of the system). Multiple iterations can be run for better accuracy (due to mutual coupling), and multiple elements can be swept at once for practical calibration of larger arrays [2].

The REV method allows in-field and in-factory calibration of arrays requiring only an amplitude measurements from its receiver. The array itself need not contain a receiver, and no modifications to the array hardware need to be made. However, the as the receiver should be placed in the radiative region in-field calibration would require additional setup if there is no receiver already present.

Calibration Lines

Array calibration can also be accomplished by adding couplers and transmission lines to array elements. If these passive structures are carefully controlled such that their phase and amplitude effects relative to the element are known, the excitations at each element's antenna can be determined precisely.

This method provides a direct feedback mechanism that enables accurate in-field calibration without modifying the array aperture. However, it is non-trivial to

implement as there will be significant routing costs incurred by the additional transmission lines, as well as RF hardware costs to determine element phase offsets.

Mutual Coupling

In some situations, mutual coupling can be used to calibrate an array. Assuming that the mutual coupling between elements is similar, and that the array is capable of transmitting and receiving across different elements at the same time, a transmission from one element received by 2 elements should be equal in phase and amplitude. Assuming the same mutual coupling, pairs of elements can be adjusted in sequence until the entire array is calibrated.

This approach does not require any additional RF hardware. Furthermore, it allows both in-field and in-factory calibration and does not rely on other methods. However, especially for smaller arrays, the requirement of similar mutual coupling necessitates that addition of dummy elements that modify the array aperture. Also, the simultaneous transmit and receive requirement may be impractical for some arrays (such as those that act only as transmitters).

Proximal Field Sensors

Proximal field sensors, which will be discussed in detail in Ch. 3, are sensors placed in the proximal reactive near-field region of the antenna array. While somewhat similar in structure to peripheral fixed probes, a basic implementation of proximal field sensors is most similar to usage of calibration lines: by knowing the coupling relationship between the sensor and elements, relative phase and amplitude offsets of elements can be determined. This thesis focuses on amplitude-only proximal field sensors, which minimize additional RF hardware costs.

Amplitude-only proximal field sensors allow for in-field and in-factory calibration of a transmitting array with low hardware (e.g. small areas for proximal field sensors connected to simple power detection circuits) costs. However, they require a single calibration point for determination of the array's phase skew: this can be accomplished by the REV method which requires only an incoherent receiver in the far-field.

Chapter 3

PROXIMAL FIELD PHASE SENSORS

3.1 Introduction

Due to their ability to provide high angular resolution, efficient spatial power combining, and dynamic electronic beamsteering, there are a myriad of applications utilizing phased arrays at microwave and millimeter wave frequencies: these include radar, communications, and wireless power transfer. In order to maintain element synchronization, however, it is crucial that phased arrays maintain timing and phase accuracy for each element in the array, particularly for large arrays ¹.

To provide a contextual example, the Caltech Space Solar Power Project (discussed in more detail in Chapter 4) proposes usage of a kilometer scale phased array operating at 10 GHz, corresponding to a filled aperture containing orders of magnitudes of elements in the billions. Though clearly such an ambitious project has many challenges, distributing a reference to each element and successfully synchronizing elements to minimal error for highly efficient wireless power transfer is a major complication.

Assuming uniform, independently distributed random phase errors with a range of $[-\delta_{max}, \delta_{max}]$ for each element for a large array ($n \rightarrow \infty$), it can be shown [5] that the normalized array gain Φ decreases by

$$\Phi(\delta_{max}) = \frac{\sin^2 \delta_{max}}{\delta_{max}^2}.$$

Fig. 3.1 depicts the loss in gain (for wireless power transfer systems, efficiency) as a function of the root mean square phase error. Assuming uniformly distributed phase errors in each element, such a system tolerating a 10% efficiency loss due to element phase error would allow only ~ 18 degrees RMS phase error globally across each element: at 10 GHz, this corresponds to 5 picoseconds of timing error. While planning and development of the reference distribution and synchronization of the SSPP phased array are still under consideration, preliminarily a hierarchical reference distribution scheme has been conceptualized. Such a scheme would potentially

¹Alex Ayling and Imaad Syed were partners in the Proximal Field Phase Sensors concept formation, system design, measurements, and documentation. This chapter includes material from [3] and [4].

include some combination of matched transmission lines, low frequency reference distribution, wireless reference distribution, and optical reference distribution: each of which possess trade-offs in both power dissipation and noise. For instance, while a fiber optic reference distribution would benefit over long distances due to minimal noise and lower loss compared to a matched transmission line, over shorter distances the relatively expensive and complicated requirement of optoelectronics and fiber coupling naturally suggests using electronic reference distribution over short distances paired with fiber optics over longer distances. A more detailed discussion and consideration of the reference distribution problem in SSPP (and very large scale array) context can be found in [6].




Figure 3.1: Gain degradation of (infinite approximation) phased array in presence of phase errors

For smaller arrays (or local segments of a large array), calibration is still not trivial though there exist many diverse methods as discussed in Chapter 2. Summarily, calibration methods can be roughly classified as in-field or in-factory [1]. For phase calibration, both types of methods typically rely on first establishing knowledge of coupling between elements and a probe, then subsequently determining the additional phase offset. A table comparing some basic array calibration methods

(summarizing Chapter 2) including the proposed proximal field phase sensors is given in Table 3.1.

Calibration Method	In-Factory Suitable	In-Field Suitable	Hardware Needed	Aperture Enlargement	Accuracy	Speed
Near-field Scanning Probes	Yes	No	V. High	No	V. High	V. Slow
Peripheral Fixed Probes	No	Yes	Moderate	Yes	High	Normal
Calibration Lines	Yes	Yes	High	No	High	Fast
Mutual Coupling	Yes	Yes	Moderate	Yes	High	Normal
Phase Toggling	Yes	Yes	Low	No	Low	Slow
Proximal Sensors	Yes	Yes	Low	No	Low	Fast

Table 3.1: Comparison of various array calibration methods

Along with calibration, characterization of arrays must be done. This is typically done in a controlled environment (usually an anechoic chamber with radiation absorbent material for the frequency of interest). Here, the antenna under test (AUT) is placed some known distance away from a well-characterized instrumentation antenna, upon which both antennas are rotated correspondingly to measure the AUT's far-field pattern at different angles. An example of a far-field range measurement is given in Fig. 3.2. For classic far-field ranges, the AUT is often paired with a standard gain horn. However, for AUTs with large aperture size and wavelength, a far field range may not be practical due to requirement of a correspondingly large chamber.

Near-field ranges address this by characterization of the AUT in the near-field via spatially scanning a probe (frequently in a planar, cylindrical, or spherical manner). As opposed to a far-field measurement, compensation of the probe's mutual coupling to the AUT must be done carefully which creates significant additional complexity. After compensation, the near-field profile is mathematically transformed to the far-field via corresponding Fourier transform, cylindrical harmonics, or spherical harmonics.



Figure 3.2: An example of a far-field range measurement in an anechoic chamber

3.2 Prior Art

Proximal field radiation sensors, first introduced in [7], are highly localized sensors integrated on the same substrate as a system's transmitting antennas. While technically existing in the system's near field, their primary distinction from a standard near-field probe lies in their substrate integration with the system antennas. This bestows them distinct advantages: they do not need additional external test equipment, can provide system information in-situ potentially even foregoing an anechoic environment, and can operate alongside the system enabling self-calibration, self-characterization, and self-monitoring.

A particularly attractive application of proximal field sensors lies in monitoring and calibration of millimeter-wave systems with on-chip radiators. This is for a multitude of reasons that both include challenges faced when integrating radiators as well as benefits from highly localized sensing. In particular, systems with on-chip radiators incur efficiency loss due to surface wave propagation in lossy dielectric of the die substrate, performance degradation due to sensitivities to PVT variations, and can be difficult and expensive to test and characterize. Incorporating proximal field sensors that utilize energy otherwise lost to surface modes could provide valuable

information about the system.

Unfortunately, akin to the challenges faced by near-field range systems, a proximal field system operates in a complicated electromagnetic environment that is difficult to characterize due to surface waves, element mutual coupling, and sensitivity to geometry/fabrication tolerances. Furthermore, while they may not require external test equipment, if phase information is to be desired, coherent detection must be implemented, which at the chip level, would incur significant complexity and cost precious die area. At the PCB level, the sensors might be excessive or impractical compared to existing methods such as usage of coupled transmission lines.

3.3 The Linear Case

For a system with n transmitting antennas and m proximal field sensors, a coupling matrix containing $n + m$ ports can be used to describe the relationship between the radiators and sensors. In narrow-band operation, sweeping radiator excitation phases will yield a sinusoidal interference pattern. As a toy example, for a two element patch array, in Fig. 3.3 the simulated output voltage at two particular sensors is measured as the relative phase of the two radiators is swept. Notice that due to different phase offset terms in the coupling matrix, the maximum and minimum voltages do not correspond to 0 and 180 degrees relative radiator drive phases.

Now, consider the more complicated two-by-two dipole phased array in Fig. 3.4, which depicts the radiators and proximal field sensors of a integrated on-chip transmitting system. Eight proximal field sensors (also dipoles) have been placed around the periphery of the transmitting antennas.

Assuming n element control phases $\phi_i \in [0, 2\pi)$, that each sensor coherently recovers a signal $s_k \in \mathbb{C}$, and that each antenna couples $a_{i,k} \in \mathbb{R}$ in amplitude and $\theta_i \in [0, 2\pi)$ in phase, the overall sensor reading will be (omitting an arbitrary global phase)

$$s_k = a_{1,k}e^{j\phi_1} + a_{2,k}e^{j(\theta_2+\phi_2)} + a_{3,k}e^{j(\theta_3+\phi_3)} + a_{4,k}e^{j(\theta_4+\phi_4)}. \quad (3.1)$$

In matrix form, if $A_{i,k} = a_{i,k}e^{j\theta_i}$ and $\mathbf{x} = \begin{bmatrix} e^{j\phi_1} & e^{j\phi_2} & \dots & e^{j\phi_n} \end{bmatrix}^T$,

$$\mathbf{s} = \mathbf{A}\mathbf{x}.$$

In principle, if one has control over $\mathbf{x}$ and coupling matrix $\mathbf{A}$ is sought, a set of at least n measurements and sensors $\mathbf{X} = \begin{bmatrix} \mathbf{x}_1 & \mathbf{x}_2 & \dots & \mathbf{x}_n \end{bmatrix}$ can be used to solve for $\mathbf{A}$ as

$$\mathbf{S} = \begin{bmatrix} \mathbf{s}_1 & \mathbf{s}_2 & \dots & \mathbf{s}_n \end{bmatrix} = \mathbf{A}\mathbf{X}. \quad (3.2)$$




Figure 3.3: two element patch array with proximal field sensors

With additional sensors and measurements, an overdetermined system can be formed and the pseudo-inverse could be used to give a practical result more robust to noise: $\mathbf{A} = \mathbf{S}\mathbf{X}^+$. In [7], both a PCB-based and chip-based system are built and similar reasoning is used to determine relative phase differences between elements: here, during simulation additional "sniff" antennas were added in the far-field region of the radiator for a convenient determination of the far field pattern of the system in accordance with the receiving or sensing ports described in the Lavaei-Babakhani method [8]. This can be done in a single full-wave simulation (followed by post-




Figure 3.4: 2x2 dipole array augmented by dipole proximal field radiation sensors

processing steps) as the linearity of Maxwell's equations with respect to excitation allows the simulator to calculate and use Green's functions to determine the relationship between ports.

Although in theory, sufficient placement of sensors sampling the near-fields can give direct information for determination of the far-field pattern of a radiator, the quantity of sensors necessary and corresponding overhead may not be realistically achievable. Practically speaking, for many arrays, determination of the antenna phases is key as some knowledge of the element patterns can be discerned both from simulation and measurement: the superposition of element patterns combined with the excitation phase offset can be used to predict the far field pattern.

With access to a full-wave simulator such as HFSS, a more direct and accurate (albeit less convenient) method for array far field pattern estimation is to extract the element far field patterns (available in HFSS as the far-field pattern is calculated with use of free space Green's functions) after the simulation is complete and to calculate the superposition sum externally. As this reduces the simulation volume compared to a naive implementation of sniffing ports, it also has the added benefit of reducing memory requirements and speeding up the simulation.

A natural question to pose would be the viability of using an incoherent amplitude only detector for the sensor front-ends which vastly simplify the sensor circuitry as it would remove the need to distribute a reference or down-convert the signal. This amplitude only detector could be implemented as a peak detector circuit, RMS power detector circuit, or logarithmic circuit that would output a low frequency signal that could be easily read by an ADC. As the incoherent detector is no longer beholden to reference distribution, there is added potential for the detector to be designed in a more drop-in fashion minimizing need for co-design of the transmitter

and proximal field receiver circuitry.

While discarding the need to detect phase is highly convenient from a hardware perspective, it invalidates the ability to use straightforward matrix algebra for the proximal field sensors: both to determine the radiator to sensor coupling and to determine the radiator excitations from sensor readings (once the coupling is found). With only sensor magnitudes Eq. 3.1 becomes:

$$|s_k|^2 = |a_{1,k}e^{j\phi_1} + a_{2,k}e^{j(\theta_2+\phi_2)} + \dots + a_{n,k}e^{j(\theta_n+\phi_n)}|. \quad (3.3)$$

Assuming prior knowledge of the coupling environment (e.g., from simulation), finding radiator excitation phases from measurements of m sensors is a phase retrieval problem which can be mathematically stated as

$$\begin{aligned} & \text{find } \mathbf{x} \\ & \text{subject to } s_i = |\langle \text{row}_i(A), \mathbf{x} \rangle|^2 \text{ for } i \in \{1, \dots, m\} \end{aligned} \quad (3.4)$$

Phase retrieval problems arise in a wide range of problems in engineering and applied physics including optics [9], X-ray crystallography [10], and astronomy [11], among others. With their rich history phase retrieval problems have been heavily studied in literature. In this context, the rows of the coupling matrix A form a frame while the antenna excitations ϕ are the desired unknown signal.

Without prior knowledge of the coupling environment, determination of the coupling environment can also be shown to be a phase retrieval problem with the usage of the Kronecker product and associated identities on Eq. 3.2:

$$\text{vec}(\mathbf{S}) = \text{vec}(\mathbf{A}\mathbf{X}) = (\mathbf{X}^T \otimes I_k)\text{vec}(\mathbf{A}) \quad (3.5)$$

With the corresponding magnitudes taken, the vectorized $\mathbf{A}$ coupling matrix is the desired unknown signal, while the matrix $\mathbf{X}^T \otimes I_k$ forms the frame for the phase retrieval problem.

3.5 Sensor Design and Placement

Some preliminary work on the choice and placement of proximal field sensors is discussed in [12]; much of this discussion focuses on modeling the substrate as a waveguide and designing sensors that respond to specific TE/TM modes of the substrate.

The importance of the placement of the proximal field sensors can be intuitively understood by examining the effect of sensor position on sensor output. To estimate


the response of a given sensor placement, Fig. 3.6a depicts HFSS simulated surface current densities of a chip-scale two by two dipole array when each element is driven in phase, with a proposed proximal field sensor placement marked by black dots. When driven at a random phase as shown in Fig. 3.6b, the surface current density profile changes. Note that in this example, the middle of the chip appears to form a null; placement of a sensor here would perform poorly as the small signal swing would be more heavily affected by noise. Qualitatively, after ensuring that they do not interfere with radiator operation, sensors should be placed such that they spatially sample the fields (to provide independent information) whilst having large signal swings under different phase conditions.

Measurement Matrices/Frames


In frame theory, there are a few key results that pertain to the number of sensors required for antenna phase recovery. A well known result [13] states that, in the absence of noise, for N phaseless measurements and a signal of length M , $N = 4M - 4$ measurements are sufficient to uniquely determine the signal; at least $4N + o(1)$ measurements are necessary asymptotically [14]. For proximal field sensors, this naturally suggests that an array with M antennas would need at least $N = 4M - 4$ sensors for phase reconstruction (provided of course that the resultant coupling matrix was known).

One way to view the proximal field sensor placement problem is as a measurement matrix design problem. In that context, a discussion of existing work for a few measurement matrices can be found in [15], which uses information theory to select an ideal measurement matrix (in contrast to the assumption of a given matrix). Notably, many works in the literature deal with random matrices (containing Gaussian i.i.d. matrices) [16], discrete Fourier transform matrices [17], or short-time Fourier transform matrices [18]. However, coupling matrices generated by proximal field sensors differ from these matrices as they are not necessarily random, and will exhibit sparsity as nearby sensors will couple more strongly to nearby antennas. These coupling matrices furthermore depend on openness in choice of proximal field sensor geometry and different potential far-field antenna selections.

While the coupling matrices for proximal field sensor systems may not be typical, the design freedom of proximal field sensors might enable existing phase retrieval algorithms to be applied. One such algorithm is the polarization approach [19], which is briefly described below.



(a) In phase



(b) Random phase

Figure 3.6: Simulated surface current density profiles for an on-chip array

For some finite set V , let $\mathbf{A}_V = \{\mathbf{a}_j\}_{j \in V} \subset \mathbb{C}^M$ be a measurement frame: we seek $\mathbf{x}$ subject to $|\langle \mathbf{x}, \mathbf{a}_j \rangle|^2 = b_j$. Assuming a pair of frame vectors $\mathbf{a}_i$ and $\mathbf{a}_j$ yield nonzero frame coefficients $\langle \mathbf{x}, \mathbf{a}_i \rangle \neq 0$ and $\langle \mathbf{x}, \mathbf{a}_j \rangle \neq 0$, the *relative phase* between the frame coefficients can be defined to be:

$$\omega_{i,j} = \left(\frac{\langle \mathbf{x}, \mathbf{a}_i \rangle}{|\langle \mathbf{x}, \mathbf{a}_i \rangle|} \right)^{-1} \frac{\langle \mathbf{x}, \mathbf{a}_j \rangle}{|\langle \mathbf{x}, \mathbf{a}_j \rangle|} = \frac{\overline{\langle \mathbf{x}, \mathbf{a}_i \rangle} \langle \mathbf{x}, \mathbf{a}_j \rangle}{|\langle \mathbf{x}, \mathbf{a}_i \rangle| |\langle \mathbf{x}, \mathbf{a}_j \rangle|}. \quad (3.6)$$

Note that with this definition, we can relate the relative phase to new, phase recovered frame coefficients c_j as

$$c_j = \omega_{i,j} |\langle \mathbf{x}, \mathbf{a}_j \rangle| \quad (3.7)$$

Viewing the frame vectors $\mathbf{a}_k$ as vertices in a graph and the relative phases between them as the edges, a sufficient number of relative phases forms a connected graph allowing recovery of $[\mathbf{x}]$: the equivalence class of $\mathbf{x}$, differing from $\mathbf{x}$ by an arbitrary global phase. This is done by multiplying by the corresponding vectors $\{\tilde{\mathbf{a}}_j\}_{j \in V}$ of the Moore-Penrose pseudoinverse of $\mathbf{A}_V$ (taking $b_j = 0$ and $\omega_{i,j} = 1$ when $\langle \mathbf{x}, \mathbf{a}_j \rangle = 0$):

$$\sum_{j \in V} c_j \tilde{\mathbf{a}}_j = \sum_{j \in V} \omega_{i,j} |\langle \mathbf{x}, \mathbf{a}_j \rangle| \tilde{\mathbf{a}}_j = \left(\frac{\langle \mathbf{x}, \mathbf{a}_i \rangle}{|\langle \mathbf{x}, \mathbf{a}_i \rangle|} \right)^{-1} \mathbf{x} \in [\mathbf{x}]. \quad (3.8)$$

In a more realistic situation with noise, the pseudoinverse method can yield errors due to phase wrapping and can instead be replaced with a more robust angular synchronization method such as the ones mentioned in Section 3.8.

To obtain the relative phases between frame coefficients, a cube root polarization identity is used (a special case of [20]) relating $a, b \in \mathbb{C}$:

$$\overline{a}b = \frac{1}{3} \sum_{k=0}^2 e^{j\frac{2k\pi}{3}} |a + e^{-j\frac{2k\pi}{3}} b|^2 \quad (3.9)$$

To apply this identity, take $a = \langle \mathbf{x}, \mathbf{a}_i \rangle$ and $b = \langle \mathbf{x}, \mathbf{a}_j \rangle$ and observe that

$$\overline{a}b = \overline{\langle \mathbf{x}, \mathbf{a}_i \rangle} \langle \mathbf{x}, \mathbf{a}_j \rangle = \frac{1}{3} \sum_{k=0}^2 e^{j\frac{2k\pi}{3}} |\langle \mathbf{x}, \mathbf{a}_i + e^{j\frac{2k\pi}{3}} \mathbf{a}_j \rangle|^2. \quad (3.10)$$

In other words, if we take pairs of frame vectors from the original frame $\mathbf{A}_V$, then take measurements with respect to their sum rotated by 0, 120, and 240 degrees, we can determine the relative phase of our frame coefficients. This could be accomplished physically in proximal field sensors with the use of an RF or microwave network that combined the readings of two sensors at different phases; by incorporating a switch-based phase shifter with preset angles, only one additional power detector would need to be added per pair of existing sensors. The proposed scheme does have practical limitations, however, as design of the multi-port network would need to be compact to minimize its burden. Additionally, the network would need to consider sensitive proximal field interactions, or incur additional costs if implemented as a separate

layer on the PCB for isolation. Finally, ensuring repeatability of the network despite fabrication geometry variations and characterization could be significant practical issues.

These types of difficulties for the hardware implementation of the polarization phase retrieval method indicate that for any useful real instance of amplitude-only proximal field sensors, it is imperative that they are small and unobtrusive, provide accurate and robust information for phase retrieval, and are relatively simple to design and fabricate.

Using simulation as a design aid, extracted s-parameter matrices from full wave electromagnetic simulation can be used to calculate the corresponding coupling matrix between sensors and antennas. As the proximal field sensors used are small relative to radiators typically of resonant length, they can be modeled as lumped elements and do not appear as a standard characteristic impedance. Viewing the coupling matrix in the form of Y- or Z-parameters converted from S-parameters are more convenient to work with. An approximation can be made to neglect the effect of sensor to sensor cross coupling, which, for M antennas and N sensors allows the formation of an M -by- N sub-matrix from the original $N + M$ -by- $N + M$ sub-matrix. This matrix is directly usable as the frame for proximal field phase retrieval simulations; the diagrams in Fig. 3.7 depict the conversion process from s-parameters generated by electromagnetic simulation to frame.




Figure 3.7: Process to obtain coupling matrix from a design

As the many available phase retrieval methods are often intended for different types of problems, testing a variety of them for proximal field sensors is useful. PhasePack [21] provides this functionality. Fig. 3.8 contains a comparison of the simulated performance of a few existing phase retrieval algorithms (via PhasePack) for a

coupling matrix generated from full-wave simulation of a four-by-four phased array with 96 sensors. While Fig. 3.8 indicates that both Wirtinger flow and amplitude flow methods converge, it is important to note that significant further investigation is necessary to practically apply these algorithms to a proximal field phased array system: good estimates of the coupling matrix are difficult to make due to the complex electromagnetic coupling environment in the proximal field. As shown in Eq. 3.2, determination of the coupling matrix from the phaseless sensor readings is an additional phase retrieval problem that must be addressed and is challenging to verify both due to requirement of a testbed with accurate phased sensor readings, as well as due to nonlinearities and nonidealities in real RF sensors and power detectors.




Figure 3.8: Comparison of reconstruction error of various phase retrieval algorithms from simulated data assuming knowledge of coupling matrix

Entropy

Another way to look at the proximal field sensor placement problem when using amplitude-only sensors is by using concepts from information theory.

If the radiator excitation phases are unknown, the amplitudes that a particular sensor can generate can be viewed as a random variable $I(x, y)$ —the distribution of this random variable will depend on the sensor's horizontal and vertical position. Given that this sensor amplitude $I(x, y)$ has corresponding probability density function $f_I(i)$, the differential entropy for $f_I(i)$ is defined to be

$$H(I) = \int_{-\infty}^{\infty} f_I(i) \log(f_I) di. \quad (3.11)$$

As a visual, for the geometry shown in Fig. 3.3, when elements are configured in a four by four array, Fig. 3.9 depicts the average entropy values across all positions (the antenna region has been blacked out).

Information (Entropy) on a Typical Tile




Figure 3.9: Entropy values averaged across all tile positions for a four by four array of elements from 3.3.

While choosing a single placement that maximizes $H(I)$ would be effective for a single sensor, when placing multiple sensors, there could be shared redundant information. To address the desire to maximize the total amount of *independent* information, measuring the redundant information with mutual information [22] is helpful, defined as

$$M(I_1, I_2) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_{I_1, I_2}(i_1, i_2) \log\left(\frac{f_{I_1, I_2}(i_1, i_2)}{f_{I_1}(i_1) f_{I_2}(i_2)}\right) di_1 di_2, \quad (3.12)$$

where $f_{I_1, I_2}(i_1, i_2)$ is the joint probability density function of the intensities. To extend this method for more than two sensors to m sensors, the conditional entropy $H(i_1|(i_2, \dots, i_m))$ should be calculated. By taking inspiration from the "OBO" approximation of mutual information in [23] conditional entropy can be approximated as

$$H(i_1|(i_2, \dots, i_m)) \approx H(i_1) - \sum_{k=2}^m M(i_1, i_k); \quad (3.13)$$

which enables fast computation of multivariate conditional entropy at the cost of slight underestimation.

As there are multiple sensors each possessing their own conditional entropy, the overall optimization problem is

$$\operatorname{argmax}_{S_l} \sum_{k=1}^{k=6} H(i_k|(S_l \setminus i_k)), \quad (3.14)$$

where $S_l = \{i_1, \dots, i_6\}$ is the list of sensor location indices. This optimization problem has thus far been approached with genetic optimization and a greedy algorithm. Both methods seek to find the set of 6 sensor placements yielding maximum total information.

3.6 Machine Learning

Usage of deep (supervised) neural networks are proposed to solve the associated phase retrieval problems for the amplitude-only proximal field sensors. Of course, this is not a new approach, and there are many works such as [24] that use deep learning for phase retrieval in the optics domain; a recent review containing information about deep learning and phase retrieval can be found in [25]. In the following paragraphs, both the merits of using machine learning and the ways that it enables amplitude-only proximal field sensors will be discussed.

Motivation and Background

As previously mentioned, the two phase retrieval problems encountered in the current amplitude-only proximal field sensors create challenges that do not lend themselves to obvious solutions with classical phase retrieval methods. Foremost, the coupling matrix (frame) must be determined. However, in practice, the coupling pathways between antennas and sensors could be complex due to effects from surface waves, near-field radiation, and element interaction, making them difficult to model and accurately simulate and estimate. The supervised learning approach used in

this work allows the model to simultaneously learn the coupling matrix as well as how to predict antenna phases from sensor amplitudes.

There are other advantages when using machine learning in this context. Foremost, the process of learning a model enables one to encompass many details that exist in real systems but may not be considered in a pure phase retrieval context. For example, these include the non-linearity of the power detector circuit and system dependence on PVT variations for components in the transmitter chain. Beyond that, machine learning is flexible to changes in radiator and sensor combinations. Of course, many of these advantages are effectively bestowed from a data-driven philosophy, of which the feasibility of data generation and collection should be posed. This question will be explored in more detail after first examining and discussing some simulations regarding system feasibility and performance.

Model and Training

Both model architecture and training methods were developed alongside the hardware of the proximal field sensors. At the very beginning, all neural networks failed to converge on the raw training data; there were two necessary dataset modifications that needed to be made. The first was conversion of target angles relative to a singular element into angles relative to neighboring elements; the second was conversion of the resultant angular value into trigonometric cosine, sine components.

Initially, the choice of a shallow single hidden layer feedforward network was driven by ease of use; at the time the performance was deemed sufficient such that effort towards improvements were focused on other aspects of the system. The initial model focused on using Levenberg-Marquardt backpropagation with a few hundred learnable parameters; this backpropagation method was empirically found to perform better than gradient methods, perhaps due to the second order method being better at navigating the (likely bumpy with many local minima) optimization landscape. However, shifting from a 2 by 2 array to a 2 by 4 array quickly indicated that adjustments would be needed for acceptable inference accuracy. Initially, a few simulations indicated promise with shallow convolutional neural networks, but while these models worked well on 2 by N arrays, they failed to provide adequate performance for 2D arrays of more than 2 rows. Eventually, based on prior literature [24] indicating usage of deep networks on phase retrieval problems, ResNet (with the Adam optimizer) was used to fit the training data yielding much better results than before.

The choice of ResNet versus other deep networks is admittedly somewhat (though not completely) of an arbitrary one. Because of the fixed size and arrangement of the proximal field sensors, a neural net intended for image recognition seemed most appropriate as a starting point. However, because the physically constructed arrays were much smaller in size than the typical image resolution, architectures with large output resolution such as encoder-decoder were not considered. The widths of the ResNet blocks, as well as other hyperparameters such as the learning rate, were tuned via Ray Tune.

In some regards, the ResNet choice could be considered to be one made of convenience, but there were other practical limitations for training mostly due to limited hardware resources. Training (including hyperparameter tuning) was done on a single personal consumer grade GPU; this significantly limited the number of runs and opportunity for model exploration.

Given the model architecture exploration, a practical future system utilizing amplitude-only proximal field phase sensors in a similar way should be designed with more experimentation and depth in the machine learning side.

3.7 Simulated Results

To validate the concept of using neural networks for proximal field sensors, first, a study was performed for a hypothetical two-by-two on-chip phased array system operating at D-band. The dimensions of the chip and circuit details were roughly based on the real system in [26] with radiator and sensor arrangement identical to that of Fig. 3.4. Following that, simulations for a design implemented at 2.5 GHz intended for discrete components with a printed circuit board were done to ensure a straightforward path to physical hardware implementation.

Chip Level Simulations

The chip level simulations are based on a two-by-two 140 GHz phased array system. The radiating elements are resonant-length dipoles whilst the proximal sensors are electrically short dipoles. The power detector circuit for the proximal sensors is based on a 22nm process and is a cross coupled transistor rectifier.

After full-wave simulation of the electromagnetic structure to yield an s-parameter matrix, Cadence Virtuoso was used (with circuit simulation assuming random antenna phases and equal antenna drive amplitudes) to generate a simulated training dataset consisting of 5000 data points. Each data point contains 8 sensor voltages

as features, and three phase angles (zero referencing the first antenna) as labels. Gaussian noise was added corresponding to expected levels at the ADC between $60\mu\text{V}$ and $150\mu\text{V}$.

A shallow real-valued neural network with two feedforward layers (one hidden layer) was used to fit to the training data. The hidden layer used a sigmoid activation while the output layer had no activation. The inputs and outputs of the data were normalized to $[-1, 1]$. To address loss function convergence issues due to the periodicity of the angular labels, each of the three angles θ_i was encoded as a $(\cos \theta_i, \sin \theta_i)$ pair during training. After training, the actual angular deviation between model predictions and ground truth is calculated (taking care to observe that the angular difference should never exceed 180 degrees) as a post processing step. Training was done with the Levenberg-Marquardt algorithm, which was selected empirically for its significantly faster and better fit convergence versus first-order gradient methods.

Parameters swept in these simulations include the number of neurons (width) of the hidden layer, as well as the noise, training dataset size, and performance across PVT variations. As a starting point, because preliminary simulations indicated that 20 neurons in the hidden layer was roughly sufficient for adequate performance, a comparison of the inference performance of a 20 neuron hidden width model across various dataset sizes and noise levels was used to estimate the number of data points needed to effectively train the model. As indicated in Fig. 3.10, dataset sizes above 500 points have minimal difference in training performance across different noise levels; for this reason, for this particular model architecture and problem, 1000 data points was picked as the standard size for studying effects of varying parameters.




Figure 3.10: Comparison of training performance for 20 width hidden layer network for various dataset sizes

First, operation of the chip at its typical-typical process corner was considered with an optimistic minimum noise value of $60\mu\text{V}$. Here, the temperature of the chip is randomized in simulation to a value in its full operating temperature range of $[0, 100]^\circ\text{C}$. As shown in Fig. 3.11, a minimum of 6 neurons in the hidden layer provide a reasonable fit but with more neurons the model accuracy continues to improve, with 32 neurons in the hidden layer providing the lowest test inference of 1.55 degrees RMS error. The short-tailed distribution suggests that outliers should be rare.




Figure 3.11: Model performance for optimistic noise values

Next, operation of the chip at its typical-typical process corner was considered with a pessimistic high noise value of $150\mu\text{V}$ and identical operating temperature assumptions. As shown in Fig. 3.12, similar to the $60\mu\text{V}$ case, a minimum of 6 neurons in the hidden layer provide a reasonable fit but with more neurons the model accuracy also continues to improve; 28 neurons in the hidden layer provides the lowest test inference of 2.00 degrees RMS error. The increase in RMS error is manifested by a spreading in the distribution, yielding a significant increase in higher error results as well as outliers. However, the test set does not contain outliers outside of 14 degrees error (well under the 18 degree requirement for less than 10% array gain degradation).




Figure 3.12: Model performance for pessimistic noise values

During the training process it was observed that the power detector circuit had a significant temperature dependence. Despite this, the previous results indicate that the model is compensating for the effects of temperature without explicit information of the sensor operating temperature. A natural question to ask would be if the model could make use of temperature information—this could be implemented as a PTAT temperature sensor which may already be included in the integrated circuit due to thermal monitoring/management requirements.




Figure 3.13: Model performance for optimistic noise values supplying temperature

As a rough approximation, a temperature accurate to 1°C is supplied to the neural network in Fig. 3.13 and Fig. 3.14 as the only change respective to Fig. 3.11 and Fig. 3.12. The addition of temperature information significantly bolsters the trained model's inference performance, with the best model with $60\mu\text{V}$ noise having 0.812 degrees RMS error, while the best model with $150\mu\text{V}$ has 1.25 degrees RMS error (less than that of the optimistic noise case without temperature information).




Figure 3.14: Model performance for pessimistic noise values supplying temperature

The power detector circuit was also observed to have significant dependence on process variation. A study was done across the SS and FF process corners (the power detector circuit contains only NMOS devices) of the chip. Respective to the noise values (with no supplied temperature information) in Fig. 3.11 and Fig. 3.12, Fig. 3.15 and Fig. 3.16 depict the simulated model performance for the SS corner, while Fig. 3.17 and Fig. 3.18 depict the simulated model performance for the FF corner. The SS corner performs roughly 0.3 degrees better than the TT corner with optimistic noise, while the FF corner performs roughly 0.8 degrees worse than the TT corner in both noise conditions.




Figure 3.15: Model performance for optimistic noise values at SS corner




Figure 3.16: Model performance for pessimistic noise values at SS corner




Figure 3.17: Model performance for optimistic noise values at FF corner




Figure 3.18: Model performance for pessimistic noise values at FF corner

Noise Level [μV]	Process Corner	Temperature Supplied	RMS Test Error [deg.]
60	TT	no	1.55
60	TT	yes	0.81
60	SS	no	1.27
60	FF	no	2.37
150	TT	no	2.00
150	TT	yes	1.25
150	SS	no	1.92
150	FF	no	2.73

Table 3.2: Summary of simulated chip-level results

Finally, the effects of quantization were studied to determine the ADC resolution that would be necessary for model inferences and training. In Fig. 3.19, the effects of quantization are compared by sweeping the number of quantization bits from two to twelve and observing the performance of the model (chosen to have the optimal

hidden layer width). An ADC with 8 bits appears sufficient with higher quantization levels yielding rapidly diminishing returns in model error.




Figure 3.19: Model performance for moderate $100\mu\text{V}$ noise value supplying temperature at TT corner and 20°C operation

3.8 Angular Synchronization

As mentioned in Section 3.7, model inferences for the rectangular array predict pairwise phase differences of neighboring elements in the horizontal and vertical directions. An additional step must be made to go from these pairwise values to the original set of phases (up to arbitrary global offset phase). In the literature, this problem is known as the angular synchronization problem [27], and has applications including time synchronization of distributed networks [28], surface reconstruction in computer vision [29], and optics [30]. For a full-scale SSPP system incorporating a hierarchical reference distribution scheme, angular synchronization will likely be necessary for any component proposing local, distributed calibration.

To state the angular synchronization problem more formally [27]: estimate n unknown angles $\theta_1, \dots, \theta_n \in [0, 2\pi)$ from m noisy measurements of their offsets $\theta_k - \theta_l$

mod 2π . For model inferences described in this thesis, angular synchronization is usually done via least squares. Pairwise adjacent measurements of rectangular arrays produce a measurement graph of a rectangular grid; the vertices of the graph correspond to array excitation angles relative to a global reference, while the edges of the graph represent the pairwise measurements. To formulate the problem for least squares solution, difference equations are used for each pairwise measurement (e.g. $\theta_1 - \theta_2 = \delta_{12}$ and a matrix for the full system is written. For an R by C array, $(R - 1)C + R(C - 1) = 2RC - R - C$ equations will exist and thus produce an overdetermined system (assuming arrays above some minimum R by C) allowing least squares solution.

The primary complication that will occur after applying least squares is accounting for the phase wrapping information that is lost as pairwise measurements are modulo 2π . This issue can easily be observed by multiplying a particular set of angles by the matrix and noting that the resultant pairwise angles lie outside the interval $[0, 2\pi)$. In other words, for the matrix equation $\mathbf{A}\vec{\theta} = \vec{\delta} + 2\pi\vec{x}$ we seek $\vec{\theta}$ having knowledge of only $\mathbf{A}$ and $\vec{\delta}$.

Determination of the number of phase wraps can be found by seeding the least squares calculation with a less susceptible alternative estimation method: a spanning tree using only a subset of the edges of the least squares calculation. Equations corresponding to rows in the matrix are removed, and (accounting for the arbitrary global offset phase) the submatrix $\mathbf{B}$ with corresponding pairwise measurements $\vec{\delta}'$ produces a system that does not suffer from the phase-wrapping error: $\mathbf{B}\vec{\theta}' = \vec{\delta}'$. The value of $\vec{\theta}'$ is used in the overdetermined system to calculate phase wraps: $2\pi\vec{x} = \mathbf{A}\vec{\theta}' - \vec{\delta}$

Of course, there are many alternatives to this specific least squares approach. Notably, in [27], angular synchronization is solved with a method involving eigenvectors as well as one using semi-definite programming; these methods are shown to be robust to outliers. The eigenvector method is of particular interest here due to its ease of implementation and speed (relative to semi-definite programming). A summary of the key points of the eigenvector method [27] follows. For the n unknown angles $\theta_1, \dots, \theta_n$, noisy pairwise measurements $\delta_{kl} = \theta_k - \theta_l$, and corresponding graph edge set E , form an n -by- n matrix $\mathbf{H}$ with elements

$$H_{kl} = \begin{cases} e^{j\delta_{kl}} & k, l \in E \\ 0 & k, l \notin E \end{cases},$$

First, observe that $\mathbf{H}$ is Hermitian because $\delta_{kl} = -\delta_{lk}$. Then, consider the quadratic programming problem

$$\max_{\|z\|^2=n} \mathbf{z}^* \mathbf{H} \mathbf{z}$$

The maximizer $\mathbf{z}$ is given by the normalized top eigenvector (eigenvector with largest eigenvalue) of $\mathbf{H}$. For $z_i = e^{j\theta_i}$, note that for a good set of angles

$$e^{-j\theta_k} e^{j(\theta_k - \theta_l)} e^{j\theta_l} = 1 .$$

In contrast, for bad edges with uncorrelated error, this expression will mostly cancel out; thus for good edges the sum of ones will equal the total number of edges m_{good} , while for bad edges the total contribution will only be $O(\sqrt{m_{bad}})$. In short, the argument of top eigenvector of $\mathbf{H}$ contains all of the angles that are sought after.

The choice to use the more sensitive method of least squares versus a more robust method such as the eigenvector method should be done with a few considerations. One of the main these is the trade-off between accuracy and time-complexity. This trade-off was explored in the context of the demonstrated 4-by-4 array. A simulation comparing of the relative accuracy the least squares method versus the eigenvector method is given in Fig. 3.20. The RMS of 10,000 random angle sets was calculated for varying noise levels $\epsilon = [0, 180^\circ]$ for $\delta_{kl} \sim Uniform(0, \epsilon)$ and $\delta_{kl} \sim Normal(0, \epsilon)$.

For the demonstrated hardware systems with RMS phase error usually less than 10 degrees, the least-squares method tends to perform better. However, for larger arrays, the balance would likely change as m uncorrelated pairwise errors would accumulate in the predicted angles to $O(\sqrt{m})$. A time-complexity comparison of the two methods is harder to do. The dominant computation cost of the least-squares method would be in solving the least squares problem; for an R by C array, the $(2RC - R - C)$ equations imply that the resultant time complexity will be bounded above by matrix-matrix multiplication (naively $O(n^3)$). The sparsity of the equations and their regular nature, however, suggest that the typical bound is likely lower. As for eigenvector method, while finding the top eigenvector as part of an eigensystem would also be $O(n^3)$, in practice, faster iterative methods exist such as the power iteration method costs only $O(n^2)$ per step (but the number of iterations needed to convergence is more challenging to determine). Anecdotally, naive implementation of both methods seems to favor the least-squares method speed-wise by an order of magnitude for $n = 16$; of course, an unoptimized implementation with such a small value of n is far from a definitive statement comparing the two methods. In the




Figure 3.20: Comparison of least-squares and eigenvector angular synchronization methods

proximal field phase sensors context, further investigation and adaptability for this topic is warranted.

3.9 First Generation Hardware

There were two main generations of hardware created to demonstrate an amplitude-only proximal field phase sensor system. The following section describes the first attempt's hardware, measurements, as well as its shortcomings and lessons learned.

Hardware Description

The first generation hardware was created with a few goals in mind. Foremost, it must be operable as a phased array. Second, it must allow for reconfiguration of its proximal sensors and radiating antennas, as well as be able to efficiently transfer sensor readings back to a test computer. Third, it should be dynamically tileable in both axes.

To this end, a board-level phased array system operating at 2.5 GHz was created with the intent of serving as a flexible, reconfigurable test platform. 2.5 GHz was chosen for its availability of CoTS components (increasing ease of prototyping and lowering cost). This phased array system consists of identical tiles, each of which

contain an antenna board and associated proximal sensor configuration that is driven by an electronics board behind it. The electronics board, in addition to performing beam-steering, also serves data acquisition and communications functions. Each tile is backed by a 3D-printed plastic support frame for ease of mechanical assembly. A single tile assembly is shown in Fig. 3.21.




Figure 3.21: Single tile assembly

Electronics Board

The electronics board, shown in Fig. 3.22 is a 4-layer FR4 printed circuit board that measures 6 centimeters square. For timing, it features a 78.125 MHz reference signal that is buffered via a CMOS clock buffer IC and distributed in daisy-chain fashion around the array. A sigma-delta fractional-N PLL-based frequency synthesis IC is used to generate the 2.5 GHz RF signal; the modulator provides the phase shifting functionality needed for beam-steering. An output power amplifier boosts the PLL signal level to appropriate levels. Finally, a micro-controller configures the PLL operation through its registers, reads back sensor values via its ADC, and relays information across the system.

The layout of the PLL circuit is based on that of Professor Williamstyer's second

generation origami tile described in detail in [31]. Key differences in this design include the usage side-facing connectors for cost-reduction over the previous design: MMCX connectors allow for reference distribution whilst ordinary 0.1 inch pin headers distribute the DC rail and serial communications lines. Additionally, a linear regulator is used instead of a buck converter due to reduced power consumption requirements. Finally, both to reduce cost and increase scalability, serial communications lines are not buffered with clock buffer ICs; instead, tile communication lines are daisy chained, and each tile communicates with its previous and next neighbor: as the test computer only needs to communicate with a single tile, signal integrity requirements are thus greatly reduced.

Although daisy chaining simplifies the hardware requirements, it adds additional software requirements that cannot be neglected. Despite being relatively simple compared to modern packet-switched networks, the resultant ring network topology from daisy chaining can be subject to pitfalls. Synchronization must be maintained for consistent and reliable communications; usage of binary semaphores can accomplish this task.




Figure 3.22: Single tile electronics board with labeled components

Antenna Board

The antenna board, shown in Fig. 3.23, is a 2-layer 1.5mm thick Rogers 4003C substrate board that is also 6 centimeters square. The RF board contains the transmitting patch antennas for the phased array system as well as 4 loop shaped proximal field sensors that each feed into individual commercial peak-detecting circuits. These sensors produce an analog voltage based on the input signal amplitude that is then sampled and digitized by the onboard ADC of the micro-controller on the electronics board. A snap-fit RF connector serves as a matched probe feed for the patch antenna and also allows for easy attachment and removal of the antenna and electronics boards.




Figure 3.23: Single tile antenna and proximal field sensor board, front and back

The design process of the antenna board first began with that of the antenna. A patch antenna was first length-tuned to the frequency of interest; subsequently, the width and probe feed offset were tuned to match to 50 Ohms. Once the patch was designed, the power detector IC datasheet was consulted to ensure incident sensor powers were within its detection range and that it was appropriately impedance matched. To this end, less than -20 dB of coupling between the patch and each proximal field sensor was targeted roughly corresponding to less than 1% parasitic power loss due to the sensors. This set the antenna drive strength requirement. In an effort to impedance match, the capacitive input impedance (at 2.5 GHz) of the sensor was paired with an inductive loop.

Measurements and Results

Measurement Setup

A diagram of the measurement setup is given in Fig. 3.24. Relevant radiating components are mounted in an far-field range system with an anechoic chamber for controlled measurement of the array pattern. A measurement computer oversees the system; this computer is responsible for commanding the micro-controllers to beam-steer and acquire sensor values, as well as reading back and storing these values. To generate a training set, the measurement computer first generates a set of random phase settings for the array (labels), then instructs the array to steer to these phase states and acquire the corresponding sensor values (features). After hardware and software averaging, each tile micro-controller responds with its result. The measurement computer then unpacks and saves each set of labels and features. The dataset generation process is primarily limited by the baud rate of the single UART serial line used for communication, which, at 460800 baud, takes under 20 minutes to generate a 100,000 point dataset for a 4 element phased array. To train, a separate training computer pre-processes the data and trains a desired model.




Figure 3.24: Measurement setup block diagram

An important measurement detail pertains to the daisy-chained system reference distribution. As the low-frequency reference is buffered across each tile, there is an additive delay that propagates across the system as shown in Fig. 3.25. Furthermore, the PLL chip that generates the RF tone has some phase delay. As a result, on initial

startup, the system's elements have unknown phase offsets relative to each other. As these offsets represent an arbitrary global phase to the proximal field phase sensors (it is indistinguishable from the effects of the phase shifter themselves), their relation to broadside must be determined by some mechanism. One way that this can be done is through a far field optimization that requires a single phaseless receiver as shown in [2].




Figure 3.25: Phase offsets due to PLL and reference buffering

Measurement Results

A two element array was first measured to serve as a proof-of-life for the system. A 2-layer, feed-forward neural net with single hidden layer as described in Section 3.7 was used to trivially fit to the sinusoidal interference pattern seen at the proximal field sensors.

Subsequently, a four element linear array was assembled as shown in Fig. 3.26. A 50,000 point dataset (with sensor values averaged 10 times in software with a 15-millisecond delay) was generated for the array, after which a model (also as described in Section 3.7) was fit and far-field beam patterns were taken. An overall RMS of 14.8 degrees was observed, with a detailed error histogram shown in Fig. 3.27. To demonstrate the model in action, the phased array was first externally beam-steered

without aid of the proximal field sensors. At each steering angle, the model's prediction of what the phases should be are recorded. Using individual element patterns, the predicted pattern based on the model's phases is superposed with the measured pattern in Fig. 3.28 for comparison. While the resultant predicted beam pattern is wider than measured due to phase error, the direction is largely correct.




Figure 3.26: Four element array measurement setup

Challenges

As there was a significant deviation between the measured training error and simulated training error, an investigation was conducted to find the source. This, paired with the fact that the system struggled to operate coherently with 4 elements without significant configuration of the PLL registers, and failed to operate with additional elements despite much effort, indicated significant hardware issues.

These hardware issues were eventually discovered to be injection pulling upon

Training Error Histogram for 50000 Point Dataset (4 Element Array)




Figure 3.27: Four element array training results

Predicted vs. Measured 4 by 1 Array Patterns (Normalized)




Figure 3.28: Four element array beam patterns

inspection of the output spectrum on a spectrum analyzer: the PLLs of each element were fighting each other. A few efforts were made to mitigate these issues. Foremost, upon observing that tile PLLs would lose lock in a time-dependent manner, a thermally dependent effect verified and mitigated by lowering each tile's PLL VCO

amplitude. Subsequently, to mitigate cross-coupling between adjacent tiles, copper tape was used to shield each element.

While these steps did reduce the symptoms of injection pulling, they were unfortunately unable to completely address faults in the hardware. This eventually led to the development of the second generation of demonstration hardware for the proximal field phase sensors.

3.10 Second Generation Hardware

The second generation hardware sought a platform for focusing on development of the proximal field phase sensors without distraction from the secondary, but show-stopping injection pulling issues of the first generation. To this end, reference distribution at RF was chosen instead of reference distribution at low frequency. Although the single oscillator of the new system removed concern for injection pulling between multiple oscillators, the stability of the amplifying chains of each array element were still of concern. To minimize the risk of instability, a programmable attenuator was placed between the output of the power amplifier and antenna to reduce coupling effects between neighboring active components. Finally, the new system's two RF boards were placed perpendicularly to each other to further minimize coupling pathways.

Note: The second generation hardware was later revised once for ease of reconfiguration with some minor hardware changes. These will be described herein as first revision changes.

Hardware Description

As with the first generation, the second generation features identical individual elements that allow for the array to be scaled up in size as needed. Each element also consists of an electronics board and antenna board. The electronics board serves similar roles: it shifts element phases and gain leveling through its programmable attenuator. Besides configuring the RF chain, the micro-controller each tile contains also performs serial communications through UART and sensor data acquisition directly via on-chip ADC (done indirectly via external sigma-delta ADC on the first revision).

Electronics Board

For each tile's electronics board, the RF reference first passes through a phase shifter, followed by a power amplifier, then, after the signal passes through a programmable attenuator, it finally reaches the antenna board. The phase shifter is a switch-based binary-weighted phase shifter, while the power amplifier is intended for WiFi/Zigbee applications (but does support continuous-wave operation). Fig. 3.29 depicts a labeled photo of a single tile's electronics board.




Figure 3.29: Single tile electronics board with labeled components

Antenna Board

The antenna boards of the second generation are very similar in form to those of the first generation, but feature two additional sensors for a total of six. All currently tested boards feature patch antennas with loop sensors; the choice of power detector circuit is also the same.

Measurements and Results

This section contains a description of the measurement setup and corresponding measurement results for different configurations of the second generation board




Figure 3.30: Single tile antenna and proximal field sensor board

level proximal field phase sensor system. As phase determination for a 2 by 2 array is trivial, a 2 by 4 array at 2.5 GHz was the first constructed as shown in Figure 3.31.




Figure 3.31: 2 by 4 phased array with circular loop proximal field sensors

Training data was generated by uniformly sampling phase shifter settings on each tile and recording corresponding sensor outputs. Done over a single 0.5M baud UART serial link, this dataset generation takes approximately 30 minutes to generate a 100,000 point dataset of all 48 sensors. Element 0 (Fig. 3.32) is used as a reference element for the array, with its phase shifter fixed to zero. With the training dataset comprising the sensor readouts and labeled phase data, a convolutional neural network was trained to predict each element's phase difference relative to its top and

left neighbors. The training time for the network was 50 minutes on a single Nvidia RTX 4090 GPU. The model architecture for the network is shown in the right side of Fig. 3.32; ReLU activation was used in between layers and the Adam optimizer was used. A short Bayesian hyperparameter optimization was done that for selection of the number of filters and sizes of the fully connected layers.

To recover phases referenced to element 0 from the neural network's pairwise neighboring phase predictions (i.e. solving the angular synchronization problem), two methods were used. The first method was a naive method utilizing either the top or left prediction of each element to construct a tree of relative phases. The second was least squares regression accounting for phase wrapping as previously mentioned. The results of the two methods are shown in Fig. 3.32. The total phase error (excluding the trivial zeroth element) for the more accurate second method was approximately 5.3 degrees RMS.




Figure 3.32: 2 by 4 phased array element labeling (top left), RMS phase error heatmap for both methods (bottom left), and neural network block diagram (right)

Next a demonstration to verify the beam-steering capabilities of the neural network was done by first using tuning each element to phase gradients corresponding

to $\pm 15^\circ$, $\pm 30^\circ$, and $\pm 45^\circ$ steering angles. Neural-network assisted steering was accomplished by using neural network inferences on sensor readouts to tune phase settings with respect to the broadside calibration state. The results of inference assisted beam-steering compared to an beamsteering done via far-field optimization are shown in Fig. 3.33. The neural network's inferences allow for beam-steering angles that are very close to the desired value; the phase error of the inferences generally results in a slight degradation of the side-lobe level.




Figure 3.33: Array far-field steering pattern measurements using far-field optimization (solid) and phase sensor feedback (dotted)

Following the measurements of the 2 by 4 array, a lower half was added to the existing array to complete a full 4 by 4 array. The new array's training data generation remained the same, albeit with data generation halving in speed for each data point due to continued usage of a single serial link. Initial training of the system on the previously described convolutional neural network proved challenging. In an effort to optimize the hyperparameters of the network, some limited hyperparameter optimization was performed via a combination of random and Bayesian optimization methods. A geometric stepped learning rate was also found to improve training loss. For the network shown in Fig. 3.34, the pairwise RMS (typically a few degrees lower than the total RMS) within rows was found to be 8.48° while within columns the value was found to be 9.65° . The larger error found across columns is consistent with fewer sensors being close to the pairwise elements vertically versus

horizontally.




Figure 3.34: Neural network architecture initially used to train 4 by 4 array

While significant effort was expended in obtaining networks with better inference performance, the ca. 9° pairwise error figure when using a 4-layer convolutional neural network on the 4 by 4 array was significantly higher than the ca. 3.4° error on the 2 by 4 array. For this reason, the decision was made to switch to training on a deep neural network in line with the existing usage [24] of deep networks in solving similar phase retrieval problems. Initial training was done with a ResNet architecture with 10 convolutional blocks as shown in Fig. 3.35 with the Adam optimizer and cosine annealing. This model performed best among 7 other models in a random hyperparameter search with a pairwise test error of 5.54° as shown in Fig. 3.36.




Figure 3.35: Initial 10 ResNet-10 model used to train 4 by 4 array

This result, a significant forward step towards the 3.4° pairwise error of the 2 by 4 array, suggested that usage of deeper nets might provide a straightforward way for increased inference performance. To test this hypothesis, 3 additional ResNet models were trained with increasing numbers of blocks: 20, 24, and 48. A summary of the results is given in Table 3.3. The training time refers to the time the network took to train on a consumer RTX 5090 GPU, while the number of multiply-accumulate (MAC) operations are given as a reference for inference expense.




Figure 3.36: Training loss for 8 different ResNet10 models

# of blocks	Pairwise Test Err [deg.]	Training Time [min.]	MACs [M]	# of Params. [M]
10	5.54	112	87.7	1.45
20	4.76	280	140	2.19
24	4.59	352	190	2.58
48	4.24	765	365	4.44

Table 3.3: Summary of training results for various ResNet block counts

For practical use, some trade-off could be made between inference cost and inference accuracy; smaller, lower accuracy networks could be used when accuracy is less of a concern, while larger, more expensive networks could be used for applications desiring greater accuracy. Of note additionally is the fact that inferences need not occur continuously; if done at lower duty cycle corresponding to the time constants that correspond to decoherence of the array (e.g., a day-night cycle resulting in temperature deviations for a low-earth orbit satellite array might only require a handful of calibrations per 24-hour cycle).

For the 4 by 4 array, the overall test error for the best performing ResNet48 model, calculated with least-squares regression for angular synchronization, is 5.57° RMS: a figure slightly higher than the 5.3° RMS of the 2 by 4 array, albeit at the significant computational cost increase of a 365 million MAC operations versus the 6 million required for the original array. If these results could be matched in a large array (segmented to use this method locally for timing/synchronization), the corresponding less than 6 degrees of error would result in less than 1% array gain degradation per Fig. 3.1; of course, significant additional effort would be needed to verify proximal field sensor operation on a larger array, synchronize segments, and scale up the neural network.

3.11 Conclusion and Future Work

By extending proximal field sensors to work with amplitude only, the work in this chapter seeks to lower the development cost and design difficulty of using proximal field sensors, which possess a distinct advantage in providing localized feedback that would hopefully (largely) obviate the need for full array characterization in a near- or far-field range. The information from the amplitude-only sensors generates a phase retrieval problem, of which deep learning appears to be a promising method for data interpretation.

There are many avenues for future work. Currently, measurements are in progress to determine model inference performance when using various information theory metrics for sensor placements. Beyond that, scaling to larger arrays, using sensors for additional information (such as radiator amplitude, or even potentially array shape configuration in flexible arrays), and applying this work at higher frequencies for on-chip radiators are all excellent directions to further explore proximal field sensors.

Chapter 4

MICROWAVE ARRAY FOR POWER-TRANSFER LOW-ORBIT EXPERIMENT

4.1 Introduction

Wireless power transfer (WPT) at distance could transform the nature of energy distribution and access, both on Earth and in space¹. 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. In particular, the Caltech Holistic Integrated Circuits Lab has been focused on using phased arrays for WPT. A key advantage of phased arrays is the ability to adjust the relative timing of the individual elements to coherently focus transmitted energy, allowing energy to 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 solar power (SSP). First conceived by Isaac Asimov in his short story, "Reason" [33], an SSP system is described where solar energy is converted to microwaves in space before being transmitted to Earth.

Unlike many existing renewable energy sources such as solar, wind, geothermal, and hydroelectricity, SSP has the unique advantage of providing continuous, stable, and predictable energy [34] without imposing storage demands. Moreover, unlike traditional existing power generation schemes, electronic architectures also enable dynamically dispatching power to nearly any desired location on Earth, including

¹Florian Bohn, Behrooz Abiri, Matan Gal-Katziri, Mohammed Reza Hashemi, Austin Fikes, Oren Mizrahi, Alex Ayling, Mohith Manohara, Jesse Brunet, and Raha Riazati worked on the Caltech Space Solar Power Project and contributed to the concepts, prototypes, and measurements that appear in the MAPLE payload. Additionally, system level concepts were developed with the Atwater Lab and Pellegrino lab. Richard Madonna, Damon Russel, and Sharmila Padmanabhan aided as advisors to MAPLE's development. This chapter is adapted from [32].

locations where the capacity or the infrastructure to satisfy the demand for electricity may not be available.

Caltech's vision for SSP (Fig. 4.1) comprises the deployment of massive kilometer-scale arrays of photovoltaic (PV) cells in space [35]. 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 Azimov, 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 SSP due to lower atmospheric attenuation [36], 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 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 [36]. 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% [37].

The works in [36], [39] provide excellent documentation for the significant work that has been done on SSP over the last 50+ years. This prior work can be roughly categorized into SSP system concepts, component design, and system demonstrations.

SSP system concepts presented in literature entail high-level theoretical design and analysis of a full-scale SSP system [40]–[49]. These are perhaps best exemplified by the NASA/DOE reference system [42] study done in 1979, the first majorly funded study of SSP's feasibility. Over the years, system concepts have evolved with technology—examples of subsequent studies done in the mid-1990s and early 2000s include ones done by NASA [43], ESA [44], and JAXA [45], [46]. More recently, SSP concept proposals are enjoying an accelerating push towards a more renewable future [47]–[49].

Component design literature discusses designing and testing of the specific electrical, optical, and mechanical components necessary to enable an SSP system. Examples include RF power amplifiers (PAs) [50], [51], antennas [52]–[54], recten-




Figure 4.1: The Caltech SSP concept [38]: A photovoltaic-wireless power transfer array converts solar energy to DC power and DC power to RF power which is then collected on Earth by a ground station.

nas [37], [55], [56], photovoltaic cells [57], space structures [58], etc.

Published system demonstrations present fully functional, end-to-end wireless power transfer systems [59]–[65]. Well-known early demonstrations include the Raytheon helicopter [59] and the Venus Site experiment at Jet Propulsion Laboratory’s Goldstone facility in 1975 [61] (which delivered over 30 kW of power at a 1.54 km distance). JAXA has explored nonlinear plasma-wave excitations that would arise from high-power microwave beams [63], [66], [67]. In 1992, the ISY-METS mission [64] demonstrated the first sub-orbital WPT system.

It would be imprudent to omit the fact that despite SSP’s potential benefits and a huge body of existing research, practical implementations of any SSP system do not exist to date due to considerable challenges: one such challenge is the immense launch costs of deploying a large SSP system in space. To this end, the Caltech Space Solar Power Project (SSPP) has prioritized mass (and thereby cost) reduction of SSP systems; With this in mind, Caltech SSPP has devoted research groups to investigate the three main subsystems of SSP: wireless power transfer, photovoltaics, and deployable space structures. The proposed WPT subsystem employs lightweight, flexible phased arrays containing radio frequency integrated circuits (RFICs) minimizing mass and size by combining onto a millimeter-scale

chip all necessary digital and analog circuitry for DC to RF power conversion and focusing. Power would be supplied by thin-film photovoltaic cells (on one, or both sides of the structure); the entire structure would be stow-able and deployable in modular pieces.

Development of the Caltech Space Solar Power Project (SSPP) has spanned over a decade; recently, such flexible, lightweight, coherent arrays powered by custom RFICs have been developed and tested on Earth [52], [53], [68]–[72]. However, to the best of the author’s knowledge, a dedicated free-space WPT system using this or any other technology has not been deployed and demonstrated in space—prior to the launch of the Caltech Space Solar Power Demonstration One [73], which marks an important milestone by showcasing SSPP’s key concept technologies operating in space. In particular, for WPT, The **M**icrowave **A**rray for **P**ower Transfer, **L**ow **E**arth Orbit Experiment (MAPLE) conducts long-term experiments demonstrating WPT in Earth’s orbit for the first time. MAPLE not only verifies the ability for flexible, lightweight, coherent arrays to transmit energy in space; it evaluates the strengths and weaknesses of the current state of the art technology. The results of the mission will serve to inform the next phase(s) of development for these arrays.

4.2 Payload Hardware

Overview

A labeled CAD rendering of MAPLE is shown in Fig. 4.3; a flight model (FM) photo of MAPLE is depicted in Fig. 4.4. As detailed in Section 4.4, Thermal and mechanical stability of MAPLE were key design factors for the mission: both to survive launch, as well as endure months of thermal cycling during orbit. For the sake of maximizing the possibility of mission success, redundancy was built into multiple levels of MAPLE; implemented in hardware is a combination of dual modular redundancy (when appropriate) and fail-soft redundancy (when more limited in cost/design feasibility). A block diagram illustrating the some of the nature of the redundancy is given in Fig. 4.2. Most mission goals could be achieved with functionality of only five of the eleven components blocks.

MAPLE’s key payload hardware that enables wireless power transfer experiments in space includes:

- Two custom-designed 16-channel silicon radio-frequency integrated circuits (RFICS) handling frequency synthesis, power amplification, and beam-forming via phase control




Figure 4.2: MAPLE system block diagram (flight computer is external to MAPLE, but has redundancy in parallel connections to both motherboards of the aggregation panel)

- A four-by-eight rectangular array of 32 flexible co-cured pop-up dipole antennas
- A flexible polyimide printed circuit board accommodates the RFICs and antennas ("flex array" or FA)
- A rigid PCB providing (with redundancy) voltage regulation, on-board telemetry, and system control
- Two rectenna arrays that can switch loads to either resistor or LED
- Two narrow field-of-view cameras, controlled by Raspberry Pi single board computers
- Two wide field-of-view cameras, also controlled by Raspberry Pi single board computers
- A transparent (both to visible light and the frequency of operation) sapphire window opposite to the flexible PCB

Custom radio-frequency integrated circuits (RFICs)

The custom RFIC fulfills the core functionality of the wireless power transfer array. Designed in a 65 nm CMOS process by previous efforts in the CHIC lab [51][52],




Figure 4.3: Labeled CAD render of MAPLE




Figure 4.4: MAPLE flight model (FM)

for each of the sixteen channels, the RFIC accomplishes frequency synthesis, timing for beam-forming, and DC to RF power conversion. By utilizing modern CMOS scaling, this complex system is integrated into a small, extremely low-mass 2.8×2.8 mm die possessing a footprint unattainable using discrete components. MAPLE contains two of these RFICs that drive the respective left- and right-half of the four-by-eight element phased array.

A simplified block diagram and chip micrograph of the RFIC are depicted in Fig. 4.5.

Reference distribution is accomplished by splitting the output of a 48 MHz discrete crystal oscillator to both RFICs. On chip, a single, centralized phase locked loop (PLL) operates with a x52 multiplication ratio to synthesize an intermediate 2.496 GHz signal distributed to four quadrants each comprising 4 cores (corresponding to a total of 16 channels). Each core contains an RF chain consisting of a second-stage fixed x4 multiplication ratio clock multiplier unit (CMU) synthesizing a 9.984 GHz signal that exits the chip after buffering and amplification. Integrated baluns provide an efficient transition from differential PA outputs to single-ended 50 Ω microstrip transmission lines feeding popup dipole antennas.




Figure 4.5: Simplified RFIC block diagram and die micrograph

Co-cured popup antennas

MAPLE uses flexible and lightweight co-cured popup antennas discussed in [74][75]. Antennas used in this application require flexibility and low-mass to enable efficient stowage for launch; they must also endure the harsh space environment. Coupled with the geometry dependent nature of antenna design, these goals suggest contradictory implementations. This is reconciled via integration of the electronic and structural layers through the "co-curing" described in further detail in [74][75]. By aligning, laminating, molding, and curing all layers together, mechanical strength with low-mass is achieved without increasing electromagnetic complexity through additional adhesives or mounts.

Flex Array (FA)

A flexible, lightweight, four-layer polyimide PCB serves as the substrate for the RFICs and popup antennas. This entire assembly is referred to as the "flex array" or FA. The flexibility of the FA [69][70] is intended to provide a path for eventual

stowage of a fully-fledged space solar power system; however, as a starting point for MAPLE, the FA is mounted to a rigid frame (Fig. 4.6). Roughly 16 cm in width and 9 cm in height, the single physical FA is logically into two tiles each with 16 antennas in a four-by-four grid. The antennas in each tile are driven by a single RFIC. By delegating digital control of each RFIC to separate micro-controllers, the full-health phased array can operate as a single four-by-eight array or fail-soft as separate four-by-four arrays. As the FA components require its PCB routing to contain a mix of RF, analog, and digital traces, careful layout was especially imperative. The four metal layers were utilized as follows:

1. Layer 1 provides pads for RFICs, micro-controllers, associated passives, and contains all connectors
2. Layer 2 distributes power for the multiple rails (3.3V, 1.8V, and 1.2V) required by the RFIC and micro-controller
3. Layer 3 has a reduced mass ground plane (existing in the form of a rectangular grid) for antennas and RF traces on layer 4.
4. Layer 4 uses 50 Ω microstrip transmission lines to route RFIC RF outputs to antenna inputs

Aggregation panel

Rectennas

MAPLE contains two rectenna arrays PCBs that receive and convert the flex-array's wirelessly transmitted RF power to DC. Inclusion of two boards provides both fail-soft redundancy as well as spatial diversity that enables demonstration of the transmitter's rapid beam-steering capabilities: one rectenna array PCB is positioned broadside the array, while the other PCB is offset and mounted to a sidewall. The rectenna array PCBs are four-layer multi-substrate boards utilizing Rogers 4003C dielectric for their antennas and Rogers 4350B dielectric for the rectifier circuits (containing a standard FR4 pre-impregnated layer in the middle).

Each rectenna array PCBs consists of a 5 x 5 grid of rectenna elements that combine their power in DC in parallel configuration; two maximum power point tracking boost circuits (each driving a white LED) and two 22 Ω DC loads provide redundancy at the MB level. Onboard MOSFETs enable switching generated DC power between the boost circuit powering the LED, and the 22 Ω load, which is both a stable,




Figure 4.6: Top: Front side of flexible phased array (FA) attached to aluminum frame Bottom: Back of FA with vertical bars to heat sink RFIC (translucency for visualization) Labels: (a) popup antennas (b) aluminum frame (c) RFIC 1 (obscured by frame) (d) microcontroller (e) RFIC 2 (soldered on interposer)

power-optimal load for optimization as well as a protective load preventing damage from open circuit rectenna operation. The front and back of a rectenna is shown in Fig. 4.7; the circuit schematic of a rectifying element is shown in Fig. 4.8; an rectifier efficiency plot is shown in Fig. 4.9.

Each rectenna element contains a dual polarized, probe-fed patch antenna that feeds into two identical rectifier circuits (designed based on [2]). For a nominal rectified power of 210 mW, the rectenna array has a conversion efficiency of 40% (with respect to power incident on its aperture). To note is that system operation for maximum rectenna conversion efficiency was not prioritized for MAPLE; as seen




Figure 4.7: 5×5 rectenna array of dual-polarized patch antenna receivers.




Figure 4.8: Rectifier circuit schematic with labeled harmonic terminations

in [76], X-band rectifiers can possess conversion efficiencies of around 70% under optimal load and incident RF power settings.

Cameras

Four cameras redundant in viewing perspective (2 narrow FOV and 2 wide FOV) are included on MAPLE. Wide FOV cameras are Raspberry Pi High Quality Cameras with $\approx 170^\circ$ FOV fisheye lenses attached and are located on MAPLE's sidewall; narrow FOV cameras are Raspberry Pi Camera Modules 2 and are located adjacent to the flex array. Each individual camera is connected to a Raspberry Pi Zero W via flexible printed circuit connectors. The Raspberry Pis manage camera operations and photo storage via onboard SD cards, as well as handle photo transfer to the AP.




Figure 4.9: Rectenna efficiency as a function of absolute output power.

This arrangement for the cameras was chosen as the micro-controllers that manage MAPLE do not possess sufficient computational power to manage imaging devices and process images; this allowed MAPLE’s software development to remain straightforward and reliable for risk minimization while providing a means for photos to be taken. A UART serial line provides the control and data transfer mechanism between Raspberry Pi and AP. A finite state machine implemented via Python script synchronizes operations. Photos are stored in a circular buffer alongside their thumbnails allowing previews while minimizing data downlink usage. For the flight computer to receive a photo, a link is established between the flight computer and Raspberry Pi (with MAPLE’s micro-controller serving as a middle man) where photos are transmitted in small (≈ 128) byte packets with error detection implemented via a checksum mechanism similar to that in UDP.

4.3 Software

Focusing algorithms

To focus the FA power onto the rectennas, MAPLE uses an implementation of the algorithm presented in [2]. This algorithm is designed to bypass the trade-off between accuracy and runtime that typically plagues optimizations of extremely high-dimensional systems. In the context of MAPLE, each of the 32-elements have a 9-bit digital phase shifter, implying a total of $(2^9)^{32} \approx 10^{87}$ unique phase configurations. This is done by sweeping over successively decreasing ranges of phases of a group of a large number of elements. By sweeping a group of elements, changes exceed the noise floor of the system. Groups are formed by bit-masking the

phases of the full array; the choice of orthogonal masks varies different combinations of elements in different sweeps, thus increasing the search space explored in few iterations.

The focusing algorithm utilizes the closed-loop feedback configuration of a "generating unit" (GU) and "receiving unit" (RU), as shown in Fig. 5 of [2]. For MAPLE, the flexible phased array (FA) serves as the "GU array". MAPLE's GU is formed when including the micro-controller on the MB that executes the focusing algorithm and reads back voltages from the receiver array, the RU, which is implemented by the aforementioned rectenna. Whilst the communication for the GU and RU was presented as using a wireless link in [2], MAPLE features a wired link for simplicity and reliability.

In MAPLE, the optimization algorithm starts by initializing the RFIC via programming of its registers (this is done through the local flex array micro-controllers). Subsequently, the MB requests adjustments to selected phases, and detects the resultant change at the receiver. Updates to the phases are based on the change in received power; the algorithm proceeds in such a manner for a different combination of elements until a specified number of iterations have finished running.

Running the focusing algorithm on a functional system produces curves relating rectified power and elapsed time as depicted in Fig. 4.15—rectified power increases rapidly before stabilizing (presumably to a maximum).

MAPLE also has the ability to upload an array of digital codes from the flight computer to the flex array: "recalling" a prior phase state. This recall allows MAPLE to measure rectified power at the receiver that were saved in nonvolatile program memory (e.g., to compare performance on Earth vs. in orbit) as well as in volatile memory (e.g., to seed an optimization with the result of a prior one). The former allows MAPLE to make direct comparisons of system performance (maintaining identical FA programming) while the latter provides some indication of how effective a single optimization is in the global space.

One final remark regarding the focusing algorithm is that it functions without direct measurement of the absolute phase delay of each element. This phase delay, the combination of the fixed phase delay (such as delay due to clock and RF distribution) and the effect of the phase shifter is not measured; instead, the optimization algorithm works to modify phase shifter settings attempting to maximize power (and thus phase coherence) at the receiver.

4.4 Mechanical and Thermal Design

Spacecraft interface

MAPLE is mounted to the deck of the spacecraft Vigoride-5 built by Momentus Inc. as shown in Fig. 4.10. Vigoride-5 provides power, communications, thermal sinking, and attitude control for Space Solar Power Demonstration One (SSPD-1) [73], Caltech's first in-orbit space solar power demonstration. SSPD-1 includes MAPLE for wireless power transfer, as well as ALBA for photovoltaics and DOLCE for space structures.




Figure 4.10: SSPD-1 integrated on Vigoride-5 spacecraft (a) MAPLE (wireless power transfer experiment) (b) flight computer (c) ALBA (photovoltaics experiment) (d) DOLCE (dynamic, deployable structures experiment)

Mechanical and Thermal Design

MAPLE is contained within a standard 6U aluminum CubeSat chassis. Within the frame contains the entirety of MAPLE's electronics. A single Micro-D connector mounted to a fixed back-panel of MAPLE serves as the connection point for externally supplied power and flight computer communications. Aluminum panels

provide a rigid backing to mount cameras, rectennas, and a sapphire covered (to mitigate payload damage from atomic oxygen) viewing port. The 4.4mm thickness of the sapphire window ($\lambda/2$ at 10 GHz) was chosen to minimize reflections at the air-sapphire interface, allowing it to be transparent optically and at the frequency of interest.

All components of MAPLE are eventually thermally coupled to the deck of Vigoride-5. Effectively, due to its large thermal capacitance relative to MAPLE the deck serves as an isothermal boundary. Though many components dissipate heat more indirectly (through the aggregation panel ground plane connections to the CubeSat frame), more concentrated sources of heat are addressed separately. Side panel Raspberry Pis have their main chip thermally strapped via copper strips and epoxy. Because the RFICs were the most concentrated sources of heat, special heat-sinking was designed to prevent RFIC temperatures from exceeding their 80°C maximum operating temperature.

Fig. 4.11 shows the a simulation of the RFIC heat-sinking solution. An aluminum frame providing mechanical support and thermal relief sits between the FA and AP: above the back of the FA where the two RFICs are located is two pedestals. The RFICs are bonded to these pedestals via Stycast 2850 thermal epoxy; thus, heat rejection from the RFICs primarily occurs through this aluminum frame. For a worst-case deck thermal boundary of 50°, simulation indicates a maximum temperature of around 74°C at the RFICs.

Beyond conductive thermal transfer, MAPLE is enclosed by a 12-layer blanket of multi-layer insulation (MLI) shielding it from direct sun exposure and preventing excessive radiative cooling. The MLI consists of outer layers made from aluminized polyimide and inner layers made of low emissivity aluminum blankets (containing perforations for air egress during launch).

4.5 Qualification Testing

MAPLE underwent radiation, mechanical, and thermal vacuum chamber testing. A brief description of the contextual purpose and results of the testing follow.

Radiation Testing

The effects of radiation exposure were of significant concern during the development of MAPLE: particularly those resulting from single event effects (SEE). Analysis from SPENVIS indicated that cumulative damage from total ionizing dose (TID)




Figure 4.11: RFIC thermal simulation with heat sink stackup

should not be of significant concern due to MAPLE's short duration LEO orbit [77]. As the RFIC central to MAPLE's operation had not yet been tested for SEE, limited SEE proton testing at Texas A & M University's (TAMU) radiation effects facility served as a preliminary indicator for mission feasibility, given that at the time there was effectively no alternative replacement for the custom RFIC. Some key additional electronics components such as the micro-controllers used for MAPLE were also tested, though lists of parts with already tested and, preferably with flight heritage, such as those provided by JPL[78][79], were preferred. Tested parts (typically 5 copies of each component) were subjected to a 45 MeV proton beam for either one or three minutes with a fluence of 10^{10} protons/cm²; no damage was observed. Fig. 4.12 depicts the radiation test setup used at TAMU. Digital relays provide power supply isolation in event of latch-up events; automation of the test setup was imperative in ensuring smooth and efficient testing.

Mechanical Vibration Testing

MAPLE was required by the launch provider to undergo mechanical vibration testing with levels and spectra in accordance with the spacecraft's maximum predicted environments. Serving as a readiness indicator prior to this, MAPLE's components were first shake-tested on a smaller shake table at the Graduate Aerospace Laboratory of Caltech; then, the full EM of MAPLE was tested. Subsequent testing (at the third-party Exporior Laboratories) subjected MAPLE's FM to four types of testing in all three axes:

1. Sine vibration: A single tone sweep over a frequency range.




Figure 4.12: TAMU proton beam with mounted MAPLE component test panel

2. Random vibration: Excitation via a waveform specified in terms of power spectral density.
3. Sine burst: A single low frequency, high amplitude tone intended to simulate high displacement vibration. For MAPLE, 15 cycles of a 15 Hz sine waveform was applied.
4. Shock: Sudden, large acceleration and low duration test simulating sudden mechanical transients such as collisions.

Fig. 4.13 depicts MAPLE on a shake testing table at Experior Laboratories and testing results from a sine vibration test indicating a presence of a mechanical resonant mode associated with the top of MAPLE's frame that was subsequently fixed. Following a final round of testing, MAPLE passed all testing without incurring changes to electrical performance or significant variations to its mechanical spectrum.

Thermal Vacuum Chamber (TVAC) Testing

Terrestrial demonstrations of wireless power transfer to date had benefited from natural convection to cool the RFIC. As MAPLE's configuration only allows for heat rejection to the deck of the launch provider, TVAC testing simulation the harsh thermal environment of space was critical to verify MAPLE's thermal design.




Figure 4.13: Left: MAPLE mounted on vibration table on its X axis Right: Sine vibration excitation and response

Similar to mechanical testing, TVAC testing occurred in stages over MAPLE's development—initially, components of MAPLE such as its camera assemblies and flex array were tested, following successful testing; the entire payload was tested. Testing of the full payload was done at Exporior Laboratories where MAPLE was placed in a chamber pumped to a high vacuum of less than 10^{-6} torr. The chamber's temperature was controlled subjecting MAPLE to both moderate and extreme temperature limits expected during its mission. These included:

- Predicted (power-off) temperature limits: $[-5, 15]^{\circ}\text{C}$
- Protoflight temperature limits: $[-10, 30]^{\circ}\text{C}$
- Extreme temperature limits: $[-30, 60]^{\circ}\text{C}$

TVAC testing for the entirety of MAPLE was done with the follow steps. First, MAPLE, powered off, was first soaked at the extreme minimum/maximum limits. Then, MAPLE was thermal cycled (95 minutes per period) at the protoflight temperature limits during which the payload was operated with typical command sequences. Finally, following soaks at the extreme temperature limits, MAPLE was brought to room temperature for 100 minutes. Visual inspection and electrical testing indicated that the payload survived TVAC testing. Fig. 4.14 depicts the MAPLE EM mounted for TVAC testing at JPL post-mission for additional thermal behavior studies.

4.6 In-Orbit Experiments

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° [80]. After initial qualification testing of Vigoride-5 (V5) 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



Figure 4.14: MAPLE EM mounted for TVAC testing at JPL post-mission

eight months.

4.7 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 orbit, to the broadside and sidewall rectennas, are pictured in Fig. 4.16 and 4.15, respectively. The curves correspond to the optimizations that resulted in the highest rectified power to the respective rectennas during the duration of the mission.

The system began with no programmed phase shift across the array and, thus, a random absolute phase distribution at the output. As the optimization improved coherence, the rectified power increased until it reaches a maximum: a process which took $\approx 7s$ as seen in Fig. 4.15 and 4.16. At these “maximum power points,” the flex array created coherent focal points on each of the rectennas.

Note that this relatively long stabilization time is a consequence of the design implementation. The optimization algorithm is run on one microcontroller in the motherboard (MB) to generate phase codes which are sent to another microcon-




Figure 4.15: Rectenna #1 (sidewall). Data from 2023/03/14, 03:19 UTC.




Figure 4.16: Rectenna #2 (broadside). Data from 2023/03/03, 23:51 UTC.

troller which translates the digital phase codes into commands sent to the RFIC. After the RFIC updates its phase shifter(s), the MB reads the voltage off the rectenna and reports it to the flight computer for data collection. Because the flight computer is simultaneously logging many streams of data, it adds further delay. This redundant implementation was chosen to maximize mission success even in the case of individual part failure; it also expands optimization runtime.

The peak rectified DC powers received by the the sidewall and broadside rectennas were 231mW and 251mW, respectively. Each of the 16 stacks, composed of two power amplifier (PA) cores (Fig. 4.5), 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. 4.17, demonstrates the magnitude of power transferred.

As a percentage of the total RFIC DC power consumption (6.85W, as discussed in Section 2.2), peak rectified powers of 231mW and 251mW correspond to end-to-end link efficiencies of $\eta_{\text{link}} = 3.37\%$ and $\eta_{\text{link}} = 3.66\%$ for the sidewall and broadside

rectennas, respectively. This is a decrease relative to the expected link efficiency ($\eta_{\text{link}} = 6.26\%$) calculated above, based on individual component efficiencies. The discrepancy can be attributed, in part, to a combination of the following possibilities: degraded PA efficiency, nonlinear phase rotator coupling preventing complete coherence, antenna mismatch, or added losses from interposer RF line fan-out. Current measured efficiencies are not representative of the expected efficiency of the currently viable optimized system, predicted to be around 30%, as shown in Section 2.2.

Power degraded towards the end of the focusing in Fig. 4.15 and 4.16. This result is a deviation from ground testing. Separate in-orbit measurements focusing three times back-to-back using only 16 elements controlled by RFIC 1 (Tile 1) (Fig. 4.18) and RFIC 2 (Tile 2) (Fig. 4.19) were conducted. The results indicate degradation in the recovered power from Tile 2 that is exacerbated in subsequent optimizations. Tile 1, in comparison, output power consistent with ground tests and remained stable over time. Additionally, the 32 element power levels are restored once the system is restarted from an off state.




Figure 4.17: 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.

Degradation in RFIC 2 and associated decoherence with elements driven by RFIC 1 is the most likely cause of the drop in power seen in the full 32-element focusing

experiments (Figs. 4.15 and 4.16.) More discussion of this phenomenon is presented in Section 4.9.




Figure 4.18: Rectenna #2 focusing with only 16 elements controlled by RFIC 1. Data from 2023/03/21, 23:16 UTC.




Figure 4.19: Rectenna #2 focusing with only 16 elements controlled by RFIC 2. Data from 2023/03/22, 00:47 UTC.

The total experiment duration is the product of the number of iterations and the iteration loop delay. In the case of the experiments shown in Fig. 4.15 and 4.16, this duration was 16s.

Individual Cores: Phase Shifting Performance

MAPLE is equipped with a function which generates two-core interference patterns at the receiver for the characterization of individual core performance. A core refers to the “unit cell” of a clock multiplier unit (CMU), buffers, a PA, and output balun *at the chip level*. Cores are paired in “stacks” of two cores according to the series connection of their respective PAs. Each RFIC has 16 cores; MAPLE has two RFICS, or 32 cores total.

Testing of each of MAPLE's 32 cores proceeds as follows: addressing a single stack, one core is turned on with variable phase while the other is turned on with constant phase. The phase of the first core is swept digitally through some subset of the entire programmable range, set by (effectively) 9 bits of control.

Phases are programmed using two bytes: a sign byte which is either all '1' or all '0' and a data byte. The sign byte only expresses an additional bit of information.

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. This is repeated for each of the 16 stacks across the two radio-frequency integrated circuits (RFICs) on MAPLE. An example of raw rectified power from one of these tests on Earth is presented in Fig. 4.20, in black. Data collected from tests in orbit are overlaid in red.




Figure 4.20: 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.

Assuming a simplified, far-field propagation model and assuming a single phase center for the receiver, the electric field at the receiver during a phase sweep of core

m against constant phase on core n can be modeled as follows:

$$\vec{E}_m(t, c) = e^{j\omega t} \left(A_m e^{j\phi_m(c)} P_m e^{j\phi_m^P} + A_n e^{j\phi_n(0)} P_n e^{j\phi_n^P} \right) \quad (4.1)$$

where ω is the wave frequency, $\phi_m(c)$ is core m 's clock multiplier unit (CMU) phase function for code c , $n \in \{m-1, m+1\}$ is the index of the other core in the stack, A_m is the amplitude of core m 's outgoing wave, and $P_m e^{j\phi_m^P}$ is the propagation vector from core m to the receiver (with amplitude P_m and phase ϕ_m^P). The propagation vector accounts for an "infinite" sum of reflections emanating from the core which can vary owing (mostly) to perturbations on the multi-layer insulation blanket and possibly other variable, internal physical features. This vector can vary, but is quasi-static for the measurement duration and likely changes negligibly in orbit where there is no wind or (expected) physical change to the payload. The magnitude of the electric field can be expressed more simply as:

$$\left| \vec{E}_m(t, c) \right| = \left| A'_m e^{j\phi'_m(c)} + A'_n \right| \quad (4.2)$$

where $A'_m = A_m P_m$ and $\phi'_m(c) = \phi_m(c) - \theta$ is the CMU phase function for core m , less a sum of constant phase terms, θ .

When rectified, each phase sweep should produce a pseudo-sinusoidal interference pattern at the receiver. This interference pattern can be fit to the following model:

$$V_m(c) = \sqrt{\frac{A_m'^2 + A_n'^2 + 2A'_m A'_n \cos(\phi_m(c) + \theta)}{2}} \quad (4.3)$$

where $V_m(c)$ is the voltage read by the receiver when core m is programmed to phase code c .

This model, $V_m(c)$, suggests a symmetry between the patterns generated by two cores in a stack: both patterns should be the same up to a constant phase shift and minor differences in the CMU phase function. This symmetry also prevents the determination of which core is m and which is n . By convention, we assign the larger of $\{A'_m, A'_n\}$ to be A'_1 and the smaller to be A'_2 . We cannot distinguish between the two amplitudes in a stack enough to assign them to their respective cores.

This model allows for the prediction of A'_m , A'_n , and extraction of the CMU function. Both the amplitudes and the CMU function are metrics indicative of RFIC performance and allow comparison of on-Earth and in-orbit performance.

Note that an alternative test has one core in an RFIC interfering with the other 15 cores programmed to a static phase shift. Asymmetry in the amplitudes virtually

guarantees the correct identification of a core amplitude, but introduces other confounding factors: coupling between cores and RFIC heat generation associated with full 16-core operation which can alter or degrade core performance. This “1 vs. 15” test is *not* used for results presented in this paper.

Six on-Earth tests performed after integrating MAPLE onto the flight deck of V5 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. CMU phase functions from each of these tests were individually extracted by inverting $V_m(c)$ and unwrapped; the ensemble average of these curves over the six integration tests is shown in Fig. 4.21, in black. The CMU phase functions are all roughly linear and are capable, on average, of producing $\approx 3.37\pi$ rad of phase rotation over 512 unique phase codes. 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.




Figure 4.21: In black: unwrapped CMU phase functions, averaged over $N = 6$ interference tests performed during (on-Earth) integration prior to launch, for each core. In red: unwrapped CMU phase functions from the first in-orbit test. The label on the upper left indicates the “coordinates” of the core: (RFIC, quad, stack, core).

PA amplitudes, as measured by A'_m and A'_n , were consistent between tests, with standard deviation between tests, averaged over all cores, of 2.92% of the mean for

A'_1 and 1.35% for A'_2 . PA amplitudes varied between cores, with a standard deviation between cores of 19.77% for A'_1 and 41.77% for A'_2 . Given that each antenna is positioned differently, each will interact with MAPLE's cavity differently. Thus, PA amplitude variance is possibly owing to variance between scattering vectors. A comparison between on-Earth and in-orbit amplitudes is given in Fig. 4.22.




Figure 4.22: Comparing on-Earth to in-orbit PA amplitudes for 32 cores across two RFICs. Error bars indicate spread (standard deviation) among tests performed on Earth (x-axis) and in orbit (y-axis).

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. Fig. 4.23 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 performance in only three stacks (six cores), circled on the left of Fig. 4.23. 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




Figure 4.23: 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 ($i \in [1, 32]$) for the first core in the stack. Degraded cores circled in red.

(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 *potential* 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 over the experiment lifetime.

4.8 Beam to Earth Ground Power Detection

As a demonstration of both MAPLE’s functionality and ability to guide its transmitted energy in the desired direction, a proof-of-concept experiment was performed wherein MAPLE beamed RF power to a ground station at Caltech in Pasadena, CA, USA where it was detected. (Note that “Power transfer” is defined as a net positive recovery. In contrast, “detection” is defined as a process where the recovered power is less than the external power needed for the receiver.)

In these experiments, MAPLE was scheduled to turn on, steer power out the sapphire window opposite the array, and maintain its state for several minutes while passing

over Pasadena. The ground station used for detection during the experiment was assembled on the roof of Caltech's Moore Laboratory. Discussed below are the ground station design details and results from the beaming experiment.

Ground Station Design

The ground station consists of two critical elements: an RF receiver chain and a tracking apparatus.

RF Receiver

MAPLE transmits a 9.984GHz continuous wave, sinusoidal signal with an EIRP of approximately between 26.5-28.5dBm (Fig. 4.25). With a projected free space path loss of 167dB¹, the predicted incoming signal power, assuming ideal pointing coordination, is approximately -138.5dBm to -140.5dBm, implying a requirement for a high-gain receiver. A 33.5dBi gain, 2.8° half-power beamwidth, dual-polarized dish antenna served as the receiving aperture, providing high-gain while being small enough (67cm diameter) to fit on commercial telescope tracking mounts.

The heterodyne receiver, shown in Fig. 4.24, down-converted MAPLE's 9.984GHz signal to a 5 MHz IF signal which was then filtered, amplified, and digitized. A low-noise amplifier was attached directly to the output waveguide transition on the antenna, before the mixer, to amplify the signal and reduce the noise contribution of subsequent stages. A 2-7MHz bandpass filter removed potential interference signals from the IF. The receiver has a gain of +68.5dB with a total noise figure of 5dB. Two identical receivers were used for each polarization of the dish antenna.




Figure 4.24: Block diagram of RF receiver chain for ground station. There is an additional 3dB of loss due to cabling from the low-noise amplifier to the mixer.

The IF signal from the receiver was digitized at 20Msps by a Keysight M3100A PXIe Digitizer module for additional processing and analysis. The received signal's

¹167dB free space path loss is based on a payload directly overhead, at an orbital altitude of approximately 527km.

SNR was maximized using a matched filter which, in this case, was a Fast Fourier Transform (FFT) due to the near-sinusoidal nature of the signal. A processing interval of 26ms allowed for a sufficiently high SNR without introducing frequency smearing between FFT bins owing to a continuous Doppler shift. 16 IF spectra were produced every second throughout the pass.

Tracking

Tracking involved coordination between three individual entities: the timing of MAPLE's software operations, Vigoride-5's (V5's) attitude, and on-Earth ground station tracking.

MAPLE operations were programmed in advance, timed to begin in coordination with the payload crossing the horizon relative to Pasadena, CA, USA. Timing accuracy was confirmed by other experiments and proved to be more than sufficient.

The 16-element subarrays driven by Tile 1 and Tile 2 are placed side by side in MAPLE; each combination of Tile 1, Tile 2, and Tile 1+Tile 2 corresponds to a slightly different main beam direction. These beams, characterized under nearly identical circumstances in an absorbing chamber on Earth, are shown in Fig. 4.25. The main beam is not exactly broadside, owing to the relative skew of the viewing window and the selection of active radio-frequency integrated circuits (RFICs). In all of the "Beam-to-Earth" experiments, only Tile 1 was powered on while Tile 2 remained off. This is because using only a single RFIC trades potentially higher power for higher stability owing to possible chip-to-chip interference. Tile 1 proved more stable overall, as shown in figures 4.18 and 4.19. Accordingly, V5 was programmed to accommodate a pointing correction of $+20^\circ$ relative to broadside so that MAPLE's main beam would be directed at the ground station.

To maximize received power from MAPLE on Earth, the ground station antenna continuously pointed at ("tracked") the payload during the pass. Because of the high-gain dish antenna's narrow beamwidth, pointing accuracy was critical to avoid a large power penalty for misalignment. To this end, the antenna was mounted on a modified telescope mount (the iOptron AZ Mount Pro) with a sufficiently high peak slew rate of $10^\circ/\text{sec}$.

The orbit-tracking mount coordinated pointing at MAPLE using a two-line element set collected within 24 hours of the pass from the US Space Force website [80]. Immediately before the pass, the mount's built-in GPS system was calibrated to

known celestial objects (e.g., stars, planets, and/or the moon) through an optical sight aligned with the antenna's main beam. The full ground station, assembled before the beam-to-Earth attempt, is pictured in Fig.4.26.




Figure 4.25: MAPLE antenna pattern through the sapphire window for radiation from Tile 1, Tile 2, and both tiles. Tiles are labeled in figure.




Figure 4.26: Full ground station assembled before the July 28th, 2023 UTC nighttime pass. a) Narrow beamwidth dish antenna, b) wide beamwidth horn antenna, c) RF receiver chain, power supplies, frequency source, digitizer, etc.




Figure 4.27: Beam-to-Earth experiment results from 7/28/2023 UTC. Top to bottom: power amplifier current draw for RFIC 1, IF power received at the digitizer, Doppler shift from IF center frequency (5MHz), and orbital tracking maps showing MAPLE during the pass. Orange ellipses in the bottom row depict the -3dB isopower contour for MAPLE’s main beam. Three vertical, dotted lines (labeled “1”, “2”, “3”) show the times at which the three bottom photos (1,2,3, respectively) were captured.

Beam-to-Earth Results

Three beam-to-Earth experiments were attempted on the following dates: May 23rd, 2023 UTC, June 30th, 2023 UTC, and July 28th, 2023 UTC. In all of the experiments, MAPLE’s beam was successfully detected with the expected Doppler shift and power-level at a time consistent with the orbital predictions and ephemeris data from V5. The final experiment on July 28th, 2023 UTC provided the largest continuous set of data; those results are presented in Fig. 4.27.

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 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. 4.27 for three different times. Unfortunately, the ground station was never within MAPLE's -3dB contour, due to the host'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. 4.27 account for pointing mismatch using a beam pattern measured on the ground with the Engineering Model (EM). 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.

4.9 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).



(a) On-Earth test performed directly prior to launch



(b) In-orbit test performed on March 3rd, 2023

Figure 4.28: The first part of an “Initial Operation Sequence”: Two optimizations to Rectenna 2, phase recall to Rectenna 2, one optimization to Rectenna 1, phase recall to Rectenna 1, and then three successive optimizations to Rectenna 1, each optimization starting its search at the last setting of the previous optimization.

Symptoms

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 dramatically low dips in power as the algorithm stumbles on local maxima in 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 previously computed search space optimum.

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.

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. 4.28b.

Presented in Fig. 4.28a 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. 4.28a, there should be marginal increases in rectified power, but generally a flat curve.

However, in Fig. 4.28b, a different pattern emerges. The first optimization, to 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 power amplifiers (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.

Power Amplifier (PA) Supply Current

PA supply currents measured during in-orbit tests also appeared anomalous in some




Figure 4.29: Supply current drop anomaly on RFIC 1. Left: Supply current for RFIC 1 (Line A, Line B) measured over time. Right: Supply current for Line A plotted against board temperature. t_{knee} , T_{knee} , and I_{equil} labeled.


cases. PA supply lines are routed from the flight computer to the aggregation panel (AP) where they are voltage-regulated and then passed to the flex array for routing to the two radio-frequency integrated circuits (RFICs). Four PA supply lines (1A, 1B, 2A, 2B) route to the two RFICs (1/2); for RFIC 1, 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. 4.29. This point in time, called t_{knee} , varied between RFICs (t_{knee} , RFIC 1 ≈ 100 s, t_{knee} , RFIC 2 ≈ 20 s) and between tests. For reasons discussed below, it’s likely the drop in current is triggered by increased temperature; the temperature of the AP at t_{knee} is called T_{knee} . Both quantities are labelled in Fig. 4.29.


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

Chip-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 clock multiplier unit (CMU) phase function is inverted to solve for the necessary unique phase code that will accomplish this phase shift. On Earth, these algorithms



(a) On-Earth test performed directly prior to launch



(b) In-orbit test performed on March 21, 2023

Figure 4.30: Rectified power during linear sweep of phase offset added to RFIC 2.

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. 4.30. 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 drops², 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.

These metrics were tracked during simple, long-duration tests performed over the mission lifetime to characterize long-term system degradation. In Fig. 4.31, it's clear RFIC 1 T_{knee} and I_{equil} drop over time, with cumulative PA usage. RFIC 2 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 RFIC 2.

Thermal metrics and results presented in Fig. 4.23 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. 4.31), exacerbates output power decreases and appears to worsen over time. Additionally, an $\approx 23\%$ decrease in peak achievable rectified power over eight months (Fig. 4.23) quantifies long-term power output degradation.
3. Individual, core-level degradation was observed in lower output amplitudes in six cores in Fig. 4.23.

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.

²As demonstrated in Fig. 4.29, 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 4.31: Long-term system degradation as indicated by two metrics, T_{knee} and I_{eq} , with respect to cumulative power amplifier (PA) usage.

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 beyond the designed temperature when PAs are turned on. The rate and asymptotic chip temperature are functions of thermal bond degradation—RFIC 2 might perform generally worse because its bond is degraded further. This means that RFIC 2 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.

The 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 temperature and vacuum (TVAC) chambers on Earth using the EM. When thermistors were mounted *close* to the chip³, in an EM model of MAPLE, during on-Earth TVAC testing, temperatures in excess of 330K and 360K were measured for RFIC 1 and RFIC 2, respectively. RFIC 1, at this temperature, behaved nominally while RFIC 2 presented all of the aforementioned symptoms. All testing indicates that the fact that anomalous behavior presented on RFIC 2 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

³The RFICs are thermally heat sunk on one side and flip-chipped to an interposer installed on the backside of the flexible printed circuit board (PCB) on the other. The closest location, thermally, to the chip is on the 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.

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 reversibility of symptoms related to thermal bond damage when RFICs are either cool or in ambient pressure, the time lag associated with symptom onset, and the possible asymmetry in phenomena presentation in only one RFIC precluded our detection of anomalous behavior in TVAC testing performed prior to launch. The necessary test conditions to trigger anomalous performance were not met and, thus, symptoms were not present in the accompanying test data. It's also possible that bond damage was induced by launch vibrations. Though vibration tests were performed prior to TVAC testing, vibration-induced damage is stochastic; damage may have occurred during launch that did not occur during vibration tests.

The RFIC-epoxy bond can be degraded both by vibration-induced failure and by workmanship flaws. Additionally, ground testing indicates contact pressure can affect the integrity of the bond, as expected based on the discussion in [81]. MAPLE's design does not apply pressure to the RFIC-epoxy bond, potentially increasing sensitivity to workmanship flaws and resulting in degraded performance. Ground-testing to further understand the failure-mode, in order to inform future missions, is ongoing.

Chapter 5

MUTUAL POWER OPTIMIZATION

There are many challenges to Space Solar Power (SSP). Caltech's vision for SSP must address reference distribution and timing synchronization as discussed in Chapter 3, whilst being stow-able and having a minimal mass footprint as discussed in Chapter 4¹. An additional challenge that must be considered in an SSP system is that of area matching.

5.1 The Area Matching Problem

A key challenge in SSPP's integration of its photovoltaic (PV) and wireless power transfer (WPT) systems is areal matching of the two systems. The PV aperture should be filled to maximize available sunlight, while the RF aperture must be filled to prevent grating lobes arising from sparse arrays. For a given unit area, the amount of incident sunlight is determined by orbit. The WPT system is left with whatever power is left from the PV system's conversion efficiency. A Sankey diagram of the power flow containing estimated system efficiency figures is depicted in Fig. 5.1.

For PV system maximum efficiency, PV modules must operate at their maximum power point (MPP). The MPP will be dependent on the type of cell as well as the operating point (intensity of sunlight, temperature, etc.). An example of the IV-curve of a typical PV cell is shown in Fig. 5.2, with the cell delivering its maximum power at (V_{MPP}, I_{MPP}) .

Thus, for maximum PV efficiency, the load impedance presented to the PV cell should be (V_{MPP}/I_{MPP}) as shown in the left of Fig. 5.3; any other load impedance for the PV source, visualized as a load line in the right of Fig. 5.3, is suboptimal.

However this power, when received by the power amplifiers of the WPT system, must be also at voltages and currents that correspond to peak efficiency of the WPT system to maximize total system efficiency. This peak is mainly due to the behavior of the power amplifiers of the WPT system that dominate power usage (and thus power loss). As an efficiency-optimized WPT system typically uses switching amplifiers, only a narrow region of operation will yield high efficiency in WPT

¹Alex Ayling was a partner in the Mutual Power Optimization concept formation, system design, measurements, and documentation. This chapter is adapted from [82] and also uses material from [83].




Figure 5.1: Sankey Diagram [83] tracking losses using medium-term efficiency projections




Figure 5.2: Typical photovoltaic cell IV curve and MPP

systems in practice. For CMOS power amplifiers, this region should correspond to a large drain to source voltage. The position of the power amplifier's maximum efficiency point is adjustable at design time through impedance matching.

The areal matching challenge is thus simultaneously operating both PV and WPT systems in their narrow regions of optimal efficiency. While variances in orbital insolation are a key challenge for areal matching, there can be significant practical challenge in designing the PV and WPT systems. For instance, in a previous PV/WPT integration in [68], there was not enough available power from the PV system for the WPT system to operate optimally, resulting in a mismatch between the PV and RF aperture as shown in Fig. 5.4. This was due to discrepancies between




Figure 5.3: Left: load impedance versus power extracted for typical photovoltaic cell; Right: the same cell's IV curve with superimposed resistive load lines

the WPT system's intended input power (designed for the PV system's asymptotic efficiency) and the PV system's prototype efficiency. For continued PV system efficiency changes, each redesign of the WPT system, if desired, would need to be done at significant expense.




Figure 5.4: Ultralight integrated SSPP tile from [68] with areal mismatch: note that the RF aperture is only partially filled

Fig. 5.2 utilizes a simple model for the photovoltaic panel. This consists of a diode for the junction and a DC current source for photoelectric current generation. The MPP will lie in the middle of the IV curve of the panel which is bounded by the panel's short-circuit current and open-circuit voltage. Note that the photocurrent due to incident light determines the IV curve's vertical position. Visually, a resistive load extracting maximum power must pass through the origin and the MPP point. Operation outside of the MPP will reduce the PV system's efficiency.

Note that the PV system's baseline photocurrent (its IV curve's vertical position) will depend on the spectrum and amount of incident light. While solar panels in orbit do not suffer from large variations that terrestrial panels do due to day-night

cycles, clouds, and seasons, they still experience variations in incident light, even in a stable geostationary orbit. This variation follows a cosine dependence between the direction of the incident sunlight and the PV panel's normal vector, as depicted in Fig. 5.5.




Figure 5.5: Normalized solar irradiance as a function of PV panel's angle

While both systems should be designed such that their operating points match, the orbit variable intensity of sunlight incident on the PV aperture implies that a single point match will not be valid for the whole system operation. The traditional solution to address this type of issue is the usage of a DC-DC converter paired with a maximum power point tracking (MPPT) algorithm. The MPPT algorithm presents the optimal load impedance to the PV system while appearing as the correct source for the actual load. As an example, for a PV source with MPP corresponding to 1.1Ω that needs to supply power to a constant load of 5Ω , Fig. 5.6 shows the usage of an ideal DC-DC converter between the source and load for applicable impedance conversion.




Figure 5.6: Usage of a DC-DC converter for MPP by impedance conversion

In the SSPP context, however, the focusing arrays must be ultra-lightweight to minimize the cost of the launch required to deploy the system to be economically

viable. There is a significant challenge in utilizing this approach for SSPP, stemming from added mass. This is because SSPP's economic viability is heavily predicated on the spacecraft's mass due to launch costs; a DC-DC converter system with MPPT raises the system's aerial mass density due to added passives, ICs, packaging, and their associated radiation shielding and heat-spreading. In the case of a commercial off the shelf solution, assuming one could procure a typical BGA part such as the LTM8063 by Analog Devices (with integrated switch and inductor to minimize mass), 27 milligrams of shielding and the 128 milligram DC-DC converter IC would add 155 milligrams of mass total. Using a single DC-DC converter per power transfer RFIC would around $30\text{g}/\text{m}^2$ areal mass increase, which would correspond to 72% areal mass increase when considering the projected asymptotic areal mass of the power transfer system shown in Fig. 5.7.




Figure 5.7: Projected asymptotic mass breakdown for the wireless power transfer system

This significant increase in mass would not be the only penalty to the system, as without an application-tailored custom IC, the DC-DC converter would likely operate at lower efficiency (a rough estimate might be 75%). In the most optimal case, assuming an application specific integrated circuit carefully designed in conjunction with the PV and power transfer systems in mind, we can estimate the mass penalty by looking at current state-of-the-art solutions. One such solution, the TPSM82822 made by Texas Instruments, would use around 52 milligrams for the full converter system (10 milligrams for the converter, 15 milligrams for the passives, and 27 milligrams for radiation and thermal management). This represents a $10\text{g}/\text{m}^2$ areal mass increase. Furthermore, it will be shown shortly that even if the converter were designed to operate at a high efficiency of around 90%, in certain operating regions

system performance would be worse than forgoing the converter completely. Another idea might be to integrate the DC-DC converter and MPP tracking within the RFIC. While this would remove the mass penalty, lossy integrated DC-DC converter passives in RFIC (which relies on utilizing a CMOS process to enable integration of the high complexity system) would limit viability. There is another fundamental limit of using a DC-DC converter: if the available power from the photovoltaics is too low, the DC-DC converter will produce an output voltage regulated to an acceptably low level, degrading PA performance.

5.2 Principles of Power Amplifier Duty Cycling

For these reasons, despite their frequency of use in terrestrial PV applications, investigation of methods besides DC-DC converters were considered. Detailed consideration was given to the concept of duty-cycling the operation of the PA. As this operation has conceptual similarities to that of DC-DC converters, first consider the operation of an idealized boost converter shown in Fig. 5.8. It can be shown that, given a duty cycle D of switching, the impedance R_{in} looking in to the converter is $(1 - D)^2$ times the load impedance R_L .




Figure 5.8: Idealized boost converter circuit and switching waveforms

Now, consider the simpler circuit in Fig. 5.9. Assuming a capacitor with C_L large compared to the switching frequency, note that $R_{in} = R_L/D$. In other words, the load resistance can be modulated by duty cycle adjustment, (provided a sufficiently large capacitor and the ability to turn on and off the load). Of course, in the context of the wireless power transfer system, this merely corresponds to switching on and off the power amplifiers. This technique will be herein referred to as Mutual Power Optimization (MPO).

By duty cycling the PAs, the RF power transmitted is modulated instead of continuous-wave (acceptable in the context of power transfer, as the system is not transmitting data). One important operational difference vs. the DC-DC converter scheme is that because there is no voltage transformation occurring, the PAs will operate at




Figure 5.9: Modulation of a load with capacitive bank

the PV voltage. This would require the co-design of the PV and PA operating voltages. In practice, if the voltage of a single respective unit is mismatched, this could be addressed by stacking PV cells and PA circuitry [2]. As the duty cycle is controllable, changes to the MPP of the PV and optimal quiescent conditions of the PAs (such as those due to temperature and intensity of insolation due to orbit) could be compensated for: this "impedance match" has dynamic capability.

Compared to the DC-DC converter, the extra cost and mass of the inductor and switch are avoided. Also, with a high enough switching frequency, there is no need to supply an extra capacitor to the wireless power system beyond the standard decoupling capacitance that is already used. For $\sim 15\mu F$ capacitors available in 0402 footprint, $\sim 1\text{MHz}$ switching frequency would be required—a relatively low frequency that would be easy to implement in the wireless power transfer RFIC. Additionally, PA duty-cycling has a higher conversion efficiency than DC-DC converters that suffer from switching and inductor losses. Finally, this scheme provides a flexible solution to the aforementioned area mismatch problem. An example of area mismatch is given in Fig. 5.10, where the PV system is unable to supply enough current for the PAs to operate at their desired voltage of $\sim 1.4\text{V}$. By decreasing the duty cycle, current draw is lowered to reach the MPP of the PV whilst operating the PAs at the correct voltage, addressing areal mismatch.

5.3 Hardware and Measurement

Experimental results through demonstration hardware act as proof-of-concept. Here, a wireless power transfer system composed of old MAPLE hardware containing a 4×4 array was used, paired with real solar panels for an in-field demonstration, and a solar cell emulator for an in-lab demonstration. A system diagram is shown in Fig. 5.11. The PAs were turned on and off via a SPI communication line, which put a practical limit on the duty cycling (maximum frequency of 16.2kHz , with




Figure 5.10: Effect of duty cycle lowering on quiescent operating conditions

achievable duty cycle in range $D = [0.17, 0.87]$). The transient peak in the rising edge is possibly thermal in nature. Although proper realization of duty cycling at high frequency would require design modification to the existing RFICs, the change would be minimal as it would effectively only entail providing direct access to the PA enable lines (accessible only via SPI in the current RFIC).




Figure 5.11: Experimental setup for MPO validation

In-lab setup and results

For lab use, a solar cell emulator with good controllability not requiring sunlight was desired. To this end, a programmable DC supply operated in constant current mode in parallel with a silicon diode was used (the diode modeling the junction and the amount of current the level of illumination). A capacitance of $300\mu F$ (scaled to a large value due to the relatively low switching frequency of 16kHz) served as the capacitive load bank. Characterization was done via DC electronic load; the experimental setup is shown in Fig. 5.12.

Fig. 5.13, plots, as a function of photocurrent, the maximum power received whilst duty-cycling. As comparisons, direct connection without duty cycling and power received by using an extrapolated DC-DC converter are also shown. Extrapolation for the DC-DC converter was chosen due to lack of commercial DC-DC converters capable of source MPP tracking and load tracking for the range of operating voltages




Figure 5.12: a) Solar cell emulator b) solar cell emulator IV curve c) lab experiment setup

and currents. The extrapolation assumes that the converter perfectly extracts maximum power from the PV cell, and converts it to load with a constant 90% efficiency. Note that using MPO has results in a significant increase in the receiver power for a wide range of PV powers; at low photocurrents; as much as a 10-fold increase compared to the baseline is observed. MPO enables the PAs to operate efficiently at V_{PV} even at low power levels; instead, the DC-DC converter needs to regulate the PA V_{DD} voltage to a low value to match input and output powers, resulting in less efficient PA operation. With higher photocurrents, the baseline receiver and duty-cycling powers converge (as PA operation nears the MPP of the PV obviating load modulation).




Figure 5.13: Power received as a function of photocurrent, with assumed DC-DC converter efficiency of 90%

In-field setup and results

An similar experiment was done using a commercial Passivated Emitter and Rear Contact (PERC) polycrystalline silicon cell. To vary PV power, periodic measurements were taken every 30 minutes starting at 8:30 A.M. and ending at 12:30 P.M. Fig. 5.14 depicts the in-field setup and measured PV power. A rectenna board serves as the receiver with detected power in the form of RMS voltages measurements. Results analogous to the in-lab experiments are depicted in Fig. 5.15 (identical assumptions for the DC-DC converter). For the entire duration, duty-cycling yields substantial performance uplift. This is especially the case early in the morning: with little power available from the PV, relative to the calibrated power meter, significantly more power is seen at the receiver—this is due to the nonlinear conversion efficiency of the (typical) rectifier circuit used (a more detailed discussion follows in Section 5.4).




Figure 5.14: PERC cell setup and PV MPP. Experiment conducted at Caltech on Oct. 16th, 2021.

5.4 Rectifiers in Wireless Power Transfer

Notwithstanding the fixed loss incurred by the antenna, the behavior of the rectenna's efficiency curve is dictated by its rectifier circuit. RF rectifiers include both integrated transistor-based rectifiers as well as discrete diode-based ones. Due to packaging parasitics, integrated transistor-based rectifiers appear more commonly at higher frequencies where they tend to perform better and can more effectively integrate smaller reactive components. Discrete diode-based rectifiers tend to appear more often at lower frequencies. As they only require an RF diode for rectification and can benefit from efficient PCB transmission line based structures, they have a lower barrier to entry. So far, discrete single diode-based rectifiers have been used




Figure 5.15: Rectenna power received as over the course of the day for PERC cells.

for SSPP to convert RF power at the system's 10 GHz operating frequency.

As a starting point to describing the RF rectifier, we first consider its elementary low frequency analogue: the half wave rectifier circuit. Essentially, the incident sinusoid is rectified by passing the half the waveform (at a single polarity) while truncating the other half. Ripple smoothing is commonly accomplished with an capacitive output filter. While at lower frequencies, the half-wave conduction of such a circuit reduces efficiency due to high I^2R losses incurred from pulsed current draw, at RF, efficient circuit design is enabled by a combination of the availability of waveform engineering and parasitic minimization of using a single diode.

One prominent perspective for RF rectifiers is the operation of a highly efficient switching amplifier in reverse [84]. Efficient discrete RF rectifiers, including the ones used in SSPP, must operate with a voltage swing close to what corresponds to the reverse breakdown voltage of the diode. This can be understood to first order by modeling the diode as an ideal diode with constant voltage drop: the larger the voltage swing, the smaller the fraction of lost power during the conduction cycle due to the diode's fixed voltage drop; taken to the limit, the voltage swing across the diode will be close to the sum of the magnitudes of the diode forward voltage and reverse breakdown voltage. From here, it can be understood that the idealized voltage waveform of a discrete RF rectifier is clearly not a sinusoid, but a square wave as seen in Class-F amplifiers. Class-F amplifiers employ waveform engineering to maximize odd order harmonics (ideally forming the Fourier series of a square wave); discrete RF rectifiers are no exception.

As they have a low forward voltage and high speed as majority carrier devices, Schottky diodes are the de-facto choice for discrete RF rectifiers. As high rectification efficiency is achieved only in narrow operating conditions, for a particular power level and line impedance, a Schottky diode of appropriate forward and reverse voltages should be selected. There is some leeway available in the input matching network's impedance transformation, but the transformation ratio will be limited by the finite Q of the matching network elements.

Higher order effects and terms that are significant when designing discrete RF rectifiers are discussed in detail in [85]. In particular, switching loss due to the diode capacitance is a frequency dependent term that should not be neglected and is unavoidable even with optimal matching, as real diodes incur loss through their equivalent series resistance while charging and discharging their junction capacitance.

5.5 Conclusion

The power and impedance mismatches between the PV source and PA load in an SSPP system must be addressed for optimal system efficiency. Areal mismatch provides a challenge for system design and integration that cannot easily be addressed with existing solutions such as DC-DC converters. Experiments duty cycling the power amplifier act as proof-of-concept for a proposed solution allowing PA operation at a higher efficiency whilst also tracking the MPP of the PV source. Future work should entail demonstration of this concept at high enough frequency to obviate need for the additional mass of bypass capacitors. Substantial consideration must also be given to the unintended side effects of modulation such as out of band emission.

BIBLIOGRAPHY

- [1] I. Şeker, “Calibration methods for phased array radars,” in *Radar Sensor Technology XVII*, K. I. Ranney and A. Doerry, Eds., International Society for Optics and Photonics, vol. 8714, SPIE, 2013, 87140W. doi: 10.1117/12.2015694. [Online]. Available: <https://doi.org/10.1117/12.2015694>.
- [2] A. Hajimiri, B. Abiri, F. Bohn, M. Gal-Katziri, and M. H. Manohara, “Dynamic focusing of large arrays for wireless power transfer and beyond,” *IEEE Journal of Solid-State Circuits*, vol. 56, no. 7, pp. 2077–2101, 2021. doi: 10.1109/JSSC.2020.3036895.
- [3] A. Wu, I. A. Syed, A. Ayling, and A. Hajimiri, “On-board array self-calibration using amplitude-only proximal-field sensors and machine-learning-based phase retrieval,” in *2025 IEEE/MTT-S International Microwave Symposium - IMS 2025*, 2025, pp. 930–933. doi: 10.1109/IMS40360.2025.11103983,
- [4] A. Wu, I. A. Syed, A. Ayling, and A. Hajimiri, “Amplitude-only proximal-field sensors and machine-learning-based phase retrieval for array calibration and monitoring,” To be submitted to *IEEE Transactions on Microwave Theory and Techniques*, 2026.
- [5] O. M. Bakr and M. Johnson, “Impact of phase and amplitude errors on array performance,” University of California at Berkeley, Electrical and Computer Sciences, Tech. Rep., Jan. 2009. [Online]. Available: <https://www2.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-1.pdf>.
- [6] M. Gal-Katziri, “Precision at scale: System design from tiny biosensors to giant arrays,” Available at https://thesis.library.caltech.edu/13857/14/Matan_Gal_2020_Final_Rev01.pdf, PhD thesis, California Institute of Technology, Pasadena, CA, Jul. 2020.
- [7] A. Safaripour, B. Asghari, M. R. M. Hashemi, and A. Hajimiri, “Proximal-field sensing: In situ prediction of far-field radiation for integrated radiators,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 67, no. 9, pp. 3743–3756, 2019. doi: 10.1109/TMTT.2019.2929154.
- [8] J. Lavaei, A. Babakhani, A. Hajimiri, and J. C. Doyle, “Solving large-scale hybrid circuit-antenna problems,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 58, no. 2, pp. 374–387, 2011. doi: 10.1109/TCSI.2010.2072010.
- [9] A. Walther, “The question of phase retrieval in optics,” *Optica Acta: International Journal of Optics*, vol. 10, no. 1, pp. 41–49, 1963. doi: 10.1080/713817747. eprint: <https://doi.org/10.1080/713817747>. [Online]. Available: <https://doi.org/10.1080/713817747>.

- [10] R. W. Harrison, "Phase problem in crystallography," *J. Opt. Soc. Am. A*, vol. 10, no. 5, pp. 1046–1055, May 1993. DOI: 10.1364/JOSAA.10.001046. [Online]. Available: <https://opg.optica.org/josaa/abstract.cfm?URI=josaa-10-5-1046>.
- [11] J. C. Dainty and J. R. Fienup, *Image Recovery: Theory and Application*. Academic Press, 1987, pp. 231–275.
- [12] A. Safaripour, "Proximal-field radiation sensors for dynamically controllable and self-correcting integrated radiators," Available at https://thesis.library.caltech.edu/10178/19/Safaripour_Thesis.pdf, PhD thesis, California Institute of Technology, Pasadena, CA, Apr. 2017.
- [13] A. Conca, D. Edidin, M. Hering, and C. Vinzant, "An algebraic characterization of injectivity in phase retrieval," *Applied and Computational Harmonic Analysis*, vol. 38, no. 2, pp. 346–356, 2015, ISSN: 1063-5203. DOI: <https://doi.org/10.1016/j.acha.2014.06.005>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1063520314000876>.
- [14] T. Heinosaari, L. Mazzarella, and M. M. Wolf, "Quantum tomography under prior information," *Communications in Mathematical Physics*, vol. 318, no. 2, pp. 355–374, Mar. 2013, ISSN: 1432-0916. DOI: 10.1007/s00220-013-1671-8. [Online]. Available: <https://doi.org/10.1007/s00220-013-1671-8>.
- [15] N. Shlezinger, R. Dabora, and Y. C. Eldar, "Measurement matrix design for phase retrieval based on mutual information," *IEEE Transactions on Signal Processing*, vol. 66, no. 2, pp. 324–339, 2018. DOI: 10.1109/TSP.2017.2759101.
- [16] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval via wirtinger flow: Theory and algorithms," *IEEE Transactions on Information Theory*, vol. 61, no. 4, pp. 1985–2007, 2015. DOI: 10.1109/TIT.2015.2399924.
- [17] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval from coded diffraction patterns," *Applied and Computational Harmonic Analysis*, vol. 39, no. 2, pp. 277–299, 2015, ISSN: 1063-5203. DOI: <https://doi.org/10.1016/j.acha.2014.09.004>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1063520314001201>.
- [18] P. Salanevich and G. E. Pfander, "Polarization based phase retrieval for time-frequency structured measurements," in *2015 International Conference on Sampling Theory and Applications (SampTA)*, 2015, pp. 187–191. DOI: 10.1109/SAMPTA.2015.7148877.
- [19] B. Alexeev, A. S. Bandeira, M. Fickus, and D. G. Mixon, "Phase retrieval with polarization," *SIAM Journal on Imaging Sciences*, vol. 7, no. 1, pp. 35–66, 2014. DOI: 10.1137/12089939X. eprint: <https://doi.org/10.1137/12089939X>.

- 1137/12089939X. [Online]. Available: <https://doi.org/10.1137/12089939X>.
- [20] J. P. D'Angelo, *Inequalities from complex analysis*. Mathematical Association of America, 2002, vol. 28, pp. 53–54.
 - [21] R. Chandra, T. Goldstein, and C. Studer, “Phasepack: A phase retrieval library,” in *2019 13th International conference on Sampling Theory and Applications (SampTA)*, 2019, pp. 1–5. doi: 10.1109/SampTA45681.2019.9030878.
 - [22] J. González-López, S. Ventura, and A. Cano, “Distributed selection of continuous features in multilabel classification using mutual information,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 7, pp. 2280–2293, 2020. doi: 10.1109/TNNLS.2019.2944298.
 - [23] T. M. Cover and J. A. Thomas, *Chapter 8 - Differential Entropy, Elements of Information Theory, 2nd ed.* Hoboken, New Jersey: Wiley-Interscience, 2006.
 - [24] Y. Rivenson, Y. Zhang, H. Günaydın, D. Teng, and A. Ozcan, “Phase recovery and holographic image reconstruction using deep learning in neural networks,” *Light: Science & Applications*, vol. 7, no. 2, pp. 17141–17141, Feb. 2018, ISSN: 2047-7538.
 - [25] J. Dong, L. Valzania, A. Maillard, T.-a. Pham, S. Gigan, and M. Unser, “Phase retrieval: From computational imaging to machine learning: A tutorial,” *IEEE Signal Processing Magazine*, vol. 40, no. 1, pp. 45–57, 2023.
 - [26] A. Ayling and A. Hajimiri, “Flip-chip aperture coupled d-band active radiator tiles in 22-nm cmos fdsoi,” *IEEE Journal of Solid-State Circuits*, vol. 60, no. 6, pp. 1932–1946, 2025. doi: 10.1109/JSSC.2024.3485062.
 - [27] A. Singer, “Angular synchronization by eigenvectors and semidefinite programming,” *Applied and Computational Harmonic Analysis*, vol. 30, no. 1, pp. 20–36, 2011, ISSN: 1063-5203. doi: <https://doi.org/10.1016/j.acha.2010.02.001>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1063520310000205>.
 - [28] A. Giridhar and P. R. Kumar, “Distributed clock synchronization over wireless networks: Algorithms and analysis,” in *Proceedings of the 45th IEEE Conference on Decision and Control*, IEEE, 2006, pp. 4915–4920.
 - [29] A. Agrawal, R. Raskar, and R. Chellappa, “What is the range of surface reconstructions from a gradient field?” In *European conference on computer vision*, Springer, 2006, pp. 578–591.
 - [30] J. Rubinstein and G. Wolansky, “Reconstruction of optical surfaces from ray data,” *Optical review*, vol. 8, no. 4, pp. 281–283, 2001.
 - [31] D. E. Williams, “Shape-changing phased arrays,” Ph.D. dissertation, California Institute of Technology, 2022.

- [32] A. Ayling, A. Fikes, O. S. Mizrahi, *et al.*, “Wireless power transfer in space using flexible, lightweight, coherent arrays,” *Acta Astronautica*, vol. 224, pp. 226–243, 2024, ISSN: 0094-5765. DOI: <https://doi.org/10.1016/j.actaastro.2024.08.006>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S009457652400434X>,
- [33] I. Asimov, *Reason*. Street & Smith, Apr. 1941.
- [34] M. A. Marshall, “Dynamics of ultralight flexible spacecraft during slew maneuvers,” Ph.D. dissertation, California Institute of Technology, 2022.
- [35] B. Abiri, M. Arya, F. Bohn, *et al.*, *A lightweight space-based solar power generation and transmission satellite*, 2022. eprint: 2206.08373 (astro-ph.IM).
- [36] P. Jaffe and J. McSpadden, “Energy conversion and transmission modules for space solar power,” *Proceedings of the IEEE*, vol. 101, no. 6, pp. 1424–1437, 2013. DOI: 10.1109/JPROC.2013.2252591.
- [37] B. Strassner and K. Chang, “Highly efficient c-band circularly polarized rectifying antenna array for wireless microwave power transmission,” *IEEE Transactions on Antennas and Propagation*, vol. 51, no. 6, pp. 1347–1356, 2003. DOI: 10.1109/TAP.2003.812252.
- [38] A. Fikes, M. Gal-Katziri, E. Gdoutos, *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.
- [39] C. T. Rodenbeck, P. I. Jaffe, B. H. Strassner II, *et al.*, “Microwave and millimeter wave power beaming,” *IEEE Journal of Microwaves*, vol. 1, no. 1, pp. 229–259, 2021. DOI: 10.1109/JMW.2020.3033992.
- [40] P. E. Glaser, “Power from the sun: Its future,” *Science*, vol. 162, no. 3856, pp. 857–861, 1968. DOI: 10.1126/science.162.3856.857. eprint: <https://www.science.org/doi/pdf/10.1126/science.162.3856.857>. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.162.3856.857>.
- [41] W. C. Brown, “Satellite power stations: A new source of energy?” *IEEE Spectrum*, vol. 10, no. 3, pp. 38–47, 1973. DOI: 10.1109/MSPEC.1973.5216798.
- [42] 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.
- [43] J. C. Mankins, “A technical overview of the “suntower” solar power satellite concept,” *Acta Astronautica*, vol. 50, no. 6, pp. 369–377, 2002, ISSN: 0094-5765. DOI: [https://doi.org/10.1016/S0094-5765\(01\)00167-9](https://doi.org/10.1016/S0094-5765(01)00167-9). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576501001679>.

- [44] W. Seboldt, M. Klimke, M. Leipold, and N. Hanowski, "European sail tower sps concept," *Acta Astronautica*, vol. 48, no. 5, pp. 785–792, 2001, ISSN: 0094-5765. doi: [https://doi.org/10.1016/S0094-5765\(01\)00046-7](https://doi.org/10.1016/S0094-5765(01)00046-7). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576501000467>.
- [45] M. Mori, H. Kagawa, and Y. Saito, "Summary of studies on space solar power systems of Japan Aerospace Exploration Agency (JAXA)," *Acta Astronautica*, vol. 59, no. 1, pp. 132–138, 2006, Space for Inspiration of Humankind, Selected Proceedings of the 56th International Astronautical Federation Congress, Fukuoka, Japan, 17-21 October 2005, ISSN: 0094-5765. doi: <https://doi.org/10.1016/j.actaastro.2006.02.033>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576506001032>.
- [46] S. Sasaki, K. Tanaka, K. Higuchi, *et al.*, "A new concept of solar power satellite: Tethered-sps," *Acta Astronautica*, vol. 60, no. 3, pp. 153–165, 2007, ISSN: 0094-5765. doi: <https://doi.org/10.1016/j.actaastro.2006.07.010>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576506002815>.
- [47] I. Cash, "Cassiopeia – a new paradigm for space solar power," *Acta Astronautica*, vol. 159, pp. 170–178, 2019, ISSN: 0094-5765. doi: <https://doi.org/10.1016/j.actaastro.2019.03.063>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576518320708>.
- [48] Y. Yang, Y. Zhang, B. Duan, D. Wang, and X. Li, "A novel design project for space solar power station (ssps-omega)," *Acta Astronautica*, vol. 121, pp. 51–58, 2016, ISSN: 0094-5765. doi: <https://doi.org/10.1016/j.actaastro.2015.12.029>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576515300680>.
- [49] X. Li, B. Duan, L. Song, Y. Yang, Y. Zhang, and D. Wang, "A new concept of space solar power satellite," *Acta Astronautica*, vol. 136, pp. 182–189, 2017, ISSN: 0094-5765. doi: <https://doi.org/10.1016/j.actaastro.2017.03.017>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576516311092>.
- [50] K. Yamanaka, Y. Tuyama, H. Ohtsuka, S. Chaki, M. Nakayama, and Y. Hirano, "Internally-matched GaN HEMT high efficiency power amplifier for space solar power stations," in *2010 Asia-Pacific Microwave Conference*, 2010, pp. 119–122.
- [51] F. Bohn, B. Abiri, and A. 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](https://doi.org/10.1109/RFIC.2017.7969054).

- [52] M. R. M. Hashemi, A. C. Fikes, M. Gal-Katziri, *et al.*, “A flexible phased array system with low areal mass density,” *Nature electronics*, vol. 2, no. 5, pp. 195–205, 2019.
- [53] M. Gal-Katziri, A. Fikes, and A. Hajimiri, “Flexible active antenna arrays,” *npj Flexible Electronics*, vol. 6, no. 1, p. 85, 2022.
- [54] O. S. Mizrahi, A. Fikes, A. Truong, F. Wiesemüller, S. Pellegrino, and A. Hajimiri, “Popup arrays for large space-borne apertures,” *IEEE Transactions on Antennas and Propagation*, vol. 72, no. 8, pp. 6387–6400, 2024. doi: 10.1109/TAP.2024.3414593.
- [55] W. Brown, “Optimization of the efficiency and other properties of the rectenna element,” in *1976 IEEE-MTT-S International Microwave Symposium*, 1976, pp. 142–144. doi: 10.1109/MWSYM.1976.1123673.
- [56] J. McSpadden, L. Fan, and K. Chang, “Design and experiments of a high-conversion-efficiency 5.8-ghz rectenna,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 46, no. 12, pp. 2053–2060, 1998. doi: 10.1109/22.739282.
- [57] E. C. Warmann, P. Espinet-Gonzalez, N. Vaidya, *et al.*, “An ultralight concentrator photovoltaic system for space solar power harvesting,” *Acta Astronautica*, vol. 170, pp. 443–451, 2020, issn: 0094-5765. doi: <https://doi.org/10.1016/j.actaastro.2019.12.032>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576519314651>.
- [58] E. Gdoutos, C. F. Sommer, A. Truong, *et al.*, “Development of the deployable on-orbit ultralight composite experiment (dolce) for the space solar power project demonstration mission,” in *AIAA SciTech 2022 Forum*, 2022, p. 1266.
- [59] W. C. Brown, “Experimental airborne microwave supported platform,” *Raytheon Co Burlington, MA Microwave and Power Tube Div*, 1965.
- [60] R. Dickinson and W. C. Brown, “Radiated microwave power transmission system efficiency measurements,” Jet Propulsion Laboratories, Tech. Rep., 1975.
- [61] R. Dickinson, “Evaluation of a microwave high-power reception-conversion array for wireless power transmission,” Jet Propulsion Laboratories, Tech. Rep., 1975.
- [62] W. Brown, “The technology and application of free-space power transmission by microwave beam,” *Proceedings of the IEEE*, vol. 62, no. 1, pp. 11–25, 1974. doi: 10.1109/PROC.1974.9380.
- [63] M. Nagatomo, N. Kaya, and H. Matsumoto, “Engineering aspect of the microwave ionosphere nonlinear interaction experiment (minix) with a sounding rocket,” *Acta Astronautica*, vol. 13, no. 1, pp. 23–29, 1986.

- [64] J. McSpadden, N. Kaya, A. Brown, and K. Chang, "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.
- [65] C. T. Rodenbeck, B. B. Tierney, J. Park, *et al.*, "Terrestrial microwave power beaming," *IEEE Journal of Microwaves*, vol. 2, no. 1, pp. 28–43, 2022. DOI: 10.1109/JMW.2021.3130765.
- [66] H. Matsumoto, "Research on solar power station and microwave power transmission in japan : Review and perspectives," *IEEE Microwave Magazine*, pp. 36–45, Dec. 2002.
- [67] H. Matsumoto, "Microwave power transmission from space and related non-linear plasma effects," *The Radio Science Bulletin*, no. 273, pp. 11–35, 1995.
- [68] M. D. Kelzenberg, P. Espinet-Gonzalez, N. Vaidya, *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.
- [69] M. Gal-Katziri, A. Fikes, F. Bohn, B. Abiri, M. R. Hashemi, and A. Hajimiri, "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.
- [70] A. C. Fikes, M. Gal-Katziri, O. S. Mizrahi, D. E. Williams, and A. Hajimiri, "Frontiers in flexible and shape-changing arrays," *IEEE Journal of Microwaves*, vol. 3, no. 1, pp. 349–367, 2022.
- [71] A. Fikes, O. S. Mizrahi, and A. Hajimiri, "A framework for array shape reconstruction through mutual coupling," *IEEE Transactions on Microwave Theory and Techniques*, vol. 69, no. 10, pp. 4422–4436, 2021. DOI: 10.1109/TMTT.2021.3097729.
- [72] O. S. Mizrahi, A. Fikes, and A. 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.
- [73] A. Fikes, E. Gdoutos, M. Klezenberg, *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.
- [74] A. Fikes, O. S. Mizrahi, A. Truong, F. Wiesemüller, S. Pellegrino, and A. Hajimiri, "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.

- [75] O. S. Mizrahi, A. Fikes, A. Truong, F. Wiesemüller, S. Pellegrino, and A. Hajimiri, “Popup arrays for large space-borne apertures,” *IEEE Transactions on Antennas and Propagation*, vol. 72, no. 8, pp. 6387–6400, 2024. doi: 10.1109/TAP.2024.3414593.
- [76] Y. Kim, Y. J. Yoon, J. Shin, and J. So, “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.
- [77] A. C. Fikes, “Future microwave arrays take shape,” Ph.D. dissertation, California Institute of Technology, 2022.
- [78] G. R. Allen, L. Z. Scheick, F. Irom, *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)*, IEEE, 2015, pp. 1–14.
- [79] G. R. Allen, L. Z. Scheick, F. Irom, *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)*, IEEE, 2015, pp. 1–14.
- [80] USSF, *Space-track.org*, United States Space Force, Arlington County, VA, USA, 2023.
- [81] D. Gluck, V. Baturkin, and D. Gilmore, “Mountings and interfaces,” *Spacecraft thermal control handbook*, vol. 1, 2002.
- [82] A. Ayling, A. Wu, and A. 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,
- [83] O. S. Mizrahi, P. Jahelka, E. Gdoutos, *et al.*, “Space solar power generation: A viable system proposal and technoeconomic analysis,” *Joule*, vol. 9, no. 6, Jun. 2025, issn: 2542-4785. doi: 10.1016/j.joule.2025.101928. [Online]. Available: <https://doi.org/10.1016/j.joule.2025.101928>,
- [84] D. Hamill, “Time reversal duality and the synthesis of a double class e dc-dc converter,” in *21st Annual IEEE Conference on Power Electronics Specialists*, 1990, pp. 512–521. doi: 10.1109/PESC.1990.131231.
- [85] J. Guo, H. Zhang, and X. Zhu, “Theoretical analysis of rf-dc conversion efficiency for class-f rectifiers,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 62, no. 4, pp. 977–985, 2014. doi: 10.1109/TMTT.2014.2298368.

