Blind Wireless Beamforming to Power, Heat, and Move

Submitted in partial fulfillment of the requirements for

the degree of

Doctor of Philosophy

in

Electrical and Computer Engineering

Jingxian Wang

M.Sc., Electrical and Computer Engineering, Carnegie Mellon University
B.Sc., Telecommunications Engineering with Management, Beijing University of Posts and Telecommunications

Carnegie Mellon University
Pittsburgh, PA

August 2022

Acknowledgements

This five-year Ph.D. journey has been full of learning and discovering. I am massively indebted to my amazing advisor, Swarun Kumar, without whom my entire Ph.D. journey would not have been possible. Swarun supported me and inspired me. Swarun always encouraged me to think outside the box, which led to many exciting projects and best paper awards. I am so grateful to have had such a wonderful mentor guiding me through my graduate career. Thank you for being such a great role model. I hope that I can help students overcome obstacles and inspire them as Swarun has done for me.

I am profoundly grateful to my amazing thesis committee: Peter Steenkiste, Carmel Majidi, and Vaishnavi Ranganathan. I am thankful to Peter, who has been immensely helpful in shaping this thesis. I have also had an amazing experience working with Carmel, who exposed me to mechanical engineering, material science, and soft robotics. It was a truly enjoyable experience working with Vaishnavi, who has been tremendously supportive throughout my internship and job search.

I am likewise greatly indebted to my amazing collaborators, without whom much of this research would not have been possible [1, 2, 3, 4, 5, 6, 7]. I am truly grateful to work with Haojian Jin on a number of projects in RFID. I have truly enjoyed brainstorming with Haojian at our dinner in the first year of my Ph.D. It has been an incredibly fun experience working with Haojian, who I have learned so much from. Junbo Zhang was the powerhouse through our collaboration on PushID. Junbo and I spent many sleepless nights in the CIC garage to do experiments. I had great fun working with Chengfeng Pan and Mason Zadan, who helped me fabricate hundreds of stretchable and soft robot samples.

I am grateful to members of the WiTech lab who I have journeyed with through my Ph.D.: Akshay Gadre, Diana Zhang, Vaibhav, Yiwen Song, Akarsh Prabhakara, Atul Bansal, Kuang Yuan, and Mohamed Ibrahim Ahmed. In addition, I would like to thank National Science Foundation and Microsoft for their generous funding support.

Most of all, to my parents, Xinyu Wang and Bo Gao, thank you for all the love and support. This thesis is dedicated to you.

Abstract

The past decades witnessed a proliferation of wireless systems research that has focused on designing more efficient communication and novel sensing applications. However, another well-known capability of wireless has arguably received less attention: the ability to heat, and more generally, *provide energy*. As an example, consider the microwave oven, a nearly 50-year-old invention that to this day continues to remain largely similar in capability, blasting energy blindly to re-heat food without advancing significantly in terms of fine-grained heating capabilities. Similarly, the existing energy delivery systems for powering battery-free IoT tags energize its surroundings blindly, constraining the battery-free tags to be detected only at specific checkpoints – a range of 5-15 meters at best. While precise control of wireless signals, also known as beamforming, has been an important solution in improving wireless communication systems' throughput and signal-noise ratio, these systems have not directly translated to innovations in wireless systems for heating and energy delivery more broadly.

In this dissertation, we build a suite of *blind beamforming* systems that deliver precisely controllable wireless energy to battery-free devices (e.g., RFID) and objects (e.g., pizza). First, this thesis describes PushID, a beamforming system that powers battery-free RFID tags at scale. Second, inspired by PushID, we present SDC, a redesigned microwave oven system to heat food in a software-defined and fine-grained thermal trajectory. Further, we build on the same heating principle to develop WASER, which moves electronics-free soft robots using wireless beamforming. Besides these, this thesis demonstrates novel battery-free soft materials that can efficiently harvest wireless energy and be capable of advanced sensing capabilities.

Contents

Contents v

List of Tables vii

List of Figures viii

1 Introduction 1

1.1 Overview . . . 1
1.2 Connecting Commercial Battery-free Tags at Scale . . . 2
1.3 Software-defined Cooking . . . 3
1.4 Wireless Actuation for Electronics-free Soft Robots . . . 4
1.5 Speech Recognition Using Battery-free Temporary Tattoos . . . 5
1.6 Organization of Thesis . . . 5

2 Pushing the Range Limits of Battery-free Tags 6

2.1 Related Work . . . 8
2.2 Overview of PushID . . . 9
2.3 Blind Distributed Beamforming . . . 11
2.4 Synchronization for Distributed RFID Readers . . . 19
2.5 Implementation and Evaluation . . . 22
2.6 Results . . . 24
2.7 Summary . . . 29

3 Software Defined Cooking 30

3.1 Background on Microwave Ovens . . . 34
3.2 An Overview of SDC . . . 36
3.3 SDC's Heat Sensing . . . 38

3.4 SDC's Heat Actuation . . . 43
3.5 Implementation and Evaluation . . . 47
3.6 Results . . . 48
3.7 Limitations . . . 53
3.8 Discussion . . . 54
3.9 Summary . . . 55

4 Wireless Actuation for Electronics-free Soft Robots 57
4.1 Related Work . . . 61
4.2 Overview and Chapter Outline . . . 64
4.3 High-power Blind Beamforming . . . 64
4.4 Design of Wirelessly Actuated Electronics-free Soft Robots . . . 70
4.5 Implementation . . . 74
4.6 Results . . . 77
4.7 Discussion and Limitations . . . 85
4.8 Summary . . . 86

5 Speech Sensing Using Wirelessly Powered Tattoos 87
5.1 Related Work . . . 90
5.2 Overview of RFTattoo . . . 91
5.3 Processing RFTattoo Signals . . . 93
5.4 Tattoo Design and Fabrication . . . 97
5.5 From RF Signals to Speech . . . 101
5.6 Implementation . . . 104
5.7 Results . . . 106
5.8 Limitations . . . 112
5.9 Discussion . . . 113
5.10 Summary . . . 114

6 Conclusion and Future Directions 115
6.1 Conclusion . . . 115
6.2 Future Directions . . . 117

Bibliography 119

List of Tables

5.1 Phoneme to viseme clustering mapping [8] 101
5.2 Feature Set used by RFTattoo. 107

List of Figures

2.1 (a) In traditional multi-antenna beamforming, a beam forms towards a particular direction; (b) PushID's distributed MIMO beams create multiple energy peaks and troughs over space resulting in complex energy distributions. 7

2.2 Architecture of PushID 10

2.3 (a) Energy pattern when four transmitters focus beamforming at the center. (b) Energy pattern when we apply a phase shift ($\frac{2}{3}\pi$) to one of the transmitters. (c) Energy pattern when a strong reflector is placed along y-axis. 13

2.4 We simulate 10 transmitters deployed on a circle with a radius of 75 meters. Red points are the location of omni-directional transmitters. White points indicate the area that can activate the RFID tags (signal strength of energized area $\leq$ -12.8dBm), black points are the opposite. From left to right, the plots represent energized pattern with 1, 330 and 450 beamforming vectors. The corresponding energized area is 34%, 80%, 97% of the enclosed area of transmitters. 16

2.5 A reflected path can be modeled by a virtual source (r_1, θ_1, ψ_1) – the mirror image of the tag about the reflector. 17

2.6 PushID constructs a spanning tree of the RFID readers in a distributed manner to guide synchronization. 20

2.7 Depicts one round of PushID's search for the unknown tags. Without a tag response, PushID starts another round by applying the next beamforming vector B_{i+1}. 23

2.8 (a) **Outdoors**: We deploy PushID in a football field with one case of the transmitters placed as shown above. (b) **Indoors**: We deploy a four-transmitter based PushID on a floor (20 $\times$ 40 meters) of an office building covering multiple rooms and cubicles (c) Plots maximum reading range vs. # of readers. M-A MIMO – Multi-antenna MIMO. 23

2.9 Time and Frequency Synchronization Accuracy 24

2.10 Throughput vs. (a) Distance of closest reader; (b) number of readers. 27

2.11 (a) CDF of Throughput with changing tag orientation; (b) Distance of closest reader. 27

2.12 (a) % of tags covered; (b) CDF of throughput with mobility 28

3.1 Results of microwaving wet thermal paper (black spots indicate high heat). (Top) A traditional microwave without/with a turntable. A turntable can mitigate uneven heating, but cold/hot spots remain. (Bottom) SDC, for uniform heating (fewer black spots show heat is spread uniformly) or heating to write "MobiCom". We use patterned susceptors and let SDC make sure that the text area has been heated in hot spots. 31

3.2 SDC's Hardware. 32

3.3 SDC recipe is a progression of desired temperature vs. time per-pixel of food. 33

3.4 A turntable with 32 neon lights (left) and a plate cover with 32 neon lights (right). 38

3.5 The turntable inside a running oven. 39

3.6 We place neon lights with different lengths of wire extensions (red: 0mm, green: 7.5mm, and blue: 15mm) under the turntable. We then measure the percentage of glowing time to quantify the sensitivity. 40

3.7 The optical fiber carries the signal outside of the chamber. 42

3.8 Workflow of SDC's heat actuation: ① SDC first computes heating gap by comparing the desired heating pattern with the current status from the thermal camera. ② SDC then adjusts the actuation plan accordingly and ③ updates the distribution continuously. 44

3.9 Programmable microwave accessories. 46

3.10 We color rice grains with thermal-chromatic pigment, which turn pink in a predictable manner as their temperature increases. 48

3.11 Visualization of heating of rice as a function of time for no rotation, default rotation and SDC. SDC results in the most uniform heating. 49

3.12 (left) Temperature variance of SDC is low vs. baselines; (right) points on food closely follow heating recipe. 50

3.13 Left: An input recipe for stress test. Middle: SDC without susceptors. Right: SDC with susceptors. Susceptors can help build more skewed thermal distributions. 51

3.14 Mean (ΔP) and Standard Deviation (σ_P) of thermal delta for arbitrary heating. The final column ($\frac{\Delta P}{d}$) denotes the temperature gradient per unit distance that can be achieved. 51

3.15 The raw bacon, and slices of bacon cooked with SDC and the original turntable. These three slices of bacon are from the same package, so the original fat patterns are nearly identical. Heated meat and fat will shrink. SDC applies heat to meat and fat discriminatively, so the fat shrinks more than the meat. 52

3.16 Mean and Standard Deviation of thermal delta for cooking milk and meat. We don't have σ_P for milk since the thermal camera can only sense the surface temperature. 53

4.1 The soft robot can potentially be in the non-line-of-sight (NLoS) from the RF antennas due to various blockages, e.g., a plastic pipe. WASER estimates the wireless channel information of the soft robot in NLoS using the battery-free LED harvesters distributed outside the pipe. The entire system is enclosed by a flexible Faraday sleeve illustrated by the green color above. Sec. 4.5 describes the actual deployment of the WASER system. Sec. 4.2 describes the overview of WASER system. 58

4.2 WASER's Battery-free LED Harvesters. 69

4.3 (a) The cross-section of WASER's layered material structure. WASER uses a layered material structure to build the soft robot including a polymer layer and a conductive layer (see details about the materials in Sec. 4.4.2). 2.4 GHz wireless signals can induce heat using dielectric and induction heating for the polymer and conductive layer respectively. The conductive layer will generate more heat than the polymer layer due to the strong induced eddy current. Thus, the heat induced on the conductive layer will also further transfer to the polymer layer to facilitate the shape change of the polymer layer (LCE). Note that wireless signals at 2.4 GHz can penetrate through dielectric materials such as LCE. (b) The S11 frequency response of our soft robot design simulated by CST Studio. 71

4.4 (a) A rectangular sample with uniformly distributed liquid metal. (b) Red lines show the path of induced eddy currents on the rectangular sample. (c) Heat distribution captured by a thermal camera: the figure shows the heat distribution of the sample with uniformly distributed liquid metal when the 2.4 GHz wireless signals are applied. The brighter color shows a higher temperature. We note that edge effects occur at the corners and edges of the uniform liquid metal conductive layer. 72

4.5 (a) A sample with patterned/textural liquid metal layer. (b) The path of induced eddy currents on a patterned liquid metal layer. (c) The heat distribution of the patterned liquid metal layer shows that the edge effect is significantly mitigated, and the high temperature is uniformly distributed through the entire sample. 72

4.6 The WASER Platform. 75

4.7 If the robot is deployed in **non-metallic spaces**, e.g., a plastic pipe, the system can be placed within a flexible Faraday sleeve. (a) & (b) The flexible Faraday enclosure. The Faraday sleeve can be scaled to volumes as large as 100 cubic-feet. The camera views the LED harvesters inside of the enclosure via a small opening (6 mm wide) of the Faraday sleeve that results in minimal leakage. (c) Top view of the system setup inside the Faraday enclosure. The robot is deployed inside the plastic pipe. Four antennas and ten LED harvesters are placed outside the plastic pipe. 75

4.8 If the robot is deployed in **metallic spaces**, e.g., a metallic pipe, WASER reuses the existing metallic structures as RF shields. For example, a soft robot deployed to inspect metallic pipes does not require additional shields around the pipes themselves – only with minor modifications to the ends of the pipes covered by Faraday fabrics. The LED harvesters would need to be placed within the pipe with the camera viewing the lights from one end. 76

4.9 (a) Relative Positions: the antennas (orange) are placed with different orientations to provide diverse spatial diversity. The soft robot (silver) is placed at $(220, 110, 0)$ and LEDs (red) surround the soft robot with a spacing of 25 mm. (b & c) The Normal Mean Squared Error of WASER's channel probing. WASER's channel probing algorithm is resilient to various types of impairments, and the error converges to -4.5 dB after seven probes. Simulation-based estimation gives 1dB higher error than LED-based estimation. 78

4.10 Empirical NMSE of channel interpolation. (a) shows that the error converges to -7.7 dB for more than 7 points in simulation. (b) shows that the error converges to -7.9 dB for more than 10 points on our testbed. 78

4.11 WASER implements different textures on the conductive layer and evaluates their heat distributions. We show that the thin lamination design significantly mitigates the edge effects and achieves a uniform heat distribution. 80

4.12 Heating efficiency with different number of antennas N under total transmitted power 20W. We show that WASER's multi-antenna beamforming approach achieves a higher heating efficiency in the speed of heating soft robots as the number of antennas increases, and outperforms the baseline considering various locations of the soft robot in both LoS and NLoS. 81

4.13 (a) WASER achieves heating a soft robot to 60 °C at 1.5 seconds when applying a 20 W total transmitted power. (b) WASER is resilient to various materials of blockage. 81

4.14 (a) The theoretical heating resolution that WASER can achieve. We define the heating resolution as the distance that the energy density attenuates 3 dB from the peak. (b)&(c) We evaluate the heating resolution of WASER by focusing energy towards the two ends of a 45 mm long soft robot. From the experiments, we show a 17.8 mm heating resolution on average. (d) We evaluate the robot movement accuracy. Specifically, we use WASER to apply beamformed energy towards the three individual parts (leftmost, rightmost) of the soft robot. We show that the most heated point deviates from the target point by an average of 7mm, 4mm, respectively. 83

4.15 Snapshots of WASER Bending. 84

4.16 Snapshots of WASER Folding. 85

5.1 RFTattoo's Architecture: (1) measures the wireless channel of RFID tattoos; (2) infers stretch bits and tongue position based on the reflected signal power and phase of multiple three-chip RFID tattoos; (3) feeds the features into machine learning models to recognize corresponding facial gestures (viseme – see Table.5.1); images show the corresponding viseme from the GRID dataset [9]; (4) combines visemes to form meaningful words and sentences by natural language processing. 92

5.2 Simulated data from HFSS (upper) and real data collected by VNA (down) of the resonant frequency shift with stretch steps of 1mm. Red patch shows the ISM bandwidth from 902 to 928 MHz. An average of 8 MHz frequency shift is seen. After just 3.25 mm of stretch, the resonant frequency goes out of the ISM band. Even within the band, the gain drops up to 6 dB. 94

5.3 The i^{th} chip on an RFTattoo tag responds with a maximum signal power when the i^{th} bit of s is one, where s is tag stretch. For example, Chip 3 responds when the 3^{rd} bit of stretch is on which corresponds to the possible stretch of 4, 5, 6, 7 in the figure. 94

5.4 RFTattoo's Antenna Design: (1) Stretchable substrate. (2) Half-wavelength Dipole. (3) Impedance matching networks with a sets of high impedance modules. (4) Coupled spacing. 98

5.5 Under stretched RFTattoo sampled from scanning electron microscope. (a) 300× magnified. (b) 5000× magnified. Small cracks appear between the silver flakes embedded withing the elastomer composite when the tag has been stretched. 98

5.6 (a) Equivalent circuit of dipole antenna. (b) Equivalent circuit of multiple impedance matching networks composed of parallel inductively coupled loops. Overall, this figure shows that multiple parallel coupled loops intend to match with the impedance of feed terminals (chip) at different frequencies which in turn leads the chip to resonate at many frequencies (three frequencies shown above). 99

5.7 Samples of visemes [10]. 12 visemes are used to map to all possible phonemes including the rest pose. 101
5.8 (a) Rotary system for stretch evaluation of RFTattoo. (b) RFID Tattoo. 106
5.9 Unwrapped phase of the reflected signal from four RFID chips in different RFID tattoos. The phase relationship changes with different tongue positions due to the impedance variation. We feed this as one of the raw data into feature extraction. See details in Sec. 5.3.2 107
5.10 (a) Stretch Accuracy: RFTattoo achieves a 1.4 mm median accuracy in stretch inference (Facial muscle movement has an average stretch of 25 mm.) (b) Tongue positioning: RFTattoo achieves 92% mean accuracy in classification of five tongue positioning. Index 0 to 4 represent tongue positions at: normal resting position, upper jaw front, upper jaw back, lower jaw front, lower jaw back. (c) Confusion matrix of 11 visemes: RFTattoo achieves an average test accuracy of 90%. 108
5.11 (a) Accuracy in predicting viseme classes for selected classifiers: Random Forest, Decision Tree, Naive Bayes, kNN and SVM. (b) Confusion matrix of 100 words classification. RFTattoo achieves an average test accuracy of 86%. (c) Accuracy in sentence construction. 109
5.12 Accuracy (%) vs. (a) Vocabulary Size; (b) user's orientation; (c) number of reader antennas. . . 111

Chapter 1

Introduction

1.1 Overview

Over the past decades, research in wireless systems has extensively focused on building efficient communication [11, 12, 13, 14, 15, 16, 17, 18] and new sensing systems [19, 20, 21, 22, 23, 24, 25, 26] with many research innovations along the way through next-generation cellular, WiFi systems, radar-based human-computer interfaces. For example, Apple Inc. uses ultra-wideband for localizing battery-powered AirTag [27]; Google introduces a gesture recognition solution based on 60 GHz millimeter-wave [28] radar technology, to name a few.

Yet, another well-known capability of wireless signals has arguably received less attention: the ability to heat and, more generally, provide energy. For example, consider the microwave oven, a nearly 50-year-old invention that to this day continues to remain largely similar in capability, blasting energy blindly to re-heat food without advancing significantly in terms of fine-grained heating capabilities [29, 30, 31, 32]. Recent work attempted to control the heating pattern by deploying specialized radio-reflective stirrer blades or a customized turntable to distribute heat evenly [33, 34, 35]. However, despite these advances, undesired hot and cold spots within a microwave remain a known problem even in industry-grade microwaves. Similarly, a Radio-frequency identification (RFID) reader powers battery-free RFID tags by blindly broadcasting the wireless energy, resulting in a short communication range – less than 5-10 meters (for commercial RFID tags) [36, 37]. **Ideally, to realize fine-grained heating and power battery-free tags at scale, we need a technology that allows for precise and programmatic control of delivering wireless energy instead of blasting the energy blindly.**

Indeed, precise control of wireless signals, also known as *beamforming*, has been widely used in today's wireless networks to mitigate interference and improve network throughput [12, 14, 38, 39]. Beamforming

carefully coordinates multiple transmitters and focuses the transmitted signals toward a targeted receiving device rather than diffusing them in all directions. Multiple copies of the transmitted signals constructively interfere with each other at the receiving device. Thus, they provide a much higher signal power which improves the signal-to-noise ratio (SNR) at the receiver [40, 41].

To perform beamforming—for example, to a cellphone from multiple cellular towers—active feedback (e.g., a message of wireless channel information) from the cellphone is a crucial prerequisite so that cellular towers can use the channel reciprocity [15] to know where they need to beamform [42, 43, 44]. However, in the context of energy delivery, the feedback is completely absent from the target. For instance, a battery-free RFID tag has no energy to provide any feedback before the tag is beamed energy [45], and a pizza cannot tell the microwave oven where its hot and cold spots are. Thus, today's beamforming techniques have not directly translated into innovation in wireless systems for heating and delivering energy.

This dissertation resolves this dilemma by developing a suite of *distributed blind beamforming* systems that deliver precisely controllable wireless energy from multiple *distributed* transmitters to battery-free devices and objects *without* any active feedback from them (*blind*). These systems facilitate the connectivity of extremely low-power Internet-of-Things (IoT) devices and enable new wireless heating applications such as software-defined cooking and soft robots actuation. Besides these, we develop novel sensing platforms through co-optimizing soft materials and wireless energy-harvesting systems.

The rest of this dissertation describes multiple contributions in three directions: 1) Connecting Commercial Battery-free Tags at Scale (Chapter 2); 2) Software-defined Cooking and Beyond (Chapter 3 and 4); 3) Novel Sensing Systems through Co-optimizing Soft Materials and Wirelessly Powered Systems (Chapter 5). The following sections briefly describes these systems and present their core contributions.

1.2 Connecting Commercial Battery-free Tags at Scale

This first part of the dissertation describes our core contribution on beamforming commercial battery-free wireless devices without any active feedback from them. It is well-known that the communication range of battery-free wireless systems is limited [36, 37]. This is especially true of low-cost battery-free IoT devices, such as RFID and NFC tags, that billions of them have been deployed in warehouses and airports [46, 47, 48]. However, today's RFID reader blasts wireless energy blindly and waits for a response from the tag. As a result, the tag can only be detected in the vicinity of an RFID reader (5–15 meters). The first part of this dissertation introduces *distributed blind beamforming* algorithms to extend the communication range of battery-free devices.

A long-held precept concerning battery-free devices, like RFID tags, is that the communication range

between the devices and the energy source (RFID readers) is limited to within 5–15 meters at best [36, 37]. PushID [3] breaks this barrier. It allows commercially available and unmodified passive RFID tags to respond from eight times farther than commercial RFID systems. Our initial observation while tackling this problem was that the sinusoidal waves transmitted by multiple distributed RFID readers could combine constructively and destructively in the air. PushID coordinates all distributed individual RFID readers to beamform the signal coherently to provide sufficient energy at targeted RFID tags. Traditional approaches to beamform targeted devices involve channel reciprocity, which requires the response of the passive RFID tags in the first place. However, this is a chicken-or-egg problem in the context of passive RFIDs since the tags cannot respond before they are beamformed from multiple distributed RFID readers.

Overcome the "Blind" Situation for the Battery-free Tags: PushID resolved this dilemma by designing a novel distributed blind beamforming technique to discover all the passive tags within a targeted region without a single response from the tag in the first place. Instead of applying various beamforming vectors randomly, PushID finds the smallest number of beamforming vectors to provide sufficient energy that can activate the tag to every geographic points in the desired area. By doing this, PushID formulates an energy cover problem and uses a gradient based approach to reduce the search space of the beamforming vectors. Overall, PushID's maximum communication range is 64 meters, an improvement of eight times compared to commercial RFID systems.

1.3 Software-defined Cooking

Similar to how blasting wireless energy blindly from commercial RFID readers results in degrading the range performance of battery-free IoT systems, the microwave oven re-heats foods randomly, resulting in a non-uniform and unpredictable heat distribution. In contrast, this second part of this dissertation presents a closed-loop heating machine to heat the food with a software-defined thermal trajectory and enables a fine-grained spatial resolution [5].

Microwave ovens blast wireless energy blindly to re-heat food, which results in random cold and hot spots on foods such as pizza [49]. While multiple solutions attempted to distribute heat evenly, undesired hot and cold spots within a microwave remain a known problem even in industry-grade microwaves. We built Software-defined Cooking (SDC) [MobiCom 2019], a novel low-cost closed-loop system that can augment existing consumer microwaves to sense and control heating at a fine-grained resolution.

Overcome the "Blind" Situation for a Pizza: At the heart of SDC is a novel approach to sensing and analyzing the constructive and destructive interference of the microwave energy inside the microwave oven. The key idea is to deploy microwave-safe passive electronics that senses the underlying physical

phenomenon producing heat, i.e. the electromagnetic field. Specifically, we deploy neon lights that harvest EM energy in the microwave to glow in proportion to its intensity. We then measure the intensity of the light to infer field intensity and estimate the temperature distribution that would result. We do this via optical fibers that carry the visible light signal outside the chamber, allowing SDC to measure light signals from lights that are hidden from view.

Upon sensing the energy distribution, SDC finely controls heat according to the software-defined thermal plan using a programmable approach. Overall, SDC improves the thermal heating uniformity by 633% compared to commercial microwaves and creates an arbitrary temperature delta of 183°C with a resolution of 3 cm. SDC explored a new category of wireless systems—systems that actively use wireless signals to heat objects of interest in a fine-grained controllable way—and was awarded Research Highlights in SIGMOBILE and Communications of ACM.

1.4 Wireless Actuation for Electronics-free Soft Robots

We found that using a similar wireless heating concept from SDC is a promising option for addressing a new research domain: moving or actuating thermally driven soft robots. Existing actuation methods for these soft robots are rudimentary; they require an external DC power supply, laser beams, or a manually controlled heat gun [50, 51, 52, 53, 54]. Although laser beams provide high-resolution projections onto the soft robot [55, 56, 57, 58, 59, 60], they suffer from obstacles blocking the laser beam from reaching the robot. In contrast, this dissertation presents WASER, a novel approach to actuate soft robots with high precision, even in the non-line-of-sight radio environment.

However, wireless channel feedback from the electronics-free soft robots is unavailable, meaning that it is challenging to accurately actuate the desired spots of them via RF beamforming. The situation gets even worse when the robot operates in non-line-of-sight scenarios.

Overcome the "Blind" Situation for an Electronics-free Robot: To address this challenge, we introduce a high-power *blind* beamforming algorithm that can precisely actuate the soft robot without any response of the robot. Specifically, WASER achieves this by (1) acquiring accurate estimates of channel information indirectly from multiple desired locations using battery-free LED harvesters surrounding the robot; (2) interpolating the channel distribution to estimate the wireless channel at the soft robot.

We build a prototype of a multi-antenna high-power beamforming platform for actuating battery-free soft robots. Our experiments show that WASER can heat the targeted region from 20°C to 60°C within 1.5 seconds.

1.5 Speech Recognition Using Battery-free Temporary Tattoos

While there is a lot of literature on assistive input-to-speech technology for people with speech impairments both in computer vision and HCI community [61, 62, 63], there remains a gap for an everyday intuitive assistive technology for speech impairments that do not require external infrastructures like cameras [64, 65], bulky sensors (EEG) [66, 67] on the face or manual hand input. This dissertation presents RFTattoo [UbiComp 2020], the first wireless speech recognition platform for voice impairments through skin-friendly, wafer-thin, stretchable, and battery-free RFID tattoos. This work is a collaboration with mechanical engineers and materials scientists to produce these batteryless tags so that they are stretchable, flexible, wafer-thin, extremely light, and made with hypoallergenic materials. Each RFID tattoo is attached to the skin surface of the face at known locations. RFTattoo tracks the strain on individual tattoos over time as they deform in response to facial motions generated by different intended sounds. However, it is often the case that certain distinct sounds produce similar facial movements. To this end, RFTattoo natural language processing models that combine identified facial gestures in context to construct meaningful words and sentences. A detailed user study with 10 users reveals 86% accuracy in recognizing the top 100 words in the English language. RFTattoo introduced a new bridge between two different areas: materials science and battery-free IoT systems and was awarded the Best Long Wearables Paper in UbiComp / ISWC 2020.

1.6 Organization of Thesis

The rest of this thesis is organized as follows. Chapter 2 presents PushID and describe how it extends the communication range of commercial batteryless RFID systems in greater detail. The chapter presents the first study of the benefits of applying distributed beamforming techniques on today's wireless energy delivery infrastructure for batteryless IoT. Chapter 3 presents Software-defined Cooking, the first closed-loop microwave heating system that can augment existing consumer microwaves to sense and control heat at a fine-grained resolution. The chapter presents a novel approach to enable programmable fine-grained food heating using a commercial microwave. In Chapter 4, we describes a high-power distributed beamforming system to actuate the fine-grained movement of an electronics-free soft robot. The chapter details our algorithm and presents experimental results of WASER. Chapter 5 presents a novel sensing system that we build on top of the wireless energy delivery system. This chapter describes soft materials that can harvest wireless energy and be capable of speech recognition.

Chapter 2

Pushing the Range Limits of Battery-free Tags

Conventionally, passive commercial RFID tags have a maximum range of about 5-15 meters. Passive RFIDs are limited in range owing to their limited cost, form-factor, and the FCC-mandated power limits of the RFID readers they harvest energy from. Indeed, in much of today's factories [46] and warehouses [47, 48], RFID-tagged products can only be detected around specific checkpoints in the vicinity of an RFID reader, and are virtually undetectable at other points in between [68, 69]. Further, recent innovation on RFID-based localization and sensing [25, 26, 1, 70] remain constrained to a few meters around each reader in these large spaces.

In this chapter, we ask – "Can we push the range limits of today's commercial passive RFID tags without increasing the prevailing density of deployment of RFID readers?". In particular, we aim to do so without modifying the RFID tags in any way by adding to cost and complexity [71, 72, 73], relays [74] or requiring batteries [75]. We further avoid using sophisticated multi-antenna or directional RFID readers that can indeed expand range [76, 77], yet are vulnerable to obstacles and limited by FCC transmit power limits [78]. More importantly, such systems require commercial RFIDs that are linearly polarized to be carefully oriented towards their location [79] to harvest sufficient energy.

We present PushID, the first system that uses distributed MIMO to increase the communication range of commercial passive RFID tags. PushID synchronizes both transmissions and receptions from multiple, distributed RFID readers to beamform power to RFID tags that are several tens of meters away from any individual reader antenna. In addition, PushID exploits the diversity in location and polarization of reader antennas to further improve range. Our detailed experimental evaluation on an eight-distributed multi-antenna reader testbed reveals that PushID achieves a 7.4$\times$ improvement in range compared to the state-of-the-art commercial readers and 20% over the mean of distance improvement even when compared to 8-antenna MIMO reader arrays [77], all while remaining compliant with FCC power limits for the

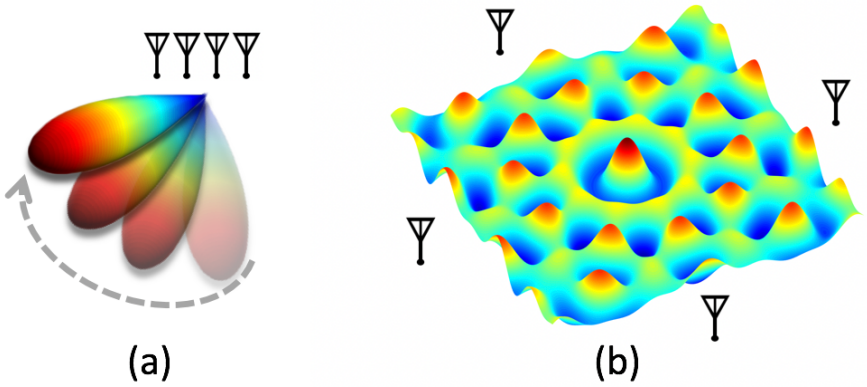

Figure 2.1: (a) In traditional multi-antenna beamforming, a beam forms towards a particular direction; (b) PushID's distributed MIMO beams create multiple energy peaks and troughs over space resulting in complex energy distributions.

readers.

PushID's main goal is to find optimal beamforming weights to beam power to an RFID tag at an unknown location and orientation. Since RFID systems are back-scatter [80], optimal beamforming weights amplify *both* the transmitted and received signals to/from RFID tags. At first blush, one may consider using channel reciprocity [15], where one infers the optimal beamforming weights based on the wireless channels of signals from the RFID tag to the reader antennas. However, in the context of passive RFIDs, this leads to a chicken-or-egg problem. To emit a response, a passive RFID tag needs to harvest sufficient energy from the beamformed signal of the RFID reader antennas. Yet to perform accurate beamforming, the readers need a response from the RFID tag in the first place. Indeed, naively iterating over all possible beamforming weights would take days to simply beamform power to one tag from a long distance.

PushID resolves this dilemma by developing a novel distributed blind beamforming approach to efficiently search through the space of beamforming weights without a response from the tag. At a high level, PushID models the received signal power at each point in space for different beamforming weights applied across reader antennas. It then identifies subsets of 3-D space where RFIDs placed would receive sufficient energy to respond. Unlike traditional beams of a directional antenna, in the context of distributed MIMO, these regions of space that receive sufficient energy are quite complex and they span the entire 3-D space as shown in Fig. 2.1. PushID shows that finding the optimal beamforming weights while minimizing overlap is analogous to the weighted set-cover problem, a well-known NP-complete problem [81]. PushID then develops heuristic approximation algorithms to efficiently search the entire space for RFID tags under a limited time budget.

Limitations: We note that PushID has two important limitations common to RFID-systems: (1) First, despite significant range improvements, a small fraction of RFID tags ($< 5\%$) are missed due to extreme

shadowing or poor orientation; (2) Second, while PushID can handle modest mobility of tags (walking speeds), it struggles at higher speeds due to high dynamism in the multipath characteristics of the tags.

We implement PushID on eight USRP N210 software radio readers, each connected to separate Jackson Lab Fury clocks and commercial Alien passive RFID tags. We perform our experiments on a 140 x 140 meter outdoor space and a 20 x 40 meter indoor space in both line-of-sight and non-line-of-sight scenarios. Our experimental results reveal that:

- PushID achieves a maximum communication range of 64 m, an improvement of 7.4× that of commercial RFIDs and 20% over expensive 8-antenna MIMO.
- Our system detects over 95% of the tags in a 140× 140 m area, while commercial readers can detect tags no further than 8.5 m at best.

Contributions: To our knowledge, PushID [3] presents the first distributed MIMO system to power commercial passive RFID tags. We present a novel blind distributed beamforming algorithm to efficiently search through the space of beamforming weights as well as novel phase synchronization for RFIDs. A detailed prototype evaluation on an eight-antenna distributed reader testbed reveals a 7.4× improvement in range compared to state-of-the-art commercial readers.

2.1 Related Work

2.1.1 Wireless Power Transfer

Recent advances in wireless power transfer take two approaches: non-radiative near-field coupling and far-field RF radiation. Near-field coupling [82, 83] uses multiple coils to charge low-power devices, yet is restricted to several tens of centimeters.

Far-field radiation promises longer distance for wireless power transfer on the order of meters through innovative hardware design. Energy-harvesting RFID tags such as the WISP [84] promise a range of 18 meters. Recent work [85, 86] demonstrates distances of kilometers by building custom backscatter tags with small batteries. Low-power WAN (LP-WAN) technologies [87, 88] including LoRa [89], SIGFOX [90] have explored battery-powered tags that can communicate over miles. In contrast, PushID strives to improve the range of current battery-free commercial RFID tags that cost a few cents by innovating at the RFID readers.

2.1.2 Blind Beamforming

Blind beamforming [91] is a set of theoretical beamforming solutions developed by the signal processing community in the presence of poor quality or even no channel state information available from the clients. Past work on blind beamforming with weak channel responses leverages its statistical features, for example, the cyclostationary property [92] and spectral self-coherence [93]. In contrast, this dissertation does not assume any channel information to leverage, given that the targets (e.g., battery-less tags) are not a priori powered up. Other work on blind beamforming with zero channel feedback builds solutions without the need for carrier frequency synchronization [94, 95, 76]. PushID builds upon this past work but overcomes unique system-level challenges pertaining to the RFID context: (1) It accounts for feedback from neighboring RFID tags that are charged during the exploration of beamforming weights to refine the search in multipath-rich settings; (2) It deals with various synchronization challenges in distributed beamforming with long RFID packets.

2.1.3 Connecting and Sensing RFIDs

There has been much past research on RFID tags including ways to mitigate collisions [96], improve communication speed [97] and design a variety of localization and sensing solutions [**?**, 23]. However, all these solutions are limited to a range of at most 5-15 meters [36, 37] around the readers and thus have limited ability to locate, sense or communicate with RFID tags. Closely related to our system are recent solutions that use multi-antenna arrays connected to RFID readers to improve range [76, 77]. While these systems improve range to commercial RFID tags (distances of up to 38 meters), our system varies in two ways: (1) First, systems with directional antennas are vulnerable to obstacles between the reader antennas and tags which can significantly attenuate the signal; (2) Second, they fail when the RFID tags are oriented poorly relative to the reader antennas. We see our system as an alternative approach using diversity of spatially distributed antennas within RFID ISM frequencies to extend the range of RFID and as complementary to [76]. We do think that both PushID and multi-antenna MIMO [98] as attractive solutions with different deployment cost/requirements. We show how by synchronizing signals across multiple RFID readers, PushID exploits the diversity in location and orientations of the reader antennas to significantly improve range compared to the state-of-the-art.

2.2 Overview of PushID

PushID aims to power and communicate with commercial passive RFID tags via RF-backscatter from a team of distributed commercial RFID readers, where tags are beyond the communication range of any

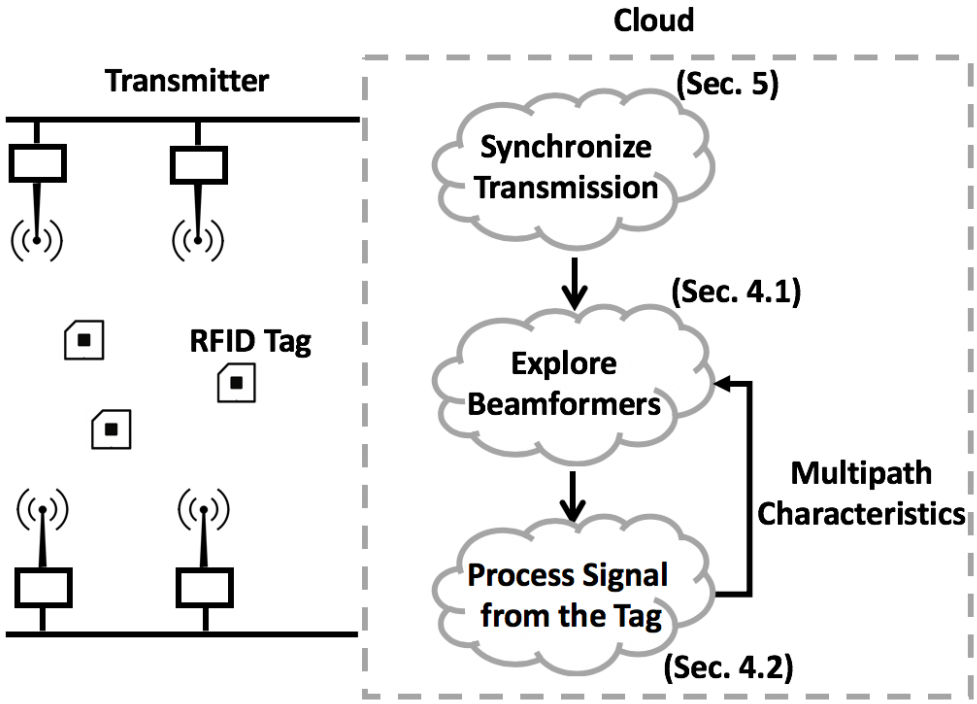

Figure 2.2: Architecture of PushID

single reader. PushID achieves this by coherently combining signals across distributed RFID readers to maximize received signal power at one or more RFID tags, whose location and orientation are a priori unknown. We note that since RFID systems are back-scatter, PushID applies beamforming weights *both* on the transmitted and received signals to maximize energy to/from RFID tags. We assume all PushID RFID readers are connected to a wired backhaul, which allows them to coordinate transmissions and data that needs to be transmitted on the downlink. We further assume that while the locations of the RFID readers are known, the number of RFID tags, their locations and their environment are unknown.

At a high level, PushID's system design is as follows (see Fig. 2.2): All RFID readers time and phase synchronize their transmissions on the air and iteratively apply various beamforming vectors in the hope of receiving responses from RFID tags in the environment. The readers then collect responses from various RFID tags in the environment and use the wireless channels they perceive in improving the search for other tags. The readers continue this process until they believe (with sufficient confidence) that they have covered the entire desired coverage area. To achieve this, PushID optimizes the following (related) properties: (1) Maximize the total number of RFIDs found within the coverage area under an overall time budget (which limits the number of beamforming vectors that can be iteratively attempted); (2) Maximize the throughput of signals from each tag.

The rest of the chapter addresses the key challenges in designing the two main aspects of PushID's

architecture:

(1) Searching through the Beamforming Space: First, our system needs to identify optimal beamforming weights to iteratively search over in order to power RFID tags in the entire space. While this problem can be trivially addressed using channel reciprocity for RFID tags that are in range, the key challenge is that tags are outside the coverage area of any single reader. Therefore, PushID must identify the smallest set of beamforming weights that can provide sufficient energy to all tags over the entire area of interest. The key to this is to effectively model multipath in the environment, which would change the set of beamforming weights to search over. Sec. 2.3 describes our approach.

(2) Synchronizing Distributed RFID Readers: Second, PushID should efficiently synchronize RFID readers that are spatially distributed indoors, without a shared clock between them. In contrast to past work in the Wi-Fi domain [99], our key challenge comes from the longer duration of RFID transmissions, over which packets can quickly lose synchronization in phase. Further, RFID transmissions, unlike Wi-Fi, are narrowband, which makes time synchronization challenging as well. Sec. 2.4 addresses these challenges.

2.3 Blind Distributed Beamforming

This section describes how PushID enables a team of RFID readers with an arbitrary geometry to find the optimal beamforming weights and beam power to all RFID tags in their coverage area, including those beyond the range of any single reader. We aim to achieve this without any response from these RFID tags to begin with or prior knowledge of their locations and orientations. For ease of exposition, this section assumes that all RFID readers experience no time, carrier frequency and phase offsets. We will explicitly deal with synchronizing distributed RFID readers in Sec. 2.4.

We make a few key observations about the scope and goals of PushID's approach:

- **Beamforming on both Downlink and Uplink:** We note that PushID seeks to amplify the received signal power from RFID readers to the tags and vice-versa. Specifically, since RFIDs operate on RF-backscatter, and owing to channel reciprocity [100], beamforming weights used on the transmit chain to power tags on the downlink can also be used to amplify their received signals on the uplink. For simplicity, the rest of this section discusses PushID in the context of maximizing energy on the downlink.
- **How much energy to beamform?:** We emphasize that the goal of PushID in this section is to simply beamform enough energy to detect an RFID tag. Once the RFID tag's response is received, we can then use the reciprocal channel to obtain the optimal beamforming vector to maximize data rate to

that tag in future transmissions. As a result, the rest of this section will favor PushID formulations that maximize the energy RFIDs require over the entire space, as opposed to focusing on individual tags.

2.3.1 Exploring the Beamforming Space

A naïve approach to perform blind beamforming is an exhaustive search of the space of beamforming vectors in hope of eliciting a response. Prior work in the context of cellular multi-antenna arrays [101] constructs codebooks that progressively steer the beam along various discrete directions in hope of covering an entire cell efficiently. For example, consider a phased array of antennas (see Fig. 2.1(a)) where a transmitter could simply beam power iteratively along discrete angles to cover the space of interest.

However, the distributed RFID context makes such an elegant design challenging. First, given that RFID readers are widely separated and they form an arbitrary geometry relative to each other, beamforming weights distribute energy over the space in very complex patterns. Further, it is challenging to find weights that both cover the entire space of interest with sufficient energy and minimize overlap.

What do beamforming energy patterns look like?: To better understand how different beamforming weights from a distributed array of antennas impact the distribution of energy over the space, we perform a simple simulation. We consider four transmitters in the corners of a square with a one meter diagonal length. For simplicity, we consider that the transmitters are in 2-D free space and use standard wireless channel models [102]. We first apply a beamforming weight that allows signals from the transmitters to add up coherently at the center of the square. We plot the distribution of energy over the entire 2-D space around the square encompassing the transmitters in Fig. 2.3 (a). We notice that while the center of the square indeed receives maximum energy (denoted by bright yellow), there are multiple spots around the center that are also energized with a similar received signal strength. This shows that applying beamforming weights in a distributed array also focuses energy on unintended points in the space. This means that simply iterating beamforming weights to focus on individual points in the space would lead to much unwanted overlap and be grossly inefficient. We therefore conclude that:

Observation 1: PushID must seek to minimize overlap between energy patterns of beamforming weights it applies.

Problem Formulation and Optimization: Based on the above observation, we will now formulate PushID's core optimization problem that seeks to find the smallest group of beamforming vectors which energizes the entire space of interest with minimum overlap between them. At least to begin with, our system cannot rely on any feedback from RFID tags in the environment, given that none of them may have sufficient

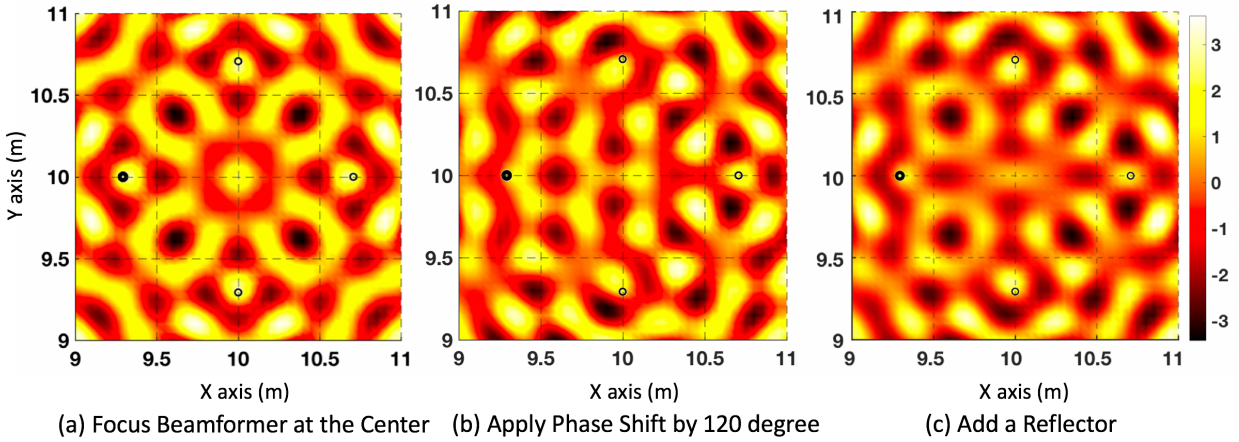

Figure 2.3: (a) Energy pattern when four transmitters focus beamforming at the center. (b) Energy pattern when we apply a phase shift ($\frac{2}{3}\pi$) to one of the transmitters. (c) Energy pattern when a strong reflector is placed along y-axis.

energy to respond. As a consequence, we have no prior information on the nature and extent of multipath in the environment. The rest of this section therefore assumes that line-of-sight paths to RFID readers dominate all other paths, and we explicitly account for multipath in Sec. 2.3.2.

At a high level, our approach shows that choosing the optimal set of beamforming vectors is analogous to a well known combinatorial problem: the weighted set coverage problem. To see how, let us imagine that the 3-D space is divided into a grid of discrete blocks. Each beamforming vector effectively supplies sufficient energy to RFIDs in some subset of these blocks. Our goal is to find the smallest set of such beamforming vectors whose union is the universal set of all blocks in the grid. This is analogous to the weighted set cover problem, which seeks to find the smallest number of sets, each containing a few integers in the range $1, \ldots, N$ whose union is precisely the universal set $\{1, 2, \ldots, N\}$. Given that weighted set cover is NP-complete [103], we propose an efficient approximation algorithm, while presenting various optimization to reduce algorithmic complexity.

Mathematically, let us assume that the space of beamforming vectors has n discrete elements $\mathbf{B} = \{B_1, \ldots, B_n\}$, and we aim to cover m discrete points in the space spanning the desired coverage area denoted by $\mathbf{G} = \{G_1, \ldots, G_m\}$. Let the variable u_{ij}, $i = 1, \ldots, n$ and $j = 1, \ldots, m$ be one if the beamforming vector B_i provides energy to the point in space G_j. Given that we assume the L base station locations are known and no information on multipath is available (we discuss multipath in Sec. 2.3.2), we can determine

u_{ij} as follows:

$$u_{ij} = \begin{cases} 1, |B_i.h_j|^2 > \tau \\ 0, \quad \text{otherwise} \end{cases} \tag{2.1}$$

$$\text{where} \quad h_j = \left[\frac{1}{d_{lj}} e^{-2\pi\sqrt{-1}\frac{d_{lj}}{\lambda}}, l = 1, \ldots, L \right] \tag{2.2}$$

Where h_j is the vector of wireless channels from base stations to grid point j, λ is the wavelength, d_{lj} is the distance between the l^{th} base station and j^{th} grid point and τ is the minimum received energy required to energize an RFID tag.

Our objective is to find the smallest set of beamforming vectors that spans all m points in $\mathbf{G}$. We can state this mathematically as the following integer linear program based on the variable x_i which is 1 if and only if the i^{th} beamforming vector is included in our optimal set:

$$\min \quad \sum_{i=1}^{n} x_i$$

$$s.t. \sum_{i=1}^{n} x_i u_{ij} \geq 1 \quad \forall j \in \{1, ..., m\} \tag{2.3}$$

$$x_i \in \{0, 1\} \quad \forall i \in \{1, ..., n\} \tag{2.4}$$

The above formulation directly resembles the well-known weighted set cover problem [103], which processes a group of sets to find the smallest sub-collection whose union is also the union of the original group of sets. While this problem is known to be NP-complete [104], a reasonable polynomial-time approximation algorithm is to relax the above integer-program formulation into a linear program (LP). Specifically, we replace Eqn. 2.4 above by the following:

$$0 \leq x_i \leq 1 \quad \forall i \in \{1, ..., n\} \tag{2.5}$$

We solve the optimization problem using standard linear programming [105] to obtain the optimal set $\{x_1^*, \ldots, x_n^*\}$. We then output the chosen set of beamforming weights by applying randomized rounding [106] on the beamforming weights. This technique interprets the fractional part of the solution to the linear program as a probability distribution and then selects a solution by sampling this distribution. Randomized rounding is known to return a set of beamforming weights that is a valid set cover with probability at least $1/2$ [106]. Mathematically, to bound the probability, let ρ be a constant that satisfies: $e^{-\rho \ log \ n} \leq \frac{1}{4n}$. Then randomized rounding proceeds for exactly $\rho \ log \ n$ iterations, and in each iteration, it picks i^{th} beamforming weight with probability dictated by its respective solution x_i to the linear program.

Prior work has shown that the above approximation algorithm results in a set of beamforming vectors whose size is within a factor $O(\log n)$ of the optimum [107]. Our implementation uses the Ellipsoid

LP-solver [108] with a worst-case complexity of $O(n^4)$, where n is the number of discrete beamforming vectors PushID's algorithm optimizes over.

Reducing Complexity and Search Space: To reduce the complexity of PushID, one must actively seek to reduce n, the number of beamforming weights that PushID considers in its optimization. Our key insight to this end is that while a large number of beamforming vectors are available, not all are created equally. To see why, let us revisit our example of the energy pattern from a beamforming vector that focuses energy at the center of a square in Fig. 2.3(a). We now slightly perturb the beamforming vector of one of the transmitters by choosing one of the transmitters and adding $\frac{2}{3}\pi$ to its phase, and we plot the updated energy pattern in Fig. 2.3(b). We make two observations about the resulting energy pattern. First, each local maximum of energy moves in different, complex ways. This is precisely why we need the optimization algorithm above to minimize overlap. Second, the size of each energized region changes with maximum diffusing energy over wider spots. In practical terms, this means that the same amount of energy is spread out over a wider area than the previous case. Spreading energy over a wider space is good in that RFID tags over a wider region can be covered by a single beamforming weight. Yet, spreading energy over too thin over a wide area is likely to make the energy per unit area insufficient to activate tags. Indeed, the most ideal beamforming weights are those who diffuse energy in a Goldilocks zone between these two extremes. We observe the following:

Observation 2: PushID must favor beamforming weights with maximal total area where RFID tags remain powered.

PushID therefore aims to search over beamforming weights that meet the above criterion of maximizing area-of-coverage for RFID tags. Our approach begins with n beamforming weights chosen randomly, where n is dictated by available computing power. For each beamforming weight, we make incremental phase shifts and measure the gradient of net increase in coverage area. We then apply a gradient-based optimization that favors phase shifts which maximize coverage area. We implement Adadelta [109] to speed up the learning rate. The below algorithm summarizes our approach.

We then feed the above n set of beamforming weights into our optimization algorithm to find the optimal set.

Design Decisions: We emphasize a few key design choices:

(1) How finely to divide space?: First, PushID must choose discrete points $\mathbf{G}$ in the space to capture the area covered by a beamforming vector. Choosing too few would lead to coverage holes, while choosing too many would waste computation. PushID therefore samples the space at an interval empirically measured to be below the minimum distance between two adjacent energized regions across beamforming vectors

Algorithm 1 Gradient-Based Beamforming Vector Pruning

1: **B** : random beamforming vectors. $t = 1, ..., N$. D_i: energy of the set of points which could be covered by the i-th beamforming vector B_i.

Loop:

2: $Q \leftarrow |\bigcup_i^n D_i|$, where Q represents the number of points covered in **G** by the beamforming vector set **B**

3: $g^{(t)} \leftarrow \nabla_{\mathbf{B}}(Q)$

4: $E[g^2]^{(t)} \leftarrow \gamma E[g^2]^{(t-1)} + (1-\gamma)g^{(t)^2}$

5: $\Delta\mathbf{B}^{(t)} \leftarrow -\frac{RMS[\Delta\mathbf{B}]^{(t-1)}}{RMS[g]^{(t)}} g^{(t)}$

6: $\mathbf{B} \leftarrow \mathbf{B} + \Delta\mathbf{B}$

while $Q < threshold$

return $arg\ max_{\mathbf{B}}\ Q$

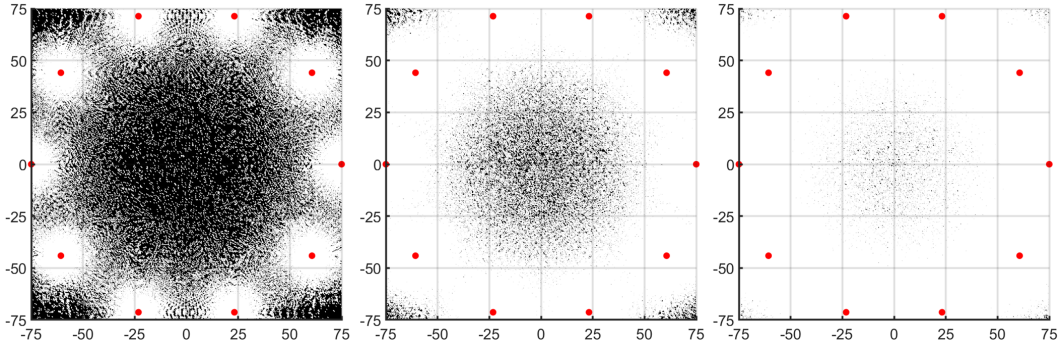

Figure 2.4: We simulate 10 transmitters deployed on a circle with a radius of 75 meters. Red points are the location of omni-directional transmitters. White points indicate the area that can activate the RFID tags (signal strength of energized area $\leq$ -12.8dBm), black points are the opposite. From left to right, the plots represent energized pattern with 1, 330 and 450 beamforming vectors. The corresponding energized area is 34%, 80%, 97% of the enclosed area of transmitters.

in **B**. We empirically find that this corresponds to a sampling distance of $\lambda/3$ in our experiments.

(2) Picking the energy threshold: PushID chooses the threshold τ empirically by measuring the smallest amount of energy needed for an RFID tag to respond at its smallest data rate. Once the RFID tag is detected, future transmissions can use the reciprocal channel measurements from this tag to speed up data rates [102]. We note that τ must be calibrated conservatively to support all RFID tag models in the space.

(3) Impact of Orientation: PushID explicitly accounts for the reader's antenna gain and polarization across spatial directions by applying a weight to each term of Eqn. 2.2: α_{lj}, which captures the attenuation in the l^{th} base station antenna when it faces the j^{th} grid point. We also account for the orientation of the tag by setting τ conservatively to the smallest amount of energy for a tag to respond should it be oriented most unfavorably relative to the readers.

(4) Run Time: We find that PushID's run-time is primarily bottlenecked by the slow beamforming switch time of transmitters (4.5 ms for our hardware) as opposed to computation. PushID's run time depends

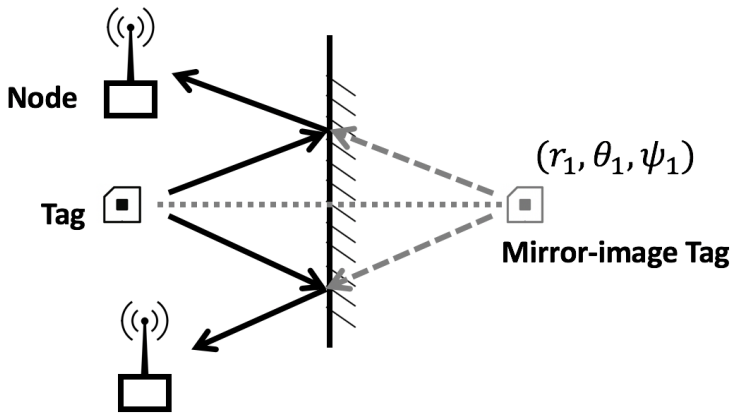

Figure 2.5: A reflected path can be modeled by a virtual source (r_1, θ_1, ψ_1) – the mirror image of the tag about the reflector.

directly on the final number of beamforming vectors PushID must iterate over. This depends on the size of the space, placement of base stations and multipath. To get a sense for expected run time, we simulate ten RFID readers in a circle of radius 75 m (Fig. 2.4). We consider a threshold of RSSI>-12.8 dBm for the RFID tags to respond. Without PushID, we find that only 33% of the total area of interest is covered. However, after only 450 iterations we find that nearly 97% of the area is covered by PushID. This maps to a total time of 2.0 s for 97% coverage. Sec. 2.6.6 discusses these tradeoffs in experiments (we observe $\approx$ 4 s run time due to multipath and change in layout).

(5) Adapting to New Information: Our approach so far arrives at a static set of beamforming weights. However PushID can benefit from the new information in the channel response of RFID tags, as they are detected. In particular, we are interested in learning about the extent and nature of multipath, which can impact the optimal set of beamforming vectors. Sec. 2.3.2 below deals with this explicitly.

2.3.2 Accounting for Multipath

While so far our discussion considers free space, the presence of multipath can considerably change the set of beamforming weights to efficiently search over a given area. To see why, we revisit our example in Fig. 2.3(a), and this time we add a strong reflector and re-evaluate the energized regions of space as shown in Fig. 2.3(c). We notice that the resulting energy heatmap varies considerably from the free-space heatmap, both in the number, size and placement of the hotspots. Note that the same set of reflectors can influence the energy perceived at different tags differently. As a result, we conclude:

Observation 3: PushID must account for multipath given its impact on the optimal set of beamforming vectors.

PushID's high-level approach to do so uses the responses from RFID tags in the environment that are progressively detected. Indeed, in the absence of any response, PushID has no information about multipath to work with, and therefore assumes a free-space channel. As responses from RFID tags are collected, PushID progressively computes information of the location and orientation of dominant reflecting surfaces in the environment. It then uses this information to update its optimization algorithm, specifically, it modifies the energy patterns corresponding to our beamforming weights.

Finding Dominant Reflectors: To compute the location of dominant reflectors, PushID adapts the MUSIC algorithm while accounting for the arbitrary geometry of the RFID readers. Specifically, the algorithm takes as input wireless channels along the various frequencies of operation of an RFID tag (note that RFID tags naturally hop through a range of frequencies in the 900 MHz ISM band). It then measures the polar coordinates: (r, θ, ϕ) representing the mirror-image of the RFID tag along dominant reflectors by measuring the power of the received signal $P(r, \theta, \phi)$ of signals received from these coordinates. Mathematically, we write:

$$P(r,\theta,\phi) = \frac{1}{|a(r,\theta,\phi)^\dagger E_n E_n^\dagger a(r,\theta,\phi)|} \tag{2.6}$$
$$\text{where: } a(r,\theta,\phi) = [e^{4\pi j|r-r_i|cos(\theta-\alpha_i)cos(\phi-\beta_i)/\lambda}]_{i=1,\ldots,N}$$

Where (r_i, α_i, β_i) are the polar coordinates of the transmitters, λ is the signal wavelength, j is the square root of -1, E_n are the noise eigenvectors of $h_{obs}h_{obs}^\dagger$, h_{obs} represents the vector of observed wireless channels of the tags and $(.)^\dagger$ is the conjugate-transpose operator. Our algorithm computes the top-s ($s = 5$ in our implementation) local maxima in $P(r, \theta, \phi)$ to define the set of polar reflector coordinates: $\{(r_k, \theta_k, \phi_k), k = 1, \ldots, s\}$.

Folding Multipath into the Optimization: At this point, we model how the energy patterns of beamforming weights change due to our knowledge of reflecting surfaces. We use a ray-tracing model [102] to account for how multipath changes received signal power. Mathematically, we rewrite Eqn. 2.2 in the definition of u_{ij} as:

$$h_j = \left[\sum_{k=1}^{s} 1/d_{ljk} e^{-2\pi\sqrt{-1}d_{ljk}/\lambda}, l = 1, \ldots, L\right] \tag{2.7}$$

Where λ denotes the wavelength and d_{ljk} represents the distance traversed by the ray emanating from base station l to grid point j when reflecting off reflector at coordinates (r_k, θ_k, ϕ_k). This formulation effectively removes the free-space assumption in our optimization to explicitly account for ambient reflectors.

Accounting for past vectors: We note that as new information about multipath emerges, one must account for how this impacts the coverage area of beamforming vectors used previously and therefore invoke the

optimization to fill gaps in coverage. Mathematically, let us denote x_i^* as an indicator function on which beamforming weights were used previously. Then we can rewrite Eqn. 2.3 as:

$$\sum_{i=1}^{n} x_i u_{ij} \geq 1 - \sum_{i=1}^{n} x_i^* u_{ij} \quad \forall j \in \{1, ..., m\} \tag{2.8}$$

Modeling fleeting and small reflectors: While the above formulation assumes reflectors impact all RFIDs in the coverage area equally, in practice, this may not be the case. Specifically, reflectors have a higher probability of impacting nearby RFIDs compared to RFID tags that are further away. Similarly, reflectors that were computed in the past may no longer exist at the same location and orientation in the future. To account for these effects, PushID employs the exponential weighting method [110] to progressively reduce the contribution of reflectors to the optimization with increasing distance from the reflector or time elapsed since detection. Specifically, we re-write Eqn. 2.7 as:

$$h_j = \left[\sum_{k=1}^{s} w_k / d_{ljk} e^{-2\pi\sqrt{-1} d_{ljk}/\lambda}, l = 1, \ldots, L \right] \tag{2.9}$$

Where $w_k = f_1^{d_{ljk}} f_2^{t_k}$ and $f_1, f_2 < 1$ are constants (empirically set to 0.9 in our implementation) and t_k is the time elapsed since the measurement of reflector k was made.

2.4 Synchronization for Distributed RFID Readers

In this section, we consider a classic challenge for distributed MIMO systems: accurate time and frequency synchronization with multiple distributed RFID readers for coherent beamforming. In particular, we actively estimate and correct for carrier frequency and timing offsets, which would otherwise cause transmissions across RFID readers to often combine incoherently. We build upon the classic distributed MIMO architecture used in the context of Wi-Fi [99, 15] while accounting for new challenges in the RFID context.

Quick Primer on Distributed MIMO: At a high level, past distributed MIMO systems for Wi-Fi [99, 15] use a master-slave architecture where multiple slave transmitters attempt to transmit in-phase with a master transmitter. Prior to transmitting each data packet, the master sends a short beacon containing a known preamble. Slave transmitters estimate their phase relative to this beacon to account for frequency offsets between the master's clock and their own clocks. In addition, slaves exploit the relatively wide bandwidth of Wi-Fi to estimate phase shifts due to timing offsets. Slaves then apply phase shifts compensating for these offsets when they transmit data packets in-tandem with the master. Of course, during the data packet transmission itself, small additional phase drifts can accumulate owing to residual time

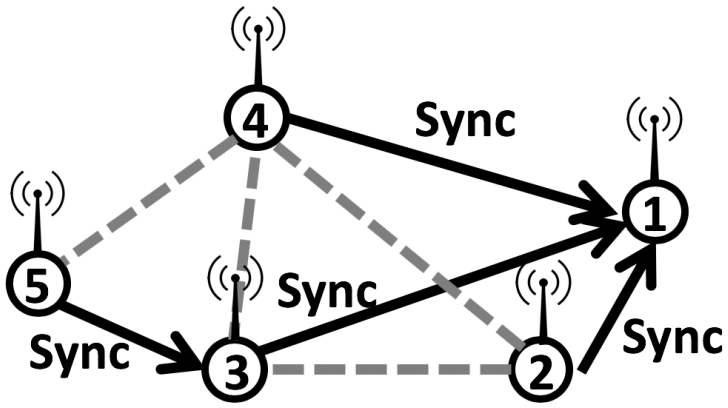

Figure 2.6: PushID constructs a spanning tree of the RFID readers in a distributed manner to guide synchronization.

and frequency offsets. As the duration of packets is short, the slope of such phase drifts can be readily corrected for.

Challenges in the RFID context: RFIDs bring two important challenges for distributed MIMO: (1) First, RFID packets last for a much longer time ($\sim$100$\times$) than Wi-Fi packets [99], ensuring greater phase drift due to frequency offset. This is because tags need to harvest enough energy in order to respond to the readers' queries, and this takes more time as the distance between the RFID tags and readers increases. (2) Second, RFID transmissions are narrowband (20 kHz), meaning that resolving timing offsets is extremely challenging.[1] The rest of this section tackles each of these challenges.

2.4.1 Frequency Offset Compensation

PushID's key idea to compensate for frequency offsets leverages the full-duplex nature of RFID readers. Specifically, each PushID slave cancels out its own signal to recover the signal from the master reader. By measuring how the phase of this node drifts over time, PushID can correct for phase drifts that accumulate since the initial synchronization.

Correcting for Drift: For simplicity, let us begin with the case of two RFID readers – one master and one slave reader. To achieve that, initial synchronization in phase, we borrow from MegaMIMO [99]'s phase synchronization protocol where the RFID slave reader applies an initial phase shift to synchronize with the master (see Algm. 2). However, as frequency offset accumulates over time, the readers will notice that the phase of the master drifts.

[1]Past proposals for Wi-Fi like SourceSync [111] cannot be directly used in the narrow-band RFID context, as they assume availability of phase measurements over wide bandwidths (20 MHz) to estimate timing offsets.

PushID forms a closed loop to compensate for residual frequency offsets. Specifically, the slave RFID reader observes the change in phase from the master over a time interval $(t, t + \Delta t)$, where t is the most recent time of synchronization and Δt is the time elapsed since then. Should we observe a change in phase of the master transmitter since then, i.e. $phase(t + T) - phase(t)$, the slave applies the negative value of this phase offset to its own transmission. Note that this would, in effect, remove residual phase errors, allowing the two readers to combine their signals coherently at RFID tags. Further, note that due to channel reciprocity, the same phase shifts would ensure coherent combining on both the downlink and uplink.

Algorithm 2 Initial Frequency Synchronization

1: f_M: master's oscillator frequency, f_{S_i}: slave i oscillator frequency.
2: $\mathbf{H(t)} = (h_1 e^{j2\pi(f_M - f_{S_1})t}, h_2 e^{j2\pi(f_M - f_{S_2})t})^T$
3: Decompose $\mathbf{H(t)}$ we have $\mathbf{H(t)} = \mathbf{R(t)HT(t)} = \begin{pmatrix} e^{-j2\pi f_{S_1} t} & 0 \\ 0 & e^{-j2\pi f_{S_2} t} \end{pmatrix} \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} e^{j2\pi f_M t}$
4: $\mathbf{H(t)} = e^{j2\pi f_M t}\mathbf{R(t)HT(t)}e^{-j2\pi f_M t} = \begin{pmatrix} e^{j2\pi(f_M - f_{S_1})t} & 0 \\ 0 & e^{j2\pi(f_M - f_{S_2})t} \end{pmatrix} \mathbf{H} = \mathbf{A}\,\mathbf{H}$
5: Apply frequency compensation $\mathbf{H}^{-1}\mathbf{A}^{-1}$ at slaves.

Scaling the system: A key challenge, however, is scaling the above system beyond two readers. This is because upon canceling one's own transmission, each reader would perceive a linear combination of all other readers in its vicinity and not that of the master alone. This means that should a phase drift occur simultaneously for multiple readers (which is likely), these readers would be misled by out-of-sync transmissions from the others among them. To make matters worse, some readers may be beyond the communication range of the master and therefore may not be in a position to synchronize directly with the master as the last resort.

To mitigate this problem, we seek to assign to each RFID reader a unique reference reader to which it may synchronize, should the master not be within its vicinity. Specifically, we assign indexes $1, \ldots, n$ to each RFID reader, where RFID reader 1 denotes the master and all other readers denote slaves. We assume that these indexes are known a priori by the readers and can be constantly broadcasted so that any reader that fails can be removed from consideration in the optimization. Each RFID reader aims to synchronize its phase relative to the neighbor with the smallest index. In effect, our synchronization scheme creates a spanning tree of RFID readers (see Fig. 2.6), provided the graph of all readers is a connected one, which we assume. This spanning tree is designed to ensure that all readers eventually synchronizes to the root – the master RFID reader.

At this point each slave RFID reader subtracts its own signal and tracks the phase of the remaining

linear combination. Should this change beyond a threshold, the RFID reader requests all its children and descendents in the spanning tree to cease transmission and then attempts to re-synchronize its transmission with the remaining active readers. We note that the master RFID reader never stops its transmission. It is easy to see that this scheme ensures that all RFID readers eventually transmit in-sync with the master.

2.4.2 Time Synchronization

PushID performs a two-step time synchronization process, a coarse synchronization to align symbols and a fine-synchronization that leverages the phase of signals from the master across frequency.

Coarse Synchronization: PushID synchronizes slaves with the master using the known preamble of the *Query* command transmitted by the reader to initialize an inventory. To maximize time resolution, slaves receive this signal from the master at a high sampling rate. Slaves then apply correlation with the known preamble to obtain the index of the master's signal. We then only consider correlation coefficients above a threshold to reject outliers (we reject the bottom 6%). We repeat this process over five preambles and choose the result with the maximum correlation coefficient.

Fine Synchronization: To compensate for drift in timing offsets, PushID exploits the phase of signals from the master RFID reader a function of frequency. Specifically, recall that RFID transmissions hop between a wide range of frequencies in the 900 MHz ISM band spanning a total of 26 MHz. Each slave RFID reader estimates the phase of signals, having subtracted its own signal, across frequencies. PushID then monitors for any change in the slope of the phase of this signal across frequencies between measurements. Specifically, recall that any time offset of Δt between the two readers results in a frequency-dependent phase shift of $\Delta\phi = 2\pi f \Delta t$. As a result, PushID can estimate timing drifts by applying a least-squares linear regression [112] of ϕ as a function of t and obtain the resulting slope m. Any drift in timing offset can simply be computed as $m/2\pi$ and corrected for. We note that our system can scale akin to frequency offset compensation above, for more than two readers in the network. Specifically, when any RFID reader goes out-of-sync in time, it informs all its descendents to stop transmitting before attempting to re-synchronize.

2.5 Implementation and Evaluation

We implement PushID on a testbed of USRP N210 software radios with SBX/WBX daughterboards operating as RFID readers. We feed an omnidirectional and planar antenna to the antenna ports of each node for full-duplex use. All our readers are SISO, unless specified otherwise. Each USRP connects with an independent Jacksonlab Fury clock which could lead frequency and timing offset among the nodes. At

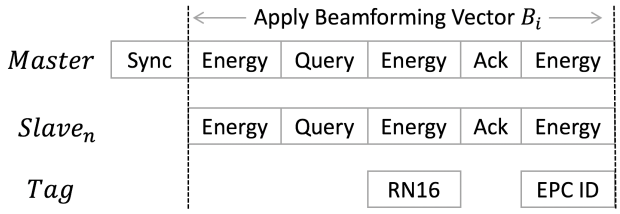

Figure 2.7: Depicts one round of PushID's search for the unknown tags. Without a tag response, PushID starts another round by applying the next beamforming vector B_{i+1}.

(a) American Football Field (b) Office Building (c) Reading distance vs. # of nodes

Figure 2.8: (a) **Outdoors**: We deploy PushID in a football field with one case of the transmitters placed as shown above. (b) **Indoors**: We deploy a four-transmitter based PushID on a floor (20 $\times$ 40 meters) of an office building covering multiple rooms and cubicles (c) Plots maximum reading range vs. # of readers. M-A MIMO – Multi-antenna MIMO.

the backend, each USRP is connected via Ethernet cables to a 64-bit Dell computer running Ubuntu 16.04. We also assume dedicated socket-based TCP connections between the reader nodes. Our RFID tags are commercial passive Alien Squiggle RFID tags. We measure a maximum range of up to 8.7 meters with our setup using one reader antenna.

PushID software: PushID is fully implemented in UHD/C++ including beamforming and distributed synchronization. In addition, we use an in-house UHD/Gnuradio based C++ RFID emulator to decode signals from the tags. We implement the set-cover based optimization in the cloud on a cluster of 64-bit core i7 Ubuntu machines and report the optimal beamforming weights to the reader nodes.

We ensure that all PushID RFID readers implement ASK modulation with PIE encoding to align with the specifications of the Gen2 RFID protocol. Apart from delivering energy, PushID readers also actively transmit messages which specify the tag's modulation format, encoding scheme and backscatter frequency. The RFID tags in our experiments employ ASK modulation and FM0 encoding. The protocol flow of PushID is shown in Fig. 2.7.

Testbed: We evaluate PushID in two testbeds indoors and outdoors. *(a) Outdoor Testbed*: We deploy PushID around a football field (140 $\times$ 140 meters) with 8 transmitters. *(b) Indoor Testbed*: We deploy four-transmitter based PushID across a floor (20 $\times$ 40 meters) of an office building covering multiple rooms and cubicles. We note both testbeds are multipath rich due to stands/partitions in the former and

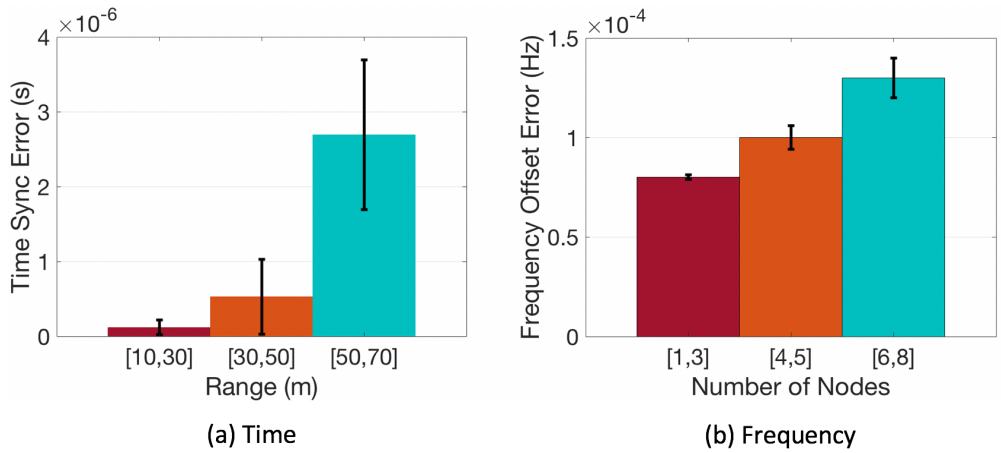

Figure 2.9: Time and Frequency Synchronization Accuracy

cubicles/furniture in the latter blocking the direct path of some readers. We mount readers in various positions including different elevations. We put RFID tags in various positions and orientations that face towards different angles. Fig. 2.8(a) and (b) shows the candidate locations of RFID tags (represented by blue dots) and readers (represented by orange dots). We note that unless specified otherwise, all results incorporate an equal amount of data points (over 1000 RFID tag locations considered among them) from both testbeds with our core results evaluating how system accuracy changes in line-of-sight vs. non-line-of-sight relative to all readers.

Baseline: We compare PushID against two competing schemes: (1) *Closest Reader:* We assume that each reader independently decodes signals and each tag receives energy from the closest reader; (2) *Multi-antenna MIMO:* We assume that all reader antennas are co-located and synchronized by an external clock [77]. We note that unless specified otherwise, error bars in graphs denote standard deviation.

2.6 Results

2.6.1 Synchronization Accuracy

We evaluate the accuracy of PushID in achieving accurate frequency and time synchronization between base stations.

Method: We consider a testbed of up to eight USRP N210-based RFID readers, one designated as the master and use PushID to synchronize the slaves to the master at high accuracy. The RFID readers are placed in various arbitrarily chosen geometries and different relative distances between the slaves and master reader. We measure two quantities of interest: (1) The error in time synchronization; (2) The error in frequency synchronization of signals at a USRP N210 receiver that compares the phase of the signals

received from the master and slave(s) post PushID's synchronization. We perform this experiment in both our indoor and outdoor testbed in which about half of the slaves on average are in non-line-of-sight relative to the master and some slaves (>50% of nodes >50 m away from master) synchronize via multiple hops using PushID's spanning tree approach.

Results: Fig. 2.9(a) shows the mean and standard deviation (error bars) in accuracy of time synchronization for different ranges of distance between a slave RFID reader and the master. We note that, as expected, the mean error with increasing distance also increases: PushID achieves a mean error of 0.12, 0.53, 2.69 μs with the range of 10 to 30, 30 to 50 and 50 to 70 meters. However, we note that even the worst-case error is much smaller than one Nyquist time sample, given the narrow bandwidth of RFIDs (20 kHz). This means that PushID achieves the required level of time-synchronization accuracy to perform efficient distributed MIMO.

Next, Fig. 2.9(b) depicts the accuracy in frequency synchronization with increase in the number of RFID readers in the network. We find as expected the mean and variance of phase synchronization errors increase modestly as more readers join the network. PushID achieves a mean error of 0.0001 Hz in frequency offset overall across experiments. We note that this error corresponds to a phase shift of 0.0005° over the duration of a typical RFID packet and therefore minimally impacts the throughput of PushID's distributed MIMO architecture, as observed in Sec. 2.6.3.

2.6.2 Range vs. Number of Nodes

In this experiment, we evaluate the maximum distance for a number of slaves that can detect the response from the tag.

Method: We deploy up to eight RFID readers in various geometries (starting from co-located and with progressively increasing spacing) with tags placed in both involving line-of-sight and non-line-of-sight relative to the reader. Note that neither the location or number of RFID tags placed in the environment are known a priori to the readers. We consider, in aggregate over 1000 RFID tag locations across experiments. Across experiments, we note the distance between the RFID tag and its closest reader. Our experiments consider distances of up to 80 meters. Our goal is to estimate the maximum distance at which an RFID tag can be detected at the readers.

Results: Fig. 2.8(c) shows the maximum range of RFID tags with respect to the number of reader nodes (up to eight). As expected, PushID amplifies the received signal power from RFID readers to the tags and vice-versa, the PushID's range increases quasi-linearly as the number of RFID readers increases. We note that the rate of increase does dip (gradually) with increasing number of readers due to the increasing

impact of time and frequency synchronization errors as reported in Sec. 2.6.1 above. We further notice a surprising decreasing trend in the standard deviation with maximum distance. We find that this stems from the robustness of PushID to orientation in the presence of multiple distributed readers all oriented in diverse spatial directions. Our results show that PushID achieves a maximum range to an RFID tag of 64 meters, a gain of about 7.4× vs. commercial RFID and 20% over multi-antenna MIMO. We also note that the standard deviation of multi-antenna MIMO is large in various experimental settings. In contrast, PushID has better resilience and stability across experiments which gives more spatial diversity that benefits the poor polarization sensitivity of the RFID dipole antennas. We also notice that sometimes the multi-antenna MIMO has better performance than PushID when tag's orientation favors the direction of co-located readers in the LOS setting.

2.6.3 Throughput vs. Distance and Scale

In this experiment, we evaluate the impact of PushID on the throughput as we vary the number of RFID readers and the distance between the tag and its closest reader.

Method: We measure the throughput by first measuring the SNR of each RFID tag measured from the eight RFID readers after coherently combining signals to and from the tag. We then adapt the ESNR metric [113] to the RFID context to estimate the maximum data rate achievable for the received SNR. Note that once PushID's algorithm is applied to detect a tag, we use channel reciprocity to maximize throughput to each detected tag in this experiment. We deploy PushID in both the outdoor (Fig. 2.8(a)) and indoor scenario (Fig. 2.8(b)) and consider tags in both line-of-sight and non-line-of-sight relative to the readers. Our RFID tags use FM0 modulation which allows for data rates over 45 kbps.

Throughput vs. Distance: Fig. 2.10a shows the increase in throughput vs. distance in line-of-sight and non-line-of-sight settings and compares it against the baseline system that connects to the closest RFID reader. We observe that for the baseline, as expected, a reader has a maximum range of about 8.7 m across both line-of-sight and non-line-of-sight settings and we see that performance drops to zero throughput beyond this distance. PushID, with a 67.5 meter maximum range on average outperforms the baseline significantly in line-of-sight and 58.9 meters in non-line-of-sight. Further, as expected the throughput of PushID drops down as the distance increases due to lower signal-to-noise ratio. We note that quite significantly, PushID's performance with eight transmitters increases the throughput of RFID tags 2.6 × when compared to the baseline closest-RFID reader scheme at its maximum range of about 8.7 m.

Throughput vs. Number of Reader Nodes: As expected, with the increasing number of reader nodes we observe a gradual (logarithmic) increase in network throughput of covered RFID tags on average for

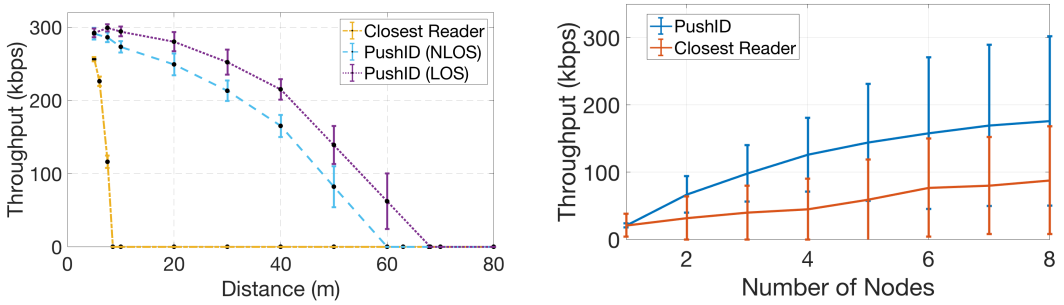

(a) Throughput vs Distance (b) Throughput vs Number of Nodes

Figure 2.10: Throughput vs. (a) Distance of closest reader; (b) number of readers.

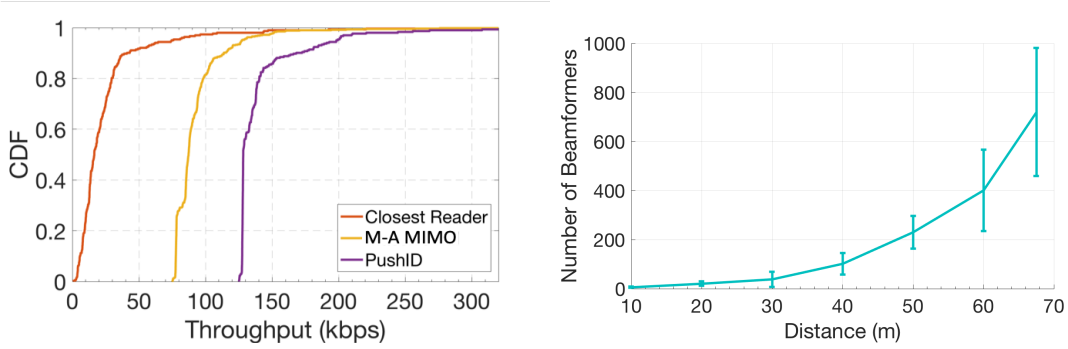

(a) CDF of Throughput (b) # of beamformers vs. distance

Figure 2.11: (a) CDF of Throughput with changing tag orientation; (b) Distance of closest reader.

PushID (see Fig. 2.10b). There is a similar, although much more modest increase with reader nodes for the baseline owing to an increase in coverage area. However, our system observes a net mean throughput gain of 2.6× over a network of 8-nodes over the baseline closest-reader system.

2.6.4 Impact of Orientation

Method: In this experiment, we model the distribution of the throughput of an RFID tag progressively oriented along various directions in 100 locations with 8 readers and compare three schemes: (1) PushID; (2) A 8-antenna MIMO scheme; (3) The closest reader baseline.

Results: Fig. 2.11a plots the CDF of throughput across schemes with changing tag orientation. We observe that PushID outperforms both multi-antenna MIMO (by 1.54× median) and the closest reader baseline (by 7.4× median). This is because, as explained in Sec. 2.3.1, PushID readers are oriented variously and therefore much more robust to change in orientations of the tags. In contrast, multi-antenna MIMO with reader antennas co-located loses performance when tags are oriented away from the MIMO reader, unlike PushID.

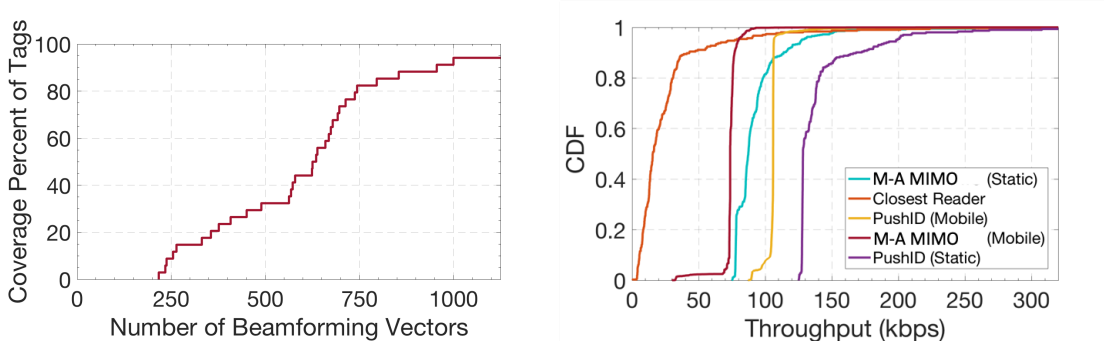

(a) coverage vs. # of beamformers

(b) Effect of Mobility

Figure 2.12: (a) % of tags covered; (b) CDF of throughput with mobility

2.6.5 Impact of Mobility

Method: We deploy two tags in the environment in our indoor testbed, each placed initially at the same 100 randomly chosen initial locations at varying distances (up to 60 m) from the closest reader: (1) static RFID tags; and (2) an RFID tag moved around at walking speeds by volunteers. We measure the performance of PushID 8-antenna MIMO scheme and our closest reader baseline from eight RFID readers and compare performance.

Results: Fig. 2.12b plots the CDF of throughput across schemes for static and mobile tags. As expected, throughput dips in the presence of mobility across schemes. However, we note that PushID achieves gains over the baseline despite mobility (1.6× for static and 1.4× for mobile over multi-antenna MIMO). PushID's robustness to mobility stems from two reasons: (1) The mobility of the RFID tag ensures that the tag is highly likely to move out of coverage holes. As a result, during its trajectory, PushID has a higher likelihood of detecting the tag, compared to a static tag. This counteracts to help recover some of the loss in performance owing to changing multipath in PushID's algorithms. (2) Once the tag is first detected, PushID can use channel reciprocity to rapidly continue beamforming to the tag and thereby respond to its mobility. We however highlight that mobility at very high speeds would significantly deteriorate PushID's performance and gains, just as it would deteriorate commercial RFID systems.

2.6.6 Convergence and Coverage

Method: In this experiment, we measure the convergence time of PushID's algorithm and how it is impacted by the distance of RFID tags and its trade-off with total area covered. We note that the initial set of beamforming vectors PushID uses can be found offline and future updates take minimal time overhead, PushID's main computational bottleneck is the rate at which beamforming weights can be applied by the USRP hardware which is 4.5 milliseconds in our testbed. We therefore measure convergence time in terms

of the number of beamforming vectors that needs to be applied. Once again, we consider RFID tags at a wide range of distances to the closest reader with eight RFID readers and run PushID.

Results: Fig. 2.12a shows that the percentage of tags discovered increases with increasing number of beamformers for distances from the closest reader over 60 m. We emphasize here that it is rare for our system to detect all tags, since some tags remain virtually undetectable due to their location, orientation or shadowing – a natural limitation of our system and indeed most wireless systems (for e.g. even the best cellular networks have deadspots [114]). Beyond 95% coverage, we see diminishing returns upon applying more beamforming vectors. Fig. 2.11b measures the trade-off between the distance of the tag and the number of beamformers needed to find $>$ 95% of tags in the area. We find that in the worst case at maximum distance, we need 980 beamformers (4.4 seconds for a USRP N210).

2.7 Summary

This chapter presents PushID, the first distributed MIMO system to power commercial passive RFID tags. PushID develops a blind distributed beamforming algorithm to efficiently search through the space of beamforming vectors. It further develops a novel phase synchronization algorithm to synchronize distributed RFIDs. A detailed prototype evaluation on an eight-antenna distributed reader testbed reveals a 7.4$\times$ improvement in range compared to state-of-the-art commercial readers. While this work focuses on using existing commercial tags without modifications, we believe future work needs to explore algorithms that innovate on radio design and light-weight computation on the tags themselves to improve PushID's performance.

Chapter 3

Software Defined Cooking

In the previous chapter, we present a distributed RFID system that extends the range of the commercial RFIDs. The key to this long-range wireless energy delivery is to precisely control the beamforming pattern – maximizing the constructive interference and minimizing the destructive interference throughout the entire space. This chapter presents another example of blindly blasting wireless energy from a microwave oven and demonstrates a re-designed microwave system to heat food in a programmable and fine-grained thermal trajectory.

Since their introduction to the consumer market in the 1970s, microwaves have seen widespread adoption and are today the third most popular domestic food heating method (after baking and grilling) [115]. Indeed, the original patents for the microwave by Raytheon Inc. in the late 1940s envisioned a universal food cooking instrument for all kinds of food ranging from meat to fish [116, 117]. While microwaves have indeed revolutionized the kitchen since their inception, today's consumer microwaves are mainly used as a blunt heating instrument (e.g., reheating pizzas) rather than a precise cooking equipment (e.g. cooking steak). The potential of microwave cooking is limited by the fact that today's microwaves heat food blindly, resulting in a non-uniform and unpredictable heating distribution [49].

State-of-the-art solutions to remedy the microwave have made advances in both how food heating is sensed and controlled. Past work has used thermal cameras just outside the microwave chamber [118, 119] to map the heat distribution within a microwave, albeit at limited resolution and refresh rates [120]. Microwaves also deploy specialized radio-reflective stirrer blades or turntables to attempt to evenly distribute heat. Yet, despite these advances, undesired hot and cold spots within a microwave remain a known problem even in industry-grade microwaves [33] (see examples in Fig. 3.1). Indeed, the fundamental challenge in building a more precise microwave for cooking is the fact that modern electronics to both sense and control microwave heating are not inherently microwave safe. The resulting need to isolate the microwave

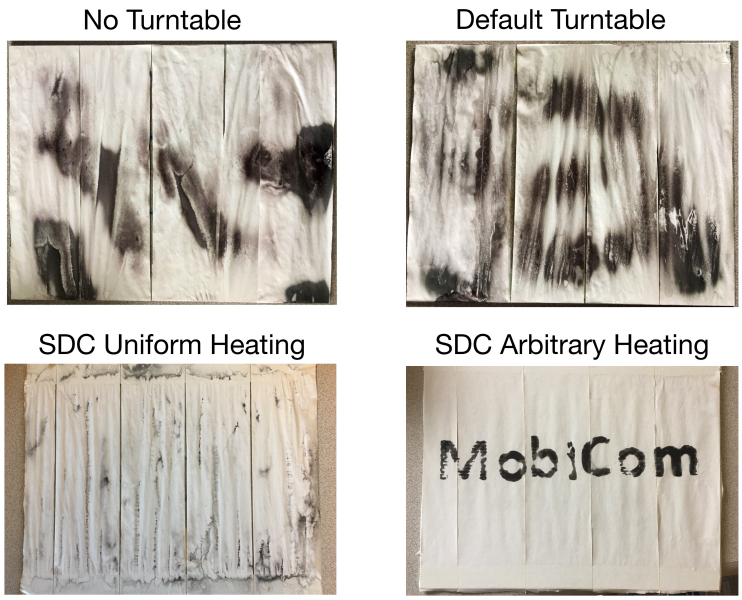

Figure 3.1: Results of microwaving wet thermal paper (black spots indicate high heat). (Top) A traditional microwave without/with a turntable. A turntable can mitigate uneven heating, but cold/hot spots remain. (Bottom) SDC, for uniform heating (fewer black spots show heat is spread uniformly) or heating to write "MobiCom". We use patterned susceptors and let SDC make sure that the text area has been heated in hot spots.

control system from the chamber itself reduces the granularity of both heat sensing and actuation.

This chapter presents SDC (software-defined cooking), a novel low-cost closed-loop system that can augment existing consumer microwaves to sense and control heating at a fine-grained resolution, all using microwave-safe components within the chamber. SDC's design can unlock numerous programmable heating opportunities (see Fig. 3.1). For example, when microwaving liquids (e.g. milk, baby formula), one need not worry about the uneven heating that may scald the mouth or destroy nutrients – a reason why microwaves are never advised to heat formula despite their convenience. Further, SDC can enable fine-grained forms of cooking, such as a computer generated Maillard reaction [121] patterns that heats the food's surface discriminatively and occurs when searing meats and pan-frying vegetables.

At the heart of SDC is a novel approach both to sense and control heat at different points in space within the microwave chamber. SDC senses heat using the phenomenon that produces heat in the first place: the electromagnetic (EM) field. SDC aims to measure the amplitude of the EM field at any given point. While battery-powered sensors exist to sense both heat and EM fields, SDC must do so using only microwave-safe components which excludes typical commercial batteries [31].

SDC achieves this by relying on the fact that the EM-field within the microwave is a natural source of energy. This means that one can simply power the sensor of the EM-field by the EM-field itself. SDC uses tiny RF-powered neon lights that glow in response to the EM-field within the microwave. Specifically, the

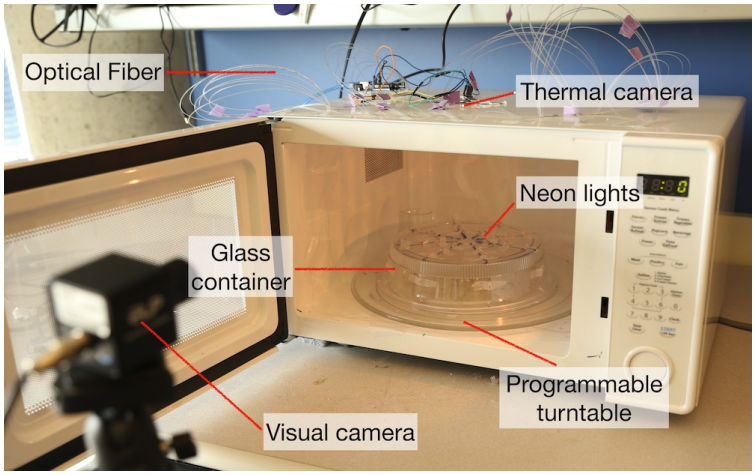

Figure 3.2: SDC's Hardware.

oscillating microwave results in a potential difference (of a few 100 V to a few kV) between two electrodes within each light bulb. Due to the potential difference, electrons are accelerated away from the cathode and give rise to collisions with the neon gas atoms or molecules, which will emit a characteristic glow in proportion to the amplitude of the field. Neon lights are inexpensive, compact and produce minimal disruption to the EM-field itself – meaning that they can be tightly packed at key locations around the chamber to sense the EM-field amplitude at high accuracy. Given that the neon lights may be obstructed from view due to the food placed in the microwave, we run optical fibers made of microwave-safe glass that carry the light signals outside the chamber to be sensed by a camera (Fig. 3.2).

Upon sensing food heating, SDC controls heat according to the user-specified thermal plan by building microwave shields that protect regions of the food that must not be over-cooked (e.g., meat). SDC achieves this through small metallic spheres placed within the microwave at key locations. While conventional wisdom says that one must not place metal in a microwave, RF-propagation is more nuanced. Specifically, metallic surfaces within the microwave only produce energetic sparks at sharp edges, found in most kitchen utensils and bowls. Metallic spheres by definition do not have edges and are thus microwave-safe [122]. SDC carefully packs metallic-spheres at specific regions of the microwave to minimize RF energy-transfer at these regions.

We implement a prototype of SDC by modifying a commercial Microwave oven (Sharp SMC1441CW). We place neon lamp arrays inside the microwave oven cavity and use a camera outside the cavity to monitor lamp flashes conducted via fiber-optic cables. We replace the coarse turntable motor with a step motor controllable via an Arduino board. Given a desired heating distribution pattern, SDC recommends the initial position where the user should place the food. During the heating, SDC continuously senses the real-time EM filed strength around the food and adjusts the actuation plan. Figure 3.2 illustrates the basic

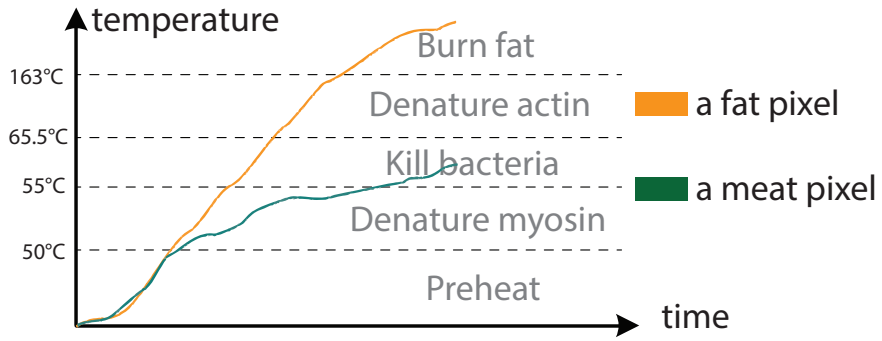

Figure 3.3: SDC recipe is a progression of desired temperature vs. time per-pixel of food.

hardware setup of SDC. We conducted detailed experiments to evaluate SDC's sensing, uniform heating, and planned heating capabilities. Our experiments reveal that:

- SDC can improve the thermal heating uniformity by 633% compared to commercial microwaves with turntables.
- SDC can create an arbitrary temperature delta of 183°C with a resolution of 3 cm.
- We demonstrate SDC in performing two unconventional cooking tasks: cooking bacon and warming milk for an infant.

Contributions: SDC is a novel redesign of the microwave oven that both senses and actuates the EM-field at fine-grained spatial resolution. SDC introduces programmable RF-powered neon lights whose signals are conducted by microwave-safe optic fibers to sense the EM-field distribution. SDC then adjusts the spatial heat distribution within the microwave chamber by moving food carefully around hot and cold spots pre-designed using microwave-safe accessories. A prototype implementation of SDC by modifying an existing commercial microwave reveals an accuracy of 7-10 degrees respectively in accuracy of temperature actuation.

We envision software defined cooking as a future extension to molecular gastronomy [123]. Among all three common heating methods (convection, conduction, and radiation), radiation is the only one that can redirect energy towards the desired location. This redirect-able feature makes Microwave oven an ideal platform to experiment software-defined cooking, as it can effectively program energy transferring without physical hardware changes. Meanwhile, while there has been a great deal of past work on novel RF applications for communication [124], sensing [125, 126] and distributed power transfer [127], novel actuation mechanisms using RF signals are less explored so far. SDC is designed to innovate in this space.

3.1 Background on Microwave Ovens

3.1.1 Microwave Oven Primer

This section provides essential background on microwave heating. A typical residential microwave oven converts a large electrical input (≈1,000W) into microwave energy and heats food using microwave radiation. The microwave oven uses a magnetron, a high-powered vacuum tube, to generate a microwave signal at around 2.45 GHz from direct current electricity supplied to the vacuum tube [29]. This signal travels through a waveguide and creates an alternating electromagnetic field inside a metal cavity where the food is cooked.

Microwave Heating 101: In a microwave, water, fat, and other electric dipoles in the food will absorb energy from the microwaves in a process called dielectric heating [128]. Namely, when an electric field is applied, the bipolar molecules tend to behave like microscopic magnets and attempt to align themselves with the field. When the electrical field changes millions of times per second (e.g., 2450 million times per second for 2.45 GHz microwave signals), these molecular magnets are unable to keep up in the presence of forces acting to slow them. This resistance to the rapid movement of the bipolar molecules creates friction and results in heat dissipation in the material exposed to the microwave radiation.

While strong direct microwave radiation can burn human body tissue as well as electronic devices, the cooking chamber works as a Faraday cage to significantly attenuate waves escaping the microwave chamber. The US federal emission standard [129] limits the amount of microwave leakage from an oven throughout its lifetime to 5 milliwatts (mW) per square centimeter at approximately 2 inches from the oven surface (a safety factor of 10,000 or more below levels that may harm people [117]),

Once the microwave signals enter the metal cavity, they are effectively reflected by the metallic walls. Original and reflected waves resonate in the cavity and form standing waves [29], which produce anti-nodes (heating hot spots) and nodes (heating cold spots). The EM field are weak at nodes and therefore nothing cooks there. In the contrast, EM fields at anti-nodes alternate at maximum amplitude to produce maximum heating. This is also the reason why microwaves have a rotating turntable so that the turntable moves food in and out of the hot spots to cook more uniformly.

Microwave Limitations: Despite efforts to make heating uniform, microwaves are unsafe for many foods and cookware [31, 32]. For example, uneven heating will cause eggs in a microwave to explode. Sharp-edged metals in a microwave (e.g. forks) spark and create a fire. Many plastics may release chemicals into the food when heated. Due to these constraints, the microwave remains inhospitable to most modern electronics.

3.1.2 Heating in Microwave Ovens

The most widely adopted microwave actuation is the turntable and the stirrer blade [34] that attempt to spread radiation uniformly. A recent patent application [35] proposed to place a smooth-edge metal body inside the chamber, which operates as a passive antenna to reflect the microwave, to achieve a better heating uniformity. However, these blind actuation approaches cannot eliminate hot/cold spots, due to the inherent unpredictability of the EM field distribution. More recent advances in the microwave generators such as an RF solid state cooker (SSC) adjust the transmitter's real-time power, frequency and phase to move the hot/cold spots around, albeit at high cost ($\sim$ \$10,000 [130]) and complexity. SDC also draws inspiration from many microwave accessories that have been developed to cook certain foods in a microwave – e.g. Corning Ware Microwave Browners [131], Microwave egg boilers [132], or the susceptors in popcorn bags [133]. Unlike this past work, SDC provides a generalized framework for heat actuation as per a user-specified thermal trajectory as well as the sensing results, without being tied to specific types of food or adding expensive components.

3.1.3 Sensing in Microwave Ovens

There has been much past related work on improving the heat sensing within microwave ovens. For example, advanced FISO Microwave Work Stations (MWS) [134] used in food research have special microwave-safe fiber-optic sensors to collect real-time fine-grained direct measurements inside the cavity, but cost \$80k+. A recent startup, Cmicro [135], is building an energy-harvesting temperature sensor that can be run inside the microwave oven [136], which can only measure the air temperature in a container.
A more inexpensive approach for direct temperature measurement is to attach regular/thermal pinhole cameras[118, 119]. Note that these holes are designed to be too small relative to microwave wavelength ($\approx$12cm) and so continue to block radiation leakage. Unfortunately, measurements outside the cavity only provide limited accuracy and refresh rate [120]. Researchers [137, 138] have also used software radios to monitor the signal strength of the microwave leakage and recognize the type of food. However, many variables, such as food type, quantity, temperature and food location inside the oven impact microwave leakage unpredictably [29]. In contrast to these systems, SDC estimates both current and future temperature distributions by directly placing low-cost microwave-safe sensors within the cavity and modeling EM propagation.

3.1.4 Computational Fabrication and Heating

Designing computational fabrication techniques [139] for digital gastronomy is an emerging topic [140, 141]. For example, Zoran et al. [141] uses a silicone modular mould to control the shape permutations in a recipe, allowing computational control of taste structures. So far, there is little work on computational heating. Phosphenes [139] allows users to experientially compose resistive heaters that generate heat spatially and temporally. The most relevant approach is laser cooking [142], which uses a computer-controlled laser cutter to heat a sequence of small spots of the food surface. While innovative, the rolling pixel-by-pixel heating process is known to be highly time-consuming. SDC overcomes the slow production time of laser heating while allowing for a high degree of flexibility in the numerous heating patterns produced.

3.2 An Overview of SDC

This section presents the overview of SDC's architecture. SDC's goal is to cook food as per software-defined specifications (i.e., software-defined recipes). Yet, this raises an important question: what is software-defined cooking and how do we describe a software-defined recipe in the context of microwave cooking?

Specifying SDC's Inputs, Outputs and Goal: at its most basic, "cooking" means applying heat to food, which can be specified in three main variables:

(1) Temperature: The most important variable in cooking is the temperature of food, which will trigger different chemical reactions (e.g., protein denaturation, Maillard reaction, and caramelization). For instance, if we want to cook a steak at least rare to kill bacteria and avoid denature protein actin, we need to heat the meat to a temperature between 55°C (the highest survival temperature for most bacteria) and 65.5°C (the denature temperature for protein actin) [143]. We note that across all three common heating methods (convection, conduction, and radiation), food is cooked from the outside in, i.e. the outer portions will warm up faster, and the heat conduction will heat the inner parts over time. So SDC focuses on controlling the surface temperature of food.

(2) Time: Time is an important factor in both cooking food accurately and killing bacteria. For example, the standard food safety rule [144] provided by the FDA states various time and temperature combinations: heating at 55°C for 89 minutes can achieve a similar effect as heating at 62.2 °C for 5 minutes to reduce Salmonella.

(3) Space: Different parts of food (e.g. meat vs. fat, egg white vs yolk) may need to be cooked with different specifications to obtain optimal tasting food. SDC therefore aims to specify heating requirements

for each spatial "pixel" of the food's surface.

In summary, we envision that the future microwave recipes specify the desired thermal trajectory, i.e. temperature vs. time, for each "pixel" of the food's surface (see Fig. 3.3 for an example steak recipe). This is precisely SDC's input, with its performance dictated by how closely it follows this specification. We discuss the creation of such recipes in §3.8.5.

SDC's Heat Sensing: While SDC can deploy infrared cameras to sense the surface temperature of the food, doing so has two key disadvantages. First, cameras cannot sense heat in the presence of objects that blocks the food from view. Second, cameras would only measure the effect of heating after-the-fact and cannot prevent heating-damage that is already done. This motivates the need to sense "intent to heat" within the microwave rather than current temperature of the food content.

SDC addresses this challenge by deploying microwave-safe passive electronics that senses the underlying physical phenomenon producing heat, i.e. the electromagnetic (EM) field. Specifically, we deploy neon lights (see Fig. 3.4 and 3.6) that harvest EM energy in the microwave to glow in proportion to its intensity. We then measure the intensity of the light to infer field intensity and estimate the temperature distribution that would result. We do this via optical fibers that carry the visible light signal outside the chamber, allowing SDC to measure light signals from lights that are hidden from view. §3.3 discusses various challenges associated with such a design: (1) How many neon lights should be placed to optimally capture the EM field?; (2) How should the system be geometrically placed to best capture diverse cooking requirements?; (3) How do we map instantaneous EM field measurements and temperature to best anticipate future temperatures at high spatial resolution?

SDC's Heat Actuation: Having sensed the current and estimated heat distribution over space, SDC must now actuate the control system to focus energy towards some specific areas of the food and away from others. A traditionally available instrument for heat actuation within the microwave is the turntable. Microwaves use turntables to even out heat distribution, ensuring that hotspots do not hover statically over specific areas of the food. Yet, turntables are coarse actuators given that they spread out energy blindly rather than focusing it on specific parts of the food.

In contrast, SDC engineers specific hot and cold regions of the food by building specialized accessories that protect or focus specific regions of the food. We then electronically control the turntable to move food in and out of these sensed hot/cold spots as needed to cook food according to the desired heat pattern. §3.4 discusses two key innovations of our approach: (1) the design of metallic radio shields that are microwave safe; and (2) an efficient algorithm to move food to achieve the desired heating specification.

Figure 3.4: A turntable with 32 neon lights (left) and a plate cover with 32 neon lights (right).

3.3 SDC's Heat Sensing

SDC's heat sensing aims to capture both the current temperature of food as well as the intent to heat. At first blush, one may assume that heat sensing can be readily achieved using a thermal camera which captures the current temperature of the food. By measuring thermal camera readings over time, one can make predictions about how food will heat in the future. Yet, thermal cameras have important limitations that limit a design that relies exclusively on them for heat sensing. First, the food thermal properties evolve on a time scale of seconds, and the carryover in cooking will continue heating even if the food is removed from the source of heat. However, thermal cameras often have limited refresh rates (<9 Hz) and modest accuracy ($\pm$ 2°C) [120]. Therefore thermal cameras only measure the effect of heating after-the-fact and cannot prevent undesired heating, often until the damage is already done. Besides, thermal cameras are limited to measuring heating on the surface of the food in direct line-of-sight.

To mitigate this, SDC complements a thermal camera that senses current temperature with microwave-safe sensors that estimate future expected temperature. Specifically, we design sensors to sense the EM field – a set of neon lights, which glow in proportion to the EM field strength within the microwave. The rest of this section discusses the various challenges in achieving this design. First, we need to design microwave-safe hardware that is extremely sensitive to the EM field and can be readily probed from outside the microwave, even in the presence of blockages (e.g. food). Second, we need to build an expected temperature model that accounts time-varying energy transfer efficiency of the EM wave, the types of ingredients, temperature, and the placement of food.

3.3.1 Sensing Hardware Design

SDC places an array of neon lights, each has a 5mm diameter and 13 mm length, inside the microwave chamber to sense EM fields. A neon light (Fig. 3.6) is a miniature gas discharge lamp, which consists of a small glass capsule that contains a mixture of neon and other gases at low pressure and two electrodes (an anode and a cathode). During microwaving, the electrodes will couple with the electromagnetic field

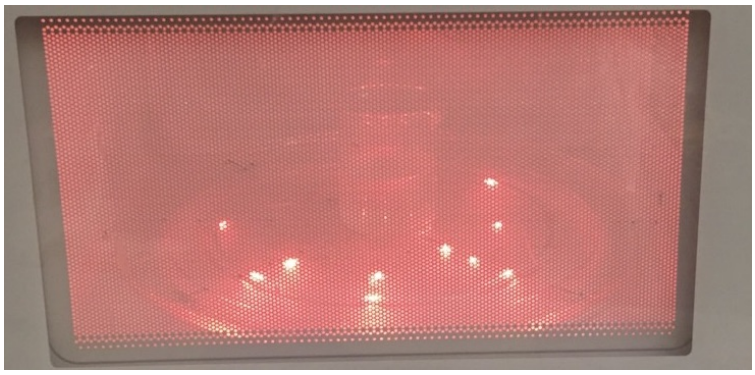

Figure 3.5: The turntable inside a running oven.

and act as antennas. The oscillating microwave applies a potential difference between two electrodes. Due to the potential difference, electrons are accelerated away from the cathode and give rise to collisions with the neon gas atoms or molecules, which will emit a characteristic glow. The brightness of the lamp is proportional to the EM field strength at the placed location, and SDC leverages that brightness to measure the EM field strength.

The glow of the light is sensed by a visible light camera outside the chamber to capture real-time EM field strength (Fig. 3.5). However, neon lights are often blocked from direct view of the camera due to obstructions such as food on the turn table. To mitigate this, SDC conducts the light from the neon lamps to the camera using optic fibers (Fig. 3.7).

Is SDC's hardware microwave safe? Neon lights are rated microwave-safe because the metal electrodes are encapsulated with a glass capsule and the gas glow discharge can avoid the energy accumulation. Each neon light consumes minimal microwave energy ($\approx$ 19.5 mW vs. the 1100 W available), producing negligible interference to the existing EM field patterns. Glass optical fibers are also microwave-safe.

The cameras used in SDC are not affected by the microwave because they are placed outside the microwave oven. The leakage through holes is negligible, since the mesh created for the optical fibers and the mesh in the front door are smaller than the 1/20 wavelength of 2.4 Ghz. Indeed, many commercial microwaves have holes of similar dimensions to support the turn table or stirrer fan.

Programming EM sensitivity: Much like EM fields in radio communication, it is important to tune neon lights to the correct range of sensitivity to obtain useful EM field strength. We define sensitivity of the neon light as the change in brightness for a given change in EM field strength. It is important to tune the sensitivity of the neon light to be in tune with the magnitude of the EM field. A highly sensitive neon light may be saturated by a strong EM field and burn the antennas, while a poorly sensitive neon lamp may not light up under a low EM field.

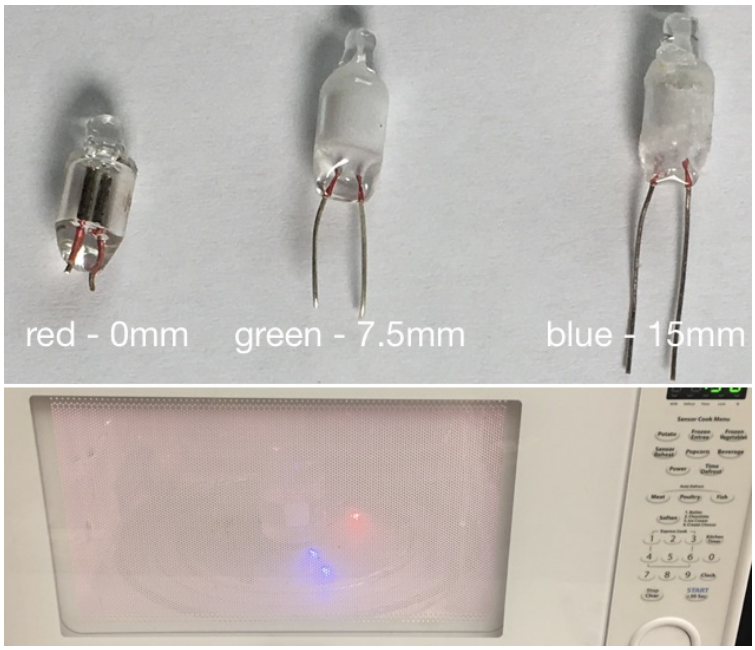

Figure 3.6: We place neon lights with different lengths of wire extensions (red: 0mm, green: 7.5mm, and blue: 15mm) under the turntable. We then measure the percentage of glowing time to quantify the sensitivity.

SDC programs the sensitivity of a neon light by changing the length of electrodes wire extensions (Fig. 3.6). Each neon light is characterized by two threshold voltages: the ionizing voltage and the maintaining voltage, which depend on the type of electrodes, their coatings, the composition of the gas, and its pressure, etc.

The neon light is dark when the gas is not ionized. When the voltage between two electrodes exceeds the ionizing voltage, the lamp switches on and the brightness is in proportion to the current when the light switches on. The ionizing voltage is the minimum voltage required to keep the gas ionized. When the lamp is on, and the voltage drops below this threshold, the gas loses its ionization, and the light turns off. The sensitivity of the neon light to the EM field is a function of how the EM field affects the voltage applied at the electrodes. We note that this process is completely dictated by the wires connected to the electrodes, which behave like RF antennas. SDC thus strives to tune the length of the wire extensions to the electrodes to achieve optimal sensitivity.

While antennas are known to resonate best at the length of one-half-wavelength, this quantity can vary as a function of antenna design and the RF frontend [145, 146]. SDC therefore chooses to measure the optimal antenna length for the neon lamps experimentally. Specifically, depending on the EM field strength, a neon light may experience one of the following three states when running the microwave: constantly off, flashing at various intervals, and constantly on. The flashing state (flashing frequency)

and consistently on (brightness) offer more fine-grained resolution of the EM field than the consistently off state. An ideal neon light sensitivity would ensure that a good percentage of the neon lights in the oven should be in the flashing and consistent on states. To tune the optimal antenna length (i.e., the wire extension length), we empirically tested various types of loads in the microwave and found that a wire extension of 8mm achieves the desired sensitivity.

Placement of Neon Lights: We place 48 T2 fixed orange neon lights (65 VAC, 0.3 mA) with a 3 cm spacing at the surface of the Microwave cavity. We set this spacing to be significantly below the wavelength of the 2.4 GHz EM signal of the microwave oven. The neon lights are placed below the turntable, as well as around the inner walls of the oven. To avoid visual occlusion, we used optical fibers to connect the lights to a ventilator scoop, carrying the light signals outside the chamber (Fig. 3.7). We experimented with optical fibers at different sizes and found that 1 mm diameter fibers provide good flexibility and operability.

While our sensing platform surrounds the food within the microwave chamber, one still needs to interpolate the EM field within the microwave in 3-D space at locations where neon lights are not present. Fortunately, the EM field within the microwave is that of a standing wave which has predictable voltages over space with a modest amount of sampling [147]. SDC applies the cubic spline interpolation [148] to approximate the expected brightness of neon lights should they be placed at remaining 3-D points in the microwave.

To further refine SDC's resolution particularly within the vicinity of the food, we develop two sensing containers: one microwave plate cover (Fig. 3.4) and one measuring cup, to measure the EM field in the air. We embedded 32 neon lights around a glass container and attached 32 neon lights to a microwave plate cover. We mainly use orange T2 neon lights intermixed with a few blue neon lights at known locations as reference points, which help SDC track the location and identity of each container as they rotate on the turn table. Since the containers are placed on the turntable, we do not connect these neon lights to the optical fibers.

3.3.2 Modeling Heat Over Time and Space

Next, we describe how measuring the brightness and flashing frequency of strategically placed neon lamps will help us model the current and future temperature of food over 3-D space and time.

Creating a Spatial Heatmap: The visual camera (ELP 4mm Lens Prototype Camera) captures the brightness of neon lights in a real-time video stream with a 640$\times$ 480 resolution at 120 fps. The visual camera is placed at the front of the microwave oven. The top part of the frame tracks the signals carried by the

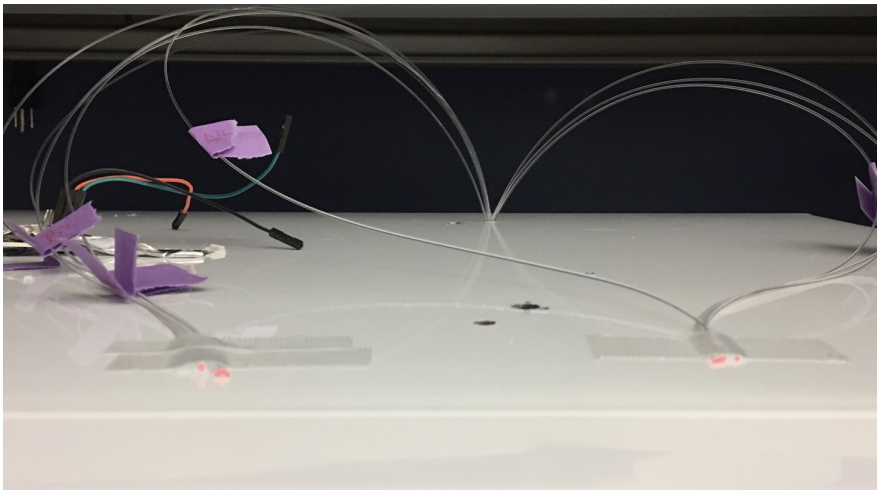

Figure 3.7: The optical fiber carries the signal outside of the chamber.

Optical Fiber which contains the brightness information of all the fixed neon lights. The camera also observes the brightness of neon lights on the sensing containers through the front door mesh (Fig. 3.6).

SDC measures the brightness of the lamps every 0.1 seconds (i.e., 12 frames of image capture). The ends of optical fibers are static. SDC uses optical flow [149] to track the movable neon lights. For each frame, SDC converts the image to grayscale, finds the pixels around the neon light or the end of an optical fiber, and sums up the pixel values as the brightness score. We manually calibrate to normalize the brightness score between the neon light in direct view and those connected to optical fibers. As mentioned previously, since the location of the neon lights are known as prior, we interpolate the brightness at remaining locations using cubic-spline interpolation. We then map neon light brightness and flashing frequencies to EM field strength empirically by comparing results from co-located neon lamps. This, coupled with spatial interpolation allows us to generate a 3D EM field intensity view within the microwave chamber. SDC can therefore estimate the EM field strength given a specific location at fine-grained spatial resolution.

Modeling Temperature over Space-Time: As mentioned previously, SDC can measure current temperature of the food surface by placing a thermal camera (AMG8833 IR Thermal Camera) on the roof of the microwave oven to sense the food surface in a top-down view. The thermal camera provides an 8x8 temperature readings with an accuracy of $\pm$ 2.5°C at 10 Hz. SDC interpolates the raw sensor input into a 240x240 square array using cubic-spline interpolation. However, using a thermal camera exclusively to model temperature has two limitations: (1) first, the camera only measures temperatures on the food's surface in its direct field-of-view; (2) second, the camera only measures current temperature and not future expected temperature.

SDC estimates future heat by integrating measurements of the EM field obtained from the neon lamps. Specifically, the heating of any given point in space of the food is proportional to the EM field intensity of that location. This means that integrating the observed field intensity, while accounting for the rotation of food over time can provide a robust estimate of its future temperature. Yet, two challenges remain in making this mapping accurate: (1) First, the temperature of the food for the very same EM field may change owing to the material composition of the food itself; (2) Second, integrating EM field intensities over time may cause errors to build up progressively as well.

SDC mitigates the limitations of both the thermal camera system and the EM field estimation by combining them and obtaining the best of both worlds. Specifically, at each point in time we reconcile the integration of EM field estimates with those of the thermal camera over space. We then use this to refine our model for mapping EM field to temperature. We repeat this process over time to continuously avoid any drift of our EM field to temperature mapping, as well as accounting for material properties.

Mathematically, SDC constructs an Extended Kalman Filter (EKF) model [150] to capture both the current temperature and its gradient over time of the microwave. Specifically, it is well known that the temperature gradient P' of the microwave for a current temperature P is proportional to the electromagnetic field strength E, i.e. $P' = kE$, where k is a constant that depends on the material properties of the food. [151]. Then we can write the following recursive equation for future temperature based on current temperature at time t:

$$P(t+1) = P(t) + P'; P' = kE; z(t+1) = P(t+1) + n$$

Where n is noise and $z(t+1)$ is the observed temperature. The above follows the formulation of EKF and can be used to estimate k as well and refine the temperatures P and gradients P' over time.

3.4 SDC's Heat Actuation

Today's microwave ovens actuate the heating process in a crude manner. The oven turntable rotates the food blindly without any precise control. The magnetron, the heating engine of the microwave oven, achieves power control by periodically turning itself on and off. SDC augments these existing blind actuation hardware into a closed-loop control system by incorporating the results of heat sensing.

3.4.1 Actuation Hardware

Smart Turntable: We modify the default turntable inside a commercial microwave oven (Sharp SMC1441CW). More specifically, we replaced the motor with a low-cost stepper motor (Yosoo 57oz-in 1Nm Nema 17 Step-

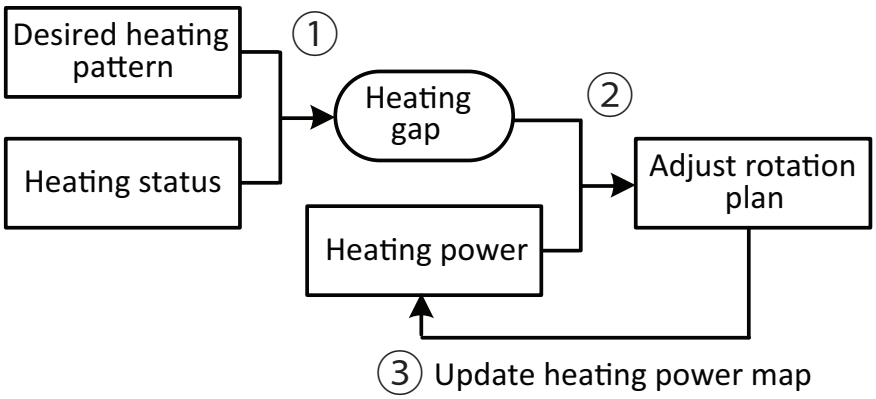

Figure 3.8: Workflow of SDC's heat actuation: ① SDC first computes heating gap by comparing the desired heating pattern with the current status from the thermal camera. ② SDC then adjusts the actuation plan accordingly and ③ updates the distribution continuously.

per Motor), and 3D printed a plastic coupler between the motor head and the glass platform to enable precise control. We also connected the magnetron to an Arduino and programmatically sent pulse-width modulation signals to control the ON /OFF of the magnetron.

SDC sets the rotating speed of the turntable to 12 RPM. The heat actuation will only start after the cycle once SDC collects an initial spatial distribution of the EM fields. Rotating the food around can manipulate the heat pattern mildly (e.g., uniform heating), but it is insufficient to create an arbitrary heat pattern.

Programmable accessories: To achieve a more skewed heating pattern, SDC develops programmable accessories (Fig. 3.9) that leverage the reflective property of microwave heating, redirecting energy towards desired locations and shield undesired locations, to achieve an arbitrary heating capability.

SDC installs a horizontal glass plane above the turntable and carefully packs metallic-spheres (3.175 mm diameter) (Fig. 3.9 right) at specific regions of the plane to minimize RF energy-transfer at these regions. Metallic spheres by definition do not have edges and are therefore microwave-safe [122]. These metallic spheres effectively form a microwave mirror to reflect microwave energy at the specific region.

The most common dielectric dipole in the food is water, so microwave heating rarely achieves temperature beyond the boiling point of water. However, some important food chemical reactions occur well above water's boiling point, such as Maillard reactions and caramelization. We introduces microwave susceptors to address this limitation. Materials like silicon carbide (Fig. 3.9 left) can effectively absorb microwave energy inside the oven and reach 200+°C within 1 minute microwaving. Attaching silicon carbide to the food surface can then trigger desired high-heat reactions.

When to use accessories? Given the desired heating pattern, SDC first determines if the accessories are needed to achieve the goal. We empirically find that using only the smart turntable, SDC can achieve a

maximum peak-to-peak temperature difference of 21°C.

To obtain a more skewed thermal distribution, SDC can either involve the accessories at the beginning (e.g., preheating a microwave susceptor to a high temperature and then attaching it to a desired region) or involve them at a later stage (e.g., placing microwave shields to protect certain areas if it's close to being overcooked). SDC's actuation algorithm will dynamically adjust the actuation plan in either case. Involving accessories does not change the nature of our sensing or actuation optimization as the neon lights still reflect the real time EM field strength around the food.

3.4.2 Recipe and Actuation Representation

Having developed programmable actuation hardware, this section formally states the actuation optimization problem that attempts to heat the food in accordance with the input heating recipe.

Heating Recipe: A SDC heating recipe will specify the desired temperature trajectory and duration for each part at different temperatures. Mathematically, we can formulate the recipes as follows. Let us imagine that the food surface is divided into a set of discrete pixels. Given n pixels $B = \{B_1, B_2, ...B_n\}$ on the surface of the food, the 3D coordination of the pixels are $\{x_i, y_i, z_i\}$ where $i \in \{1, 2, ...n\}$. The receipt is a mapping function f that maps the pixels and the timestamps to desired temperatures throughout the D minutes cooking journey:

$$f(B_i, j) = p_{ij}, \quad i \in \{1, 2, ...n\} \quad 0 < j < D \tag{3.1}$$

where j denotes the timestamp since start of the cooking process, and p_{ij} refers to the desired temperature for i-th pixel at the timestamp j.

SDC's Optimization Problem: Our goal of the smart turntable is to find a rotation plan S^* that can move food in and out of these hot and cold spots as needed to cook food according to the desired heat trajectory P, which contains collection of desired temperatures p_{ij} across the space and time. SDC defines a rotation plan S as a sequence of angle-duration and magnetron on-off-duration tuples:

$$S = \begin{bmatrix} \{\theta_1 : d_{\theta 1}\}, \{\theta_2 : d_{\theta 2}\}, \{\theta_3 : d_{\theta 3}\}, ... \\ \{o_1 : d_{o1}\}, \{o_2 : d_{o2}\}, \{o_3 : d_{o3}\}, ... \end{bmatrix} \tag{3.2}$$

$$D = \sum \{d_{\theta 1}, d_{\theta 2}, d_{\theta 3}, ...\} = \sum \{d_{o1}, d_{o2}, d_{o3}, ...\} \tag{3.3}$$

where $\{\theta_k : d_{\theta k}\}$ indicates that the turntable will stay at the absolute offset angle θ_k for a duration of $d_{\theta k}$, $\{o_k : d_{ok}\}$ describes the duration d_{ok} for keeping the magnetron on or off (o_k). Based on these definitions,

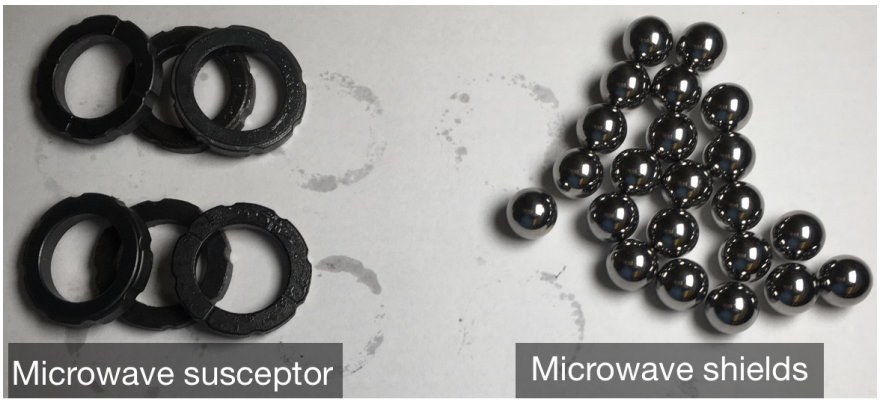

Figure 3.9: Programmable microwave accessories.

we now formulate SDC's core optimization problem as follows:

$$S^* = \arg\min_S \sum ||\bar{P}(S) - P||^2 \tag{3.4}$$

where $\bar{P}(S)$ denotes the temperature trajectory for the n pixels using a rotation plan S over time.

3.4.3 Actuation Algorithm

Solving this optimization problem is challenging for two reasons. First, microwaves heat the food through a standing wave, so they cannot heat individual pixels independently. Heating one pixel will inevitably heat other pixels as well. To achieve the desired heat pattern, we need to select a set of heat patterns whose union is equivalent to the target heat pattern. Second, the heat pattern, the result of the EM-ingredient coupling, is non-static and unpredictable. The EM field distribution changes gradually when the turntable rotates the food and when the food heats up. SDC cannot predict the output heat pattern until the food is heated.

The Stochastic Knapsack: At a high level, this problem is a variant of the stochastic knapsack problem [152, 153], a classic resource allocation problem of selecting a subset of items to place into a knapsack of given capacity. Placing each item in the knapsack consumes a random amount of the capacity and provides a stochastic reward [152], which are only observable after the item is placed. Finding exact solutions to the stochastic knapsack problem is PSPACE-hard [154]. Indeed, a cyberphysical system like SDC cannot afford long computational processing times, since once food is heated too much, it cannot reverse the heat applied.

Due to the intrinsic uncertainty of stochastic knapsack problems, adaptive and closed-loop strategies often perform better than open-loop ones in which the items chosen are invariant of the remaining time budget [154]. So different from the traditional 3D printer, which computes the motion trace in advance,

we design SDC to calculate the rotation plan on-the-fly, based on the real-time sensing feedback.

Approximation Algorithm: We propose a greedy approximation algorithm to determine the immediate rotation plan on-the-fly based on the sensing result in §3.3. Our greedy strategy is as follows: *"at each step of the journey, heat at the rotation angle whose temperature gradient is most similar to the current heating gap"*.

Fig. 3.8 illustrates the workflow of the greedy algorithm. In SDC, the thermal camera continuously senses the current food temperature at the n pixels: $C = \{c_1, c_2, ..., c_n\}$. SDC then compares the desired heating pattern $f(B_i, j)$ with the observed thermal distribution, and computes the real-time heating gap (i.e., available capacity) $G = \{g_1, g_2, ..., g_n\}$. SDC also maintains a dictionary $\{\theta : P'_\theta\}$ to document the temperature gradient P'_θ at each offset angle of the turntable. SDC then queries the vector $P'_\theta = \{p'_1, p'_2, ..., p'_n\}$ using the pixel coordinates.

When running, SDC continuously computes the similarity Sim_θ between the temperature gradient at θ and the current heating gap using the cosine of an angle between these two vectors.

$$Sim_\theta = \frac{P'_\theta \cdot G}{|P'_\theta| \cdot |G|} \tag{3.5}$$

Once computed, SDC rotates the turntable to θ^*, which has the most well-aligned temperature gradient. We note that the turntable may not reach all the targeted angles θ^*, as the most well-aligned temperature may change during the rotation and a new rotation command will override the earlier one.

$$\theta^* = \arg\max_\theta Sim_\theta \tag{3.6}$$

Simultaneously, as explained in §3.3.2, SDC updates the dictionary of temperature gradients based on the EM field and real-time observation of temperature from the IR camera at each rotation angle θ.

3.5 Implementation and Evaluation

We implement SDC by modifying a commercial Microwave oven (Sharp SM1441CW). As explained in Sec. 3.3, our setup contains an array of neon lamps on the bottom of the turntable whose signals are conducted to a ELP 4mm Lens Prototype camera outside the microwave via fiber-optic cables attached to the base of the microwave (see Fig. 2.1). We perform actuation with an Arduino board connected to a Yosoo 57oz-in 1Nm Nema 17 Stepper Motor. We implement SDC's optimization in real-time with SDC's actuation algorithms implemented in Python. We conduct experiments in an indoor space on a kitchen table with a variety of food types including meat, rice, milk and fish and evaluate SDC's heat sensing and actuation.

Performance Metric: To characterize SDC's performance, we focus on the two metrics: (1) ΔP: peak-to-peak temperature differences (i.e., the difference between the maximum and the minimum temperature)

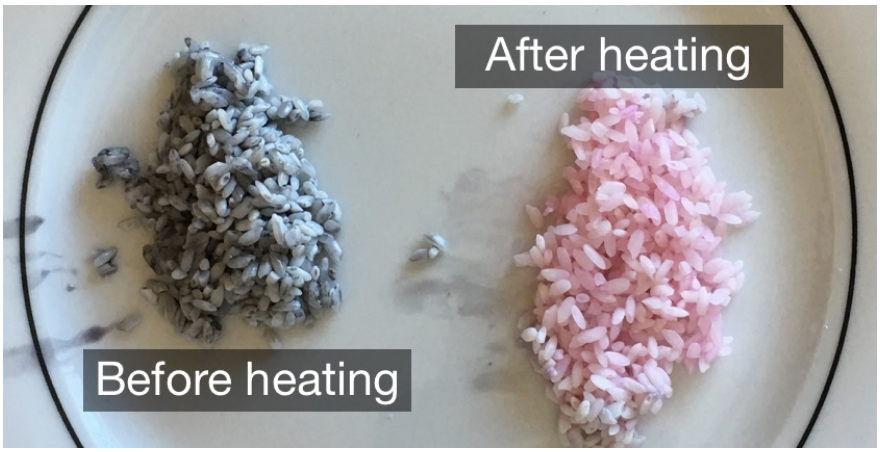

Figure 3.10: We color rice grains with thermal-chromatic pigment, which turn pink in a predictable manner as their temperature increases.

and (2) thermal delta, $P_{Dev} = \sqrt{\sum ||\bar{P}(S) - P||^2}$ (Eq. 3.4), i.e., the collective temperature deviation between the actual output and desired pattern. To quantify the thermal delta, we measure the temperature at 9 discrete points on the surface in a 3x3 grid with 2cm spacing.

Ground truth: To obtain ground-truth temperature data, we use a non-contact infrared thermometer (Etekcity LaserGrip 630), which provides $\pm 2°C$ resolution from -50°C to 580°C, as well as the thermal camera we used in the SDC.

3.6 Results

3.6.1 Uniform Heating

Non-uniform heating is a major drawback associated with today's microwaves [49], which not only affects the quality of the food but also compromises food safety when the microorganisms may not be destroyed in the cold spots. This experiment evaluates SDC when provided with a Uniform heating plan, a commonly input thermal trajectory provided to SDC, in which we aim to heat all the pixels to the same temperature at a uniform pace.

Method: We conduct our evaluation by heating raw rice grains using SDC. While the specific type of material heated does not impact our evaluation result, we choose raw rice grains for three reasons: (1) First, rice grains are known to absorb microwave energy; (2) Second, the limited contact surfaces among the grains conduct heat in a limited manner, allowing us to visualize pixel-by-pixel heat more clearly; (3) Third, rice grains are reusable across multiple evaluation sessions once cooled, so we can keep the food that we heat consistent across experiments.

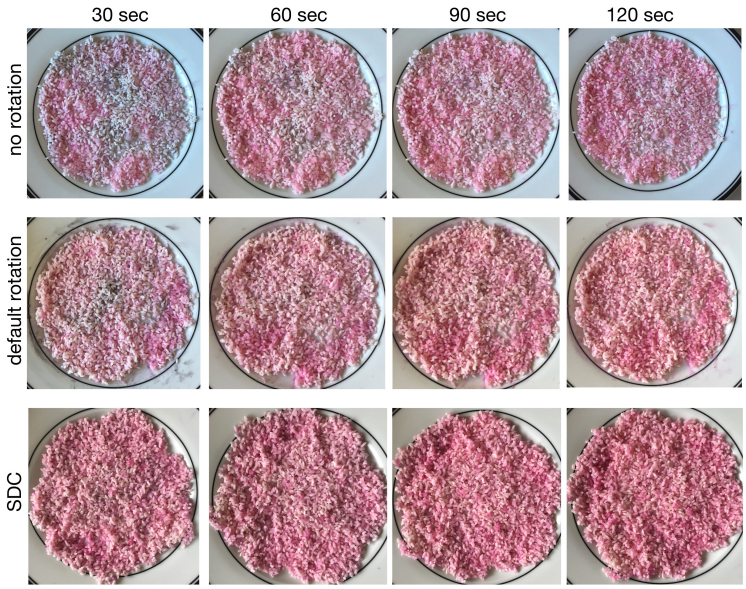

Figure 3.11: Visualization of heating of rice as a function of time for no rotation, default rotation and SDC. SDC results in the most uniform heating.

To visualize the heat pattern, we color the grains with thermal-chromatic pigment (Fig. 3.10), which will turn into pink as the temperature increases. We use the thermal-chromatic pigment approach because it can provide a rich and analog temperature visualization, while thermal cameras have a limited resolution and the final output image are based on the interpolation.

We begin our experiment at a room temperature of 20°C. We set a target temperature of 60°C over 2-minutes of microwaving, based on multiple empirical runs evaluating the typical temperature achieved over this period. We create a uniform heating recipe that requires that the food be heated uniformly from 20°C to 60°C over 2 minutes uniformly over space as per the recipe provided by the thick blue line in Fig. 3.12. We note that the thermal trajectory is identical across all pixels of the food's surface. However the temperature increase is not designed to be linear over time, instead mimicking the smoothed average temperature trajectory for raw rice within a microwave under normal microwave operation.

To characterize the benefits of SDC, we use two blind microwave heating as the baselines: the same microwave oven (1) with and (2) without turntable rotation. To collect the immediate temperature during heating, we take out the food every 30 seconds to measure the ground truth temperatures.

Results: Fig. 3.11 shows a visualization of the thermal-chromatic pigment, which changes color at 31°C and progresses to darker shades of pink with increased temperature. The rice colored (dull white) regions denote spots of food that remain below 31°C. We observe that SDC achieves a uniform pink hue that darkens over time, while the baselines (no rotation or default rotation) continue to have cold spots through time. Note that SDC visually appears to have the deepest shade of pink vs. the baselines at t=120 seconds

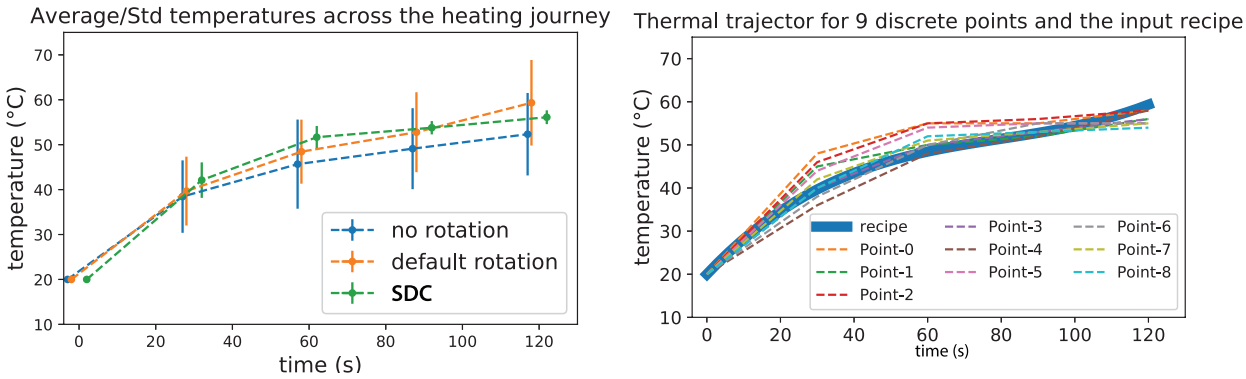

Figure 3.12: (left) Temperature variance of SDC is low vs. baselines; (right) points on food closely follow heating recipe.

as it achieves more spatially uniform temperature relative to the baselines. In actuality, there are also some hot spots of the baseline schemes that achieve even higher temperatures (over 70 °C), while SDC achieves uniform temperature closer to 60 °C as desired.

Validating Heat Sensing and Actuation: Fig. 3.12 (right) shows the trajectory of the temperature over time for nine discrete uniformly spaced points of the food using SDC. Note that all points closely follow the recipe over time, , which demonstrates SDC's high accuracy in modeling the temperature gradient through EKF. Fig. 3.12 (left) compares the average and standard deviation of the trajectory across the same discrete points measured over multiple experiments vs. time. Of particular interest here is the standard deviation of the temperature of the food where one can clearly observe that SDC achieves a lower spatial variance in temperature when compared to the baseline schemes. This validates our findings that while microwaves heat food blindly and non-uniformly, SDC can achieve significant uniformity in heating.

The table to the right measures two quantities for SDC and the baselines: (a) P_{Dev}: standard deviation of temperature across space for any given time; (b) ΔP: The difference between maximum and minimum temperature over space on the food.

Scheme	P_{Dev}	ΔP
No Rotation	9.2 °C	29.5 °C
Default	9.5 °C	24.8 °C
SDC	1.5 °C	6.5 °C

Our results validate the correctness of both heat sensing and actuation, both of which must operate correctly to achieve the desired heating objective.

3.6.2 Arbitrary heating

In real-world cooking, different ingredients often require to be cooked at different temperatures. SDC can support these activity computationally by specifying thermal trajectories for different surface pixels. In

Figure 3.13: Left: An input recipe for stress test. Middle: SDC without susceptors. Right: SDC with susceptors. Susceptors can help build more skewed thermal distributions.

Scheme	ΔP	σ_P	$\frac{\Delta P}{d}$
without accessories	21°C	99°C	3°C/cm
with accessories	183°C	42°C	61°C/cm

Figure 3.14: Mean (ΔP) and Standard Deviation (σ_P) of thermal delta for arbitrary heating. The final column ($\frac{\Delta P}{d}$) denotes the temperature gradient per unit distance that can be achieved.

this section, we aim to stress test SDC by exploring the maximum heating resolution, i.e., the maximum temperature difference that can be created in a fine-grained spatial resolution. The apparatus, performance metric, and the ground truth acquisition method are the same as the uniform heating experiments. The only difference is the input thermal recipe to SDC.

Method: We create an imaginary recipe that heats a unique thermal pattern (depicted in Fig. 3.13 – left). The recipe sets the target temperature for the inner ring area at 500°C and the rest area at 50°C. We deliberately set an unachievable goal of 500°C for SDC to stress-test the system and evaluate how well SDC can approximate to the targets. We conduct two independent experiments with and without the help of microwave accessories.

Results: Fig. 3.13(right) shows a visualization of the thermal-chromatic pigment, which aligns well with the desired pattern (Fig. 3.13 – left), without and with microwave accessories. As expected, we observe that the presence of accessories helps improve the contrast between the high temperature and low temperature rings. This is precisely why accessories are needed to improve SDC's performance during tasks such as searing, where extreme temperature gradients are needed on the food.

Figure 3.14 on the right summarizes the mean ΔP and standard deviation σ_P of the maximum temperature difference between the inner and outer ring achieved in SDC with and without microwave accessories. We also note that with microwave accessories, SDC can cause extremely high temperature gradients (up to 61°C per centimeter) at very fine spatial resolution.

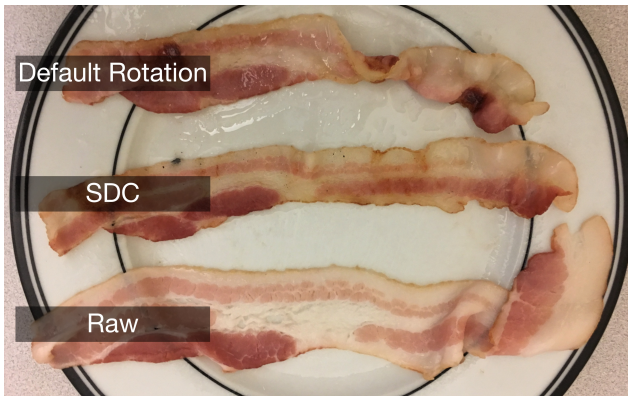

Figure 3.15: The raw bacon, and slices of bacon cooked with SDC and the original turntable. These three slices of bacon are from the same package, so the original fat patterns are nearly identical. Heated meat and fat will shrink. SDC applies heat to meat and fat discriminatively, so the fat shrinks more than the meat.

3.6.3 App – Microwaving Milk

Microwave ovens are frequently used for reheating leftover food, and bacterial contamination may not be repressed if the safe temperature is not reached across the food's surface. This results in the risk of foodborne illness, as with all inadequate reheating methods. While microwaves can destroy bacteria as well as conventional ovens, they do not cook as evenly, leading to an increased risk that parts of the food will not reach recommended temperatures [155]. More fundamentally, microwaving liquids may cause uneven heating so that some parts of the liquid may scald the mouth. This is precisely why microwaves are not recommended for heating liquids such as milk, particularly for infants. In this section, we evaluate the effectiveness of SDC in uniformly heating a cup of milk.

Method: We specify a uniform heating recipe that requires a spatially uniform final temperature of 71°C [156]. We place 200 ml milk in a 500 ml measure cup across multiple experiments. We measure the ground truth temperature at the end of the experiments using a contact-based thermometer (ThermoPro TP18). We measure the temperature at 9 different points of the cup. We use a contact based sensor for this experiment since the non-contact infrared thermometer only measures surface temperature.

Results: Figure 3.16 presents the thermal delta (deviation) between the planned and observed temperature per-pixel across experiments for the milk. We observe a temperature between 67°to 74°at all the 9 points. The whole milk is heated uniformly without parts that may scald the mouth or be too cold and therefore preserve bacteria.

Scheme (HF=SDC)	ΔP	σ_P
HF with milk	7°C	N/A
HF with bacon (meat)	10°C	6°C
HF with bacon (fat)	8°C	5°C

Figure 3.16: Mean and Standard Deviation of thermal delta for cooking milk and meat. We don't have σ_P for milk since the thermal camera can only sense the surface temperature.

3.6.4 App – Cooking Meat (Bacon)

In this section, we evaluate SDC's performance in cooking with advanced thermal recipes. We consider bacon with different heating requirements for the meat and fat.

Method: We use the heating recipe required per pixel for the meat vs. the fat as shown in Fig. 3.3. Based on the package instruction, we set the heating time to 1 minute. We place strips of bacon on the microwave plate across multiple experiments. We measure the continuous ground truth temperature using the thermal camera, and measure the final temperatures at discrete points using the non-contact infrared thermometer (Etekcity LaserGrip).

Results: Figure 3.16 presents the thermal delta (deviation) between the planned and observed temperature per-pixel for the bacon, across experiments. Fig. 3.15 depicts the initial and final product over the cooking process. Note that while default rotation heats the bacon unevenly (resulting in the uneven shape), SDC heats the microwave more uniformly while differentiating between the heat applied to meat and fat. Indeed, we observe that the cooking process does not overcook/burn the meat, while at the same time avoiding colder spots that may pose a health hazard.

3.7 Limitations

We emphasize the following limitations of our current implementation. (1) Thermal Camera: The accuracy of the commercial thermal camera we are using maintains a high accuracy only between 0°C and 80°C. Beyond these temperatures, the accuracy of thermal camera will decrease. (2) Neon Lights: Our current prototype handles up to 64 neon lights (32 on the microwave plate and 32 on the glass container) and scaling to larger numbers makes our system too bulky. We believe this can be addressed in commercial implementations. (3) Heating Model: Our heating model makes no assumption and does not use prior information on the material type of the food. Using this information (via user input or through image recognition) can greatly improve the performance of SDC. We leave this task for future work. (4) Some hardware (e.g., silicon carbide) used in the SDC prototype may not be FDA approved.

3.8 Discussion

3.8.1 Actuations beyond turntable

SDC mainly modifies the EM-field strength to desired levels by accurately moving food on a programmable 2D turntable. However, the proposed techniques, e.g., the EM field sensing (§3.3.2), the greedy algorithm (§3.4.3), can be applied to 3D actuation directly. Throughout our development, we had experimented several 3D actuation hardware designs, e.g., using strings to suspend the food container in the air [157], adding an elevation platform to the turntable. However, either approach requires a large chamber space, making the oven clumsy and hard to clean. So we decided to retain the default 2-D turntable design, and leave alternate designs for future work.

Beyond food movement, there also exists other degrees of freedom in actuating the microwave energy. First, we can dynamically adjust the chamber size or place a microwave-safe reflector inside to change the standing wave formulation. Second, we can replace the magnetron with solid-state microwave ovens and apply RF beamforming on the food. We leave these options for future work.

3.8.2 Alternative hardware design

Retrofitting existing commercial microwave ovens for new contexts allows us to re-imagine the design of future microwave ovens. Commercial microwave ovens use either 900 MHz or 2.4 GHz, which fall into the unlicensed spectrum. Future microwave ovens may consider supporting a higher frequency mode (e.g., 10 GHz). 10 GHz radio has a smaller wavelength (i.e., 3.0 cm), resulting in a lower food penetration (0.4 cm vs. 2cm at 2.4 GHz [158]), which might allow the microwave to produce cooking effects like searing or frying.

A pure electrical mechanism may also replace the neon bulbs. The bulbs currently harvest energy through small dipole antennas, which can potentially be fed into a rectifier to make direct electrical measurements of EM field strength. The challenge of such a solution is to tune the circuit to ensure it will not absorb too much RF power and burn itself.

3.8.3 Cross-technology interference

Microwave leakages create cross-technology interference on Wi-Fi communication [137], which can potentially improve the cooking process as well. Different types, positions, and quantities of food will absorb the microwave energy differently, impacting the microwave leakage [29]. Throughout our development, we run several pilot studies to understand the relationship between the content inside microwave and

microwave leakage. For example, we noticed that different food temperature and internal moisture would change the leakage frequency. This might allow future systems to determine if the food is starting to dry out.

3.8.4 Energy efficiency & cooking process

Microwave heating is a wireless power transfer application, which converts electrical energy into microwave radiation and then transferring the heat into the food. The energy efficiency of the microwave generation is relatively consistent (around 65% for a modern magnetron [159]), while the food heat absorption depends on the food size, material, and the standing wave formulation [160]. Our technology may cook faster than the blind heating approach because we can intentionally heat the cold spots while the traditional method relies on the slow heat conduction to cook cold spots.

There also exists a trade-off between heating accuracy and duration. One extreme example is that SDC can turn off the magnetron during rotation and only heat the food at the position when it has the most well-aligned temperature gradient. The heating process would then be more accurate but slower. Besides, the shape of the food will impact the performance of discriminative heating as well: too large or small food surface will make SDC hard to rotate the food in/outside of the hot/cold spots.

3.8.5 Recipe Creation

SDC assumes that the desired heat distribution is available to SDC either through manual user input or from an auxiliary food-sensing system. We envision that the future cooking recipes will detail the desired temperature and time duration for each "pixel" of the food. Since the food will be cooked from outside in, the recipe can be represented as the combination of surface temperatures and time durations. Chefs can develop such recipes through empirical experiments, or the food developers can computationally model the ingredient (e.g., using computer vision [161]) and generate a heating recipe through cooking principles (e.g., the ideal temperature to cook a medium-rare steak is between 55°C to 60°C [162]).

3.9 Summary

This chapter presents SDC, a novel next-generation microwave oven that both senses and actuates heating at fine-grained spatial resolution. SDC uses programmable RF-powered neon lights whose signals are conducted by microwave-safe optic fibers to sense the EM-field distribution. It adapts the spatial heat distribution within the microwave chamber by moving food between regions whose temperature is controlled using microwave-safe wave-guides. A prototype of SDC reveals promising accuracy in heat

sensing and actuation, opening up the microwave to new cooking applications such as searing steak and defrosting fish.

We believe this work opens up a range of directions for future work: (1) While our system considers only rotation of the turntable, allowing for more degrees of freedom could lead to finer actuation accuracy. (2) While we assume recipes are statically provided by the user, learning recipes over time and recognizing food through EM/image recognition is an exciting topic for future work. (3) Exploring inexpensive ways of performing high power RF beamforming as a more refined actuation mechanism remains an open challenge.

Chapter 4

Wireless Actuation for Electronics-free Soft Robots

Beyond the wireless actuation inside a microwave oven to heat food, this chapter reveals that using a similar wireless heating concept from Chapter 3 is a promising option for addressing a new research domain: moving or actuating thermally driven soft robots.

Soft devices and robots are emerging novel platforms used in various ubiquitous computing and human-computer interaction applications, including enabling soft surfaces [163], shape deformation sensing [164], and stretch sensing [4]. Soft robots offer unique features compared to traditional robots: miniaturization [165], shape morphing [166], and adaptability to complex and constrained terrains [167, 168].

In this chapter, we introduce a framework for engineering soft robots that move and change shape in response to *wireless actuation*. While various classes of soft robots exist, we target the development of machines that are composed entirely of soft materials and exclude any rigid components such as integrated circuits or batteries [169, 170, 171]. To achieve this, we use liquid-crystal elastomers (LCE) [172, 173], a shape memory polymer material that has become increasingly popular in soft robotics[174]. Innovations in material science have shown that LCEs – soft polymeric systems with well-defined and orientable molecular alignment of liquid crystal mesogens – can be fabricated into thin films that exhibit large, fast, complex and reversible shape changes upon *thermal* stimulation [175, 176, 177]. When heated to 60° Celsius [178], LCEs can undergo maximal contraction with strains and stresses comparable to that of natural skeletal muscle [50].

Unfortunately, existing approaches to actuate LCE-based soft robots remain rudimentary. Past work leverages resistive Joule heating [50, 51, 52, 53] to produce heat through electrical current delivered by an external DC power supply [54]. More recently, stretchable thermoelectric Peltier devices have been

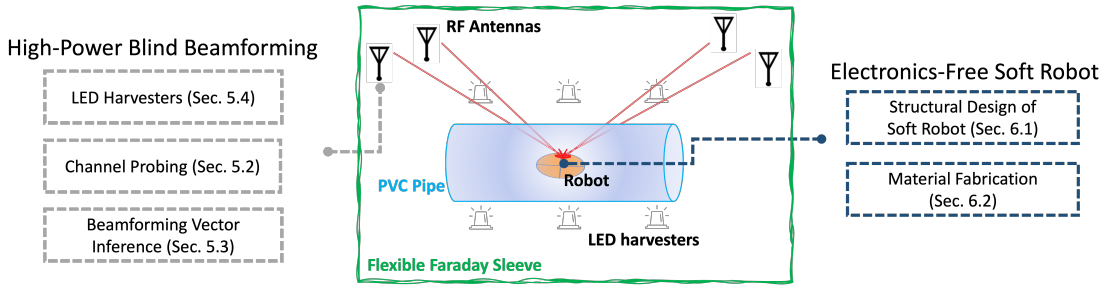

Figure 4.1: The soft robot can potentially be in the non-line-of-sight (NLoS) from the RF antennas due to various blockages, e.g., a plastic pipe. WASER estimates the wireless channel information of the soft robot in NLoS using the battery-free LED harvesters distributed outside the pipe. The entire system is enclosed by a flexible Faraday sleeve illustrated by the green color above. Sec. 4.5 describes the actual deployment of the WASER system. Sec. 4.2 describes the overview of WASER system.

integrated into LCE actuators for active LCE heating and cooling along with regenerative energy harvesting through the Seebeck effect [179]. Researchers have also demonstrated shape memory polymer and elastomer-made soft robots that can be heated by light sources [55, 56, 57, 58, 59, 60]. However, these systems require line of sight with the light source and are therefore not robust to obstacles or occlusions. In general, existing LCE actuation solutions are constrained by their requirement for a tethered electrical connection to an external power supply or line-of-sight positioning with a heat or light source. Such constraints limit the mobility of LCE-powered soft robots in various real-world and unstructured environments that require operating within confined and enclosed spaces (e.g., inside a pipe, behind a wall or within machinery).

To address this limitation with LCE-based soft robots, we have developed WASER (**W**ireless **A**ctuation for **S**oft **E**lectronics-free **R**obots), the first wireless heating platform that enables high-resolution heating actuation for soft robots within a non-line-of-sight context. Specifically, we build a programmable high-power beamforming platform using off-the-shelf RF hardware. WASER induces the heat actuation of soft robots wirelessly by applying 2.4 GHz wireless signals across an array of antennas. WASER performs wireless heating within a flexible Faraday sleeve to ensure system safety and minimize microwave leakage. The ultimate goal of WASER is to actuate the soft robot with speed and precision. For example, WASER aims to heat different parts (1 - 2 centimeter wide) of the robot body to enable various shaping changing and motion capabilities.

System Overview: WASER is developed through an inter-disciplinary effort that brings together wireless systems and material science research. On the wireless front, WASER builds on distributed beamforming technology, which focuses wireless energy towards a targeted area in a 3-D space [3]. WASER generates optimized sets of beamforming weights that beam maximized wireless power to the desired area of the robot in the non-line-of-sight while avoiding rising temperature outside the targeted area. On the ma-

terials front, we combine thin film LCE with soft, mechanically compliant electrically conductive layers. Together, these materials allow for electronics-free soft robots that harvest 2.4 GHz wireless energy and convert this stimulation into shape change and motion. Our WASER soft robot achieves a significantly improved thermal and electrical conductivity at 2.45 GHz without detrimentally altering its mechanical or shape-morphing properties. A detailed experimental evaluation shows that our WASER heating platform achieves 17.8 mm heating resolution[1]. We show that WASER can actuate desired shape deformation of the robot (starting from room temperature) within 1.5 seconds. Further, we demonstrate multiple modes of motion and shape change that can be achieved with the WASER robot.

Our primary challenge in WASER's design is to precisely actuate the desired part of a soft robot. In other words, our framework needs to find the optimal beamforming vector applied across the antenna array to beam maximized energy towards the desired areas of the soft robot while suppressing temperature rises outside the targeted area. Traditionally, we could infer the optimal beamforming vector by exchanging wireless channel information between the target and the beamforming platform. However, our target – a battery-free and electronics-free soft robot – does not explicitly provide its channel feedback to our platform. Further, the soft robot is in a non-line-of-sight configuration, so inferring the optimal beamforming vector of the robot becomes more challenging. To solve this problem, WASER designs creative alternatives to solicit indirect feedback of the robot, coupled with novel approaches to model signal propagation in the non-line-of-sight radio environment. Specifically, we estimate the channel of the WASER robot by deploying a grid of battery-free 2.45 GHz LED harvesters. A detailed description of designing and fabricating the 2.45 GHz LED harvesters can be found in Sec. 4.3.4. The brightness of the distributed LED array represents the wireless channel information around the deployed soft robot at multiple known locations. Note that WASER assumes the relative positions of the soft robot and the distributed LED array is known before any actuation is applied. WASER then can interpolate the wireless channel of the soft robot based on the captured brightness of LEDs at multiple known locations. Yet, the reflectors in the non-line-sight could generate multipath reflections and influence the channel estimation of the robot. To address this challenge, WASER applies multiple pre-generated sets of extremely low-power channel probing signals to estimate the radio environment. In effect, this allows WASER to derive the optimal beamforming vector that precisely beams to a particular body part of our soft robot. After estimating the channel, WASER further uses the distributed LED array to ensure a minimal wireless energy leakage to the environment preventing rising temperature outside the targeted area.

In addition, WASER must improve the speed of LCE-based soft robot actuation – i.e. reduce the heating time required for the robot to exhibit a maximized shape change for a reasonable total transmitted power.

[1]The heating resolution is defined as the distance that the energy density attenuates 3 dB from the peak (Sec. 4.6.6).

Previous work has shown that it takes 40 seconds to actuate an LCE-based soft actuator with a directional hundred-Watt microwave power source [180]. This slow pace of actuation limits its application in LCE-based soft robotics. The reason for this inefficacy is that prior work uses dielectric heating which is not efficient to convert electromagnetic signals into thermal energy within LCE polymers. Instead, WASER utilizes a novel material architecture and structure that combines highly conductive liquid metal (LM) alloy with the liquid-crystal composites. WASER then optimizes the geometry of soft robots so that it resonates well with 2.45 GHz RF signals. As a result, the robot can be heated based on both dielectric heating and induction heating which induces strong eddy current within the conductive component of the soft robot. Another challenge WASER must address is to avoid edge effects of the conductors in a strong EM field. The edge effects cause the electrons to concatenate on the edges and corners of the soft robot, which limits the shape changing capabilities of the soft robot. To address this problem, we design several non-uniform LM-based material structures for the robot. Sec. 4.4 describes our design choices that improve the speed of actuation to exhibit the desired shape changes.

We build a prototype of a multi-antenna high-power beamforming platform for actuating battery-free soft robots. Our prototype uses four USRPs paired with solid state amplifiers which gains 30 dB at 2.45 GHz. Our experiments show that WASER can heat the targeted region from 20°C to 60°C within 1.5 seconds. A detailed experimental evaluation on WASER shows that the heating platform achieves a high-precision heat distribution with an average spatial resolution of 17.8 mm. We also demonstrate shape morphing capabilities and motions modalities enabled by WASER which we evaluate in Sec. 4.6.8.

Potential Applications: We envision the WASER actuation platform that can function as an open-source tool available for broad uses by researchers in wireless, soft robotics, and material science in areas including smart manufacturing, food processing and machine maintenance. Indeed, we see WASER as a primitive that enables many applications, some of which we elaborate on in Sec. 4.7.

Safety: We limit the maximum transmit power of WASER to 20 Watts (43 dBm) which is the typical transmit power of a 4G cellular antenna [181]. WASER operates at 2.4 GHz whose wireless power reduces by about 40 dB in the first meter [182]. Thus, the power level beyond 1 meter is safe for general public exposure. To guarantee the safety and avoid jamming the existing Wi-Fi devices, we deploy our system in a flexible Faraday shield. We describe more details about the Faraday shield and possible methods to reduce the transmit power in Sec. 4.5 and Sec. 4.7 respectively.

4.1 Related Work

4.1.1 Battery-free Soft Robots

Inspired by nature, research in the emerging field of soft robotics proposes soft-bodied robots that exclude any rigid components. These soft robots have the potential to exhibit unprecedented adaptation to harsh or uncertain environments [183]. For example, a soft robot promises to achieve motion with high curvature and demonstrates twisting, bending, and folding. In addition, removing rigid energy sources such as batteries offers a promising path to robot miniaturization [184, 185] and thus helps the soft robot operate in confined spaces [186]. Given these unique features and advantages, soft robots require external power supplies either in the form of onboard batteries or are tethered directly to external power supplies to move and exhibit shape changes. Wired systems pose difficulties for systems trying to navigate complex terrain as power and data lines may impede movement. Soft robot actuators have poor energy efficiency [187, 188, 189], and for untethered systems using onboard batteries, this can lead to short ranges and the need to return for charging. Separately, pneumatic and fluid based soft robots require bulky pumps and compressors that must be attached by tubing to the robot [190, 191]. Previous work has performed actuation using light sources requiring line of sight to the robot [58, 59, 60]. WASER aims to build a remote microwave-based actuation platform for soft robots that simplifies precise energy delivery in confined spaces. This platform aims to develop robotic systems that can enter hibernation for a long period of time and reactivate when WASER energy is applied without the need to change batteries or retrieve the system for charging or maintenance when in hard-to-reach locations.

4.1.2 Thermally Responsive Soft Robot Actuation

Our study explores controlling thermally responsive soft materials that can exhibit complex shape changes upon thermal stimulation. We specifically focus on LCE due to its popularity in soft robotics and versatility in manufacturing, shape programmability, and device-level integration [174]. Past efforts with LCEs in soft robotics have typically relied on Joule heating to induce heat and shape change. This is accomplished by embedding or coating the LCE with electrically resistive materials that conduct electricity [192, 50, 53, 51]. However, this method requires a wired connection to an external DC power supply, which is needed to continuously provide several Amperes of current flowing through the resistive elements. This dependency on a tethered connection to external hardware significantly limits the mobility of robotic systems. Additionally, researchers in soft robotics apply the heat to the thermally responsive soft robot by placing it on a hotplate [193, 194, 195] or using a hot air gun [196, 197]. However, this ap-

proach requires manual hand input. Recently, light-based stimuli methods have been explored in the soft robotics community using scanning lasers [58], ultraviolet [56], and even infrared lights [198]. However, these constrain the operation space of the robot to be within line-of-sight of the light emitters.

More recent work [180] proposes using high-power microwave signals at 2.4 GHz to heat polymer-like materials (e.g, LCE). Different from the previous methods, this method involves an energy conversion from microwave energy to thermal energy. The energy conversion is a process in which wireless energy is absorbed by dielectric or conductive materials. This proposed method [180] can actuate the LCE actuator in non-line-of-sight (NLoS), but blasts electromagnetic energy blindly and fails to precisely and efficiently actuate the LCE (i.e., it takes 40 seconds to actuate the robot). In contrast, WASER strives to improve the precision and efficiency of actuating the thermally responsive soft robots in NLoS by developing a microwave beamforming actuation platform.

4.1.3 Primer on Microwave Heating

Over the past few years, researchers in the wireless system community have widely explored the capabilities of 2.4 GHz unlicensed microwave frequencies for wireless networking and sensing. Historically, however, the ISM (Industrial, Scientific and Medical) band was initially allocated primarily for non-communication purposes, such as microwave oven heating. This section describes the applications and mechanism of microwave heating.

4.1.4 Applications of Microwave Heating

One of the most commonly used microwave heating devices is the domestic microwave oven. The microwave oven cooks food by blasting wireless power (between 0.5 kW and 2 kW) at a center frequency of 2.45 GHz. Industrial scale heating is another major sector where microwave heating techniques are deployed, say for plastic softening [199], drying [200], and tempering deep frozen meats [201]. In addition to industrial manufacturing and food processing, microwave heating has also been used to perform sterilization, such as sterilizing food [202], medical/dental instruments [203] and N95 filtering respirators [204]. Industrial microwave systems either use 2.45 GHz or 915 MHz wireless signals within the transmitted power range of 3 kW and 120 kW. Recent work has also used microwave heating to facilitate chemical reactions [205]. The benefits of applying microwave heating are that they enable deep internal heating and offer a significant reduction in heating time.

4.1.5 Mechanisms of Microwave Heating

When high-power wireless signals interact with an object, a part of the electromagnetic energy is absorbed and converted into thermal energy or heat. This conversion can be attributed to either dielectric (polarization) or conductive losses, leading respectively to dielectric and induction heating. The relative contributions of both these heating mechanisms to an object are heavily dependent on the underlying materials of that object. We describe the details of each microwave heating mechanism below:

Dielectric Heating: For dielectric materials or electrical insulators such as pure water, there is no free electric charge flowing through the material. However, a rapid changing electric field (e.g., the E field component of 2.4 GHz electromagnetic signals) can cause dielectric polarization which generates heat within the material. This is because the molecular dipoles inside the material are rendered polarized meaning that each molecular dipole (e.g., a water molecule) have one end carrying positive charges and the other end carrying negative charges. Note that a molecular dipole is always aligned with the direction of an external electric field. Thus, emitting rapid oscillating electromagnetic fields, such as 2.4 GHz wireless signals causes the dielectric material's molecules to rotate rapidly – as they try to align themselves with the external electric fields. When molecules rotate, the collision among them creates friction, which dissipates energy and contributes to heating. The dielectric loss is governed by the relative permittivity and the loss tangent of the underlying material [160]. Dielectric heating is the main contributor of thermal energy when we cook our food using a microwave oven.

Induction Heating: Also known as eddy current heating, induction heating is the most dominant heat contributor for conductive materials. Since there are a large number of free electrons in conductive materials, a rapidly changing magnetic field (e.g., the H field component of 2.4 GHz electromagnetic signals) will induce currents inside a conductor [206] following Faraday's Law of Induction. The continuous induced current can result in the heat of the conductive materials, guaranteed by Joule's First Law. Therefore, a material which has a higher electrical resistivity will generate a higher eddy current loss. In addition to this, the oscillating magnetic fields can also contribute to hysteresis losses if the underlying materials have high values of loss tangent and permeability. This phenomenon generates heat as the magnetic moments within materials try to oppose to the rapid changing magnetic fields in the environment.

Microwave heating of real-world materials is usually a combination of the two heating mechanisms above. Dielectric objects such as plastic and LCE barely consist of any free electrons, so they are heated dominantly due to dielectric heating. Meanwhile, the reason that conductive objects can heat extremely fast in a microwave oven is owing to induced eddy current.

4.2 Overview and Chapter Outline

Fig. 4.1 summarizes the setup of our system. Our objective is to design a soft robotic platform that can precisely heat the soft robot at fine spatial granularity and achieve this heating within a few seconds from an external wireless energy source. We assume that the robot could potentially be in non-line-of-sight, e.g. within a pipe. We assume that the entire system is enclosed within a Faraday shield or enclosure that prevents microwave leakage to ensure safety. We further assume that the location of the soft robot itself is known and that it is designed to move predictably when heated to specific temperatures. We assume access to at least some locations within the Faraday shield that are in line-of-sight of our system to place additional sensors (LED harvesters). We note that as shown in Fig. 4.1, the harvesters can potentially entirely be outside the plastic pipe or other non-metallic enclosed space that the soft robot navigates. Should the enclosed space be metallic, WASER re-uses it as the Faraday shield and the LED harvesters would need to be placed within the enclosed space.

The rest of this chapter addresses two key design questions of WASER: (1) First, we develop a high-power beamforming solution at our wireless energy source that efficiently heats the desired points on the robot, without requiring electronics physically attached to the robot (Sec. 4.3); (2) Second, we present the design of the soft robotic platform's material that allows it to remain flexible while conducting heat effectively (Sec. 4.4). We conclude with a detailed system implementation and evaluation (Sec. 4.5-4.6).

4.3 High-power Blind Beamforming

LCE-based soft robots can be remotely actuated through the application of wireless heating. However, wireless channel feedback from the soft robots is unavailable, meaning that it is challenging to accurately actuate the desired spots of them via RF beamforming. The situation gets even worse when the robot operates in non-line-of-sight scenarios. To address this challenge, we introduce a high-power *blind* beamforming algorithm[2] that can precisely actuate the soft robot without any response of the robot. Specifically, WASER achieves this by (1) acquiring accurate estimates of channel state information indirectly from multiple desired locations using battery-free LED harvesters surrounding the robot (potentially at a visible place, e.g. outside a pipe); (2) interpolating the channel distribution to estimate the wireless channel at the moving robot.

[2]Blind due to the absence of direct feedback from the robot.

4.3.1 Problem Formulation

Our goal is to perform accurate beamforming towards the robot, even focusing the wireless energy at individual sub-sections of the robot body (see Sec. 4.4). As a result, WASER must infer the wireless channel state information at different parts of the soft robot.

We now mathematically characterize the relationship between wireless channels and beamforming weights in the soft robotic context. We represent the channel state information at a desired point on the soft robot by a complex vector $\mathbf{h} \in \mathbb{C}^N$, where N is the number of antennas. The n-th element of $\mathbf{h}$, denoted by $\mathbf{h}_n$, then represents the channel observed by the n-th antenna. The amplitude $|\mathbf{h}_n|$ denotes the channel strength observed by the antenna n, and the phase $\angle\mathbf{h}_n$ refers to the phase of the wireless channel between the n-th antenna and the robot. We use $\mathbf{w} \in \mathbb{C}^N$ to denote the *precoder*, which represents the known beamforming weights applied to the transmitted signals from the antenna 1 to antenna N. In order to beam maximized wireless energy to the robot, WASER must find out the optimal precoder $\mathbf{w}$ to apply across the antennas:

$$\max_{\mathbf{w}} \quad \left|\mathbf{w}^\top \mathbf{h}\right| \tag{4.1}$$

$$\text{s.t.} \quad |\mathbf{w}| \leq P \tag{4.2}$$

where Eqn. 4.2 restricts the total transmitting power P from all distributed antennas. Given an observed wireless channel $\mathbf{h}$, the solution to the optimization Eqn. 4.1 above can be derived as

$$\mathbf{w_o} = P\frac{\bar{\mathbf{h}}}{|\mathbf{h}|} \tag{4.3}$$

where $\mathbf{w_o}$ is the optimal beamforming vector that beams maximized wireless energy to the robot and $\bar{\mathbf{h}}$ is the complex conjugate of $\mathbf{h}$. Ideally, the wireless channel information $\mathbf{h}$ between the soft robot and the distributed antennas can be measured by sending a known series of signals from the specific point on the soft robot. However, LCE-based soft robots are battery-free and electronics-free and so cannot send any feedback.

4.3.2 Low-power Channel Probing

To ensure that the LCE-based soft robots remain electronics-free, WASER must infer the wireless channel $\mathbf{h}$ of the soft robot indirectly. To address this challenge, WASER's key approach is to deploy an array of distributed 2.45 GHz LED harvesters at known locations outside the work space of the soft robots, including outside of any enclosure that surrounds the robot (such as a pipe). Specifically, WASER first obtains the wireless channel information at each location that these LED harvesters are deployed. The

channel information can be measured by the brightness of the harvesters. We then apply various sets of precoding weights across all transmit antennas, and estimate the corresponding wireless channel of the LED harvesters based on the brightness of LEDs.

Specifically, we design a novel closed-loop probing algorithm using an extremely low power to probe the LEDs at desired locations. The challenge however is that the brightness of LEDs can only contain the amplitude of the wireless channel. In other words, we cannot get the phase information of the wireless channel by only relying on the brightness of LED arrays.

WASER therefore rephrases the channel state information estimation problem into a *phase retrieval* problem. At a high level, WASER has to probe the LED harvesters with multiple precoders $\mathbf{W}_k$ where k represents the k-th probe. Mathematically, WASER obtains the phase of wireless channels of LED harvesters $\mathbf{h}$ by solving the following optimization:

$$\min_{\mathbf{h}} \quad ||\left|\mathbf{W}^\top \mathbf{h}\right| - \mathbf{b}||_2^2 \tag{4.4}$$

where $\mathbf{W} = [\mathbf{W}_1, \mathbf{W}_2, \cdots, \mathbf{W}_K]$ is the concatenation of all precoders, and $\mathbf{b}_1, \mathbf{b}_2, \cdots, \mathbf{b}_K \in \mathbb{R}$ is the received signal strength (RSS) of the LED harvester in each probe. To solve the optimization above, traditional solutions like matrix factorization [207] and low-rank property [208, 209] are not suitable, since these solutions assume a free-space radio environment. Considering an enclosed multipath-rich reflective radio environment, WASER adapts a classic phase retrieval solution called PhaseLift [210] to solve the optimization Eqn. 4.4.

While Eqn. 4.4 is a well-known NP-hard problem [211, 212], it can be solved by being reformulated as a matrix recovery problem. We specifically expand $\mathbf{h}$ into a matrix $\mathbf{H}$, and solve the dual problem of Eqn. 4.4 by minimizing the rank of $\mathbf{H}$. This problem can be alternatively posed as a trace minimization problem that minimizes the trace of $\mathbf{H}$. This is a semi-definite programming problem can be solved using convex programming algorithms [213].

Number of Probing Times: As reported by [210], the PhaseLift algorithm typically has to apply $t \times N$ number of probes to converge, where $t = 2, ..., 4$ and N is the number of antennas used in beamforming. Thus, WASER uses $4 \times N$ times of probes to guarantee the convergence of its algorithm. The actual probing signals (precoders) are randomly generated complex Gaussian vectors. The probing results are normalized to have a mean of one to meet the requirement of the PhaseLift algorithm. Normalizing the probing results as well as the wireless channel $\mathbf{h}$ does not matter in our experiments, since we only want to infer the relative strength and phase difference of different elements in $\mathbf{h}$ for beamforming, rather than absolute values. In our experiments, we deploy 8 LED harvesters and apply 16 times channel probing to calculate

the corresponding wireless channels. WASER uses four RF antennas, and the overall computation latency to solve the optimization Eqn. 4.4 takes less than 1 second using embedded devices.

4.3.3 Finding the Optimal Beamforming Vector

At this point, WASER can acquire wireless channel estimates based on LED harvesters measurements at multiple desired locations around the robot. We recall that while the robot itself may be in non-line-of-sight, the harvesters are assumed to be in line-of-sight relative to WASER (e.g. outside the robot's enclosure). Hence, WASER must estimate the channel at a desired point on the robot potentially in non-line-of-sight. It can then use this to infer the optimal beamforming vector to beam maximal power at the robot. WASER achieves this by interpolating the channel state information where the robot is located based on the calculated channels of harvesters.

Interpolating Channel Information: In compliance with FCC regulations, our high-power beamforming actuation must be performed inside a larger RF enclosed space (e.g., a Faraday sleeve) to prevent extensive leakage of the microwave energy. WASER's experiments use a flexible metal Faraday shield called the *cavity resonator*. For instance, a microwave oven is a cavity resonator that prevents the leakage of extra high-power EM waves. EM theory states that electromagnetic waves, with a steady emitter, will form a standing wave inside a cavity resonator, because the waves reflect at the closed ends. However, directly solving the electric field (E-field) and magnetic field (M-field) at the standing wave involves solving a differential equation with 3 unknowns, which is challenging. To address the problem, we derive the Transverse-Electric (TE) mode of the system where we assume $E_z = 0$ and the Transverse-Magnetic (TM) mode where we assume $H_z = 0$. The final solution would be a combination of the TE mode and TM mode. Due to the properties of the standing waves, the x, y, z partitions of the electric field and magnetic field can be represented by a sinusoid product of the coordinates. An analytical solution can thus be carried out using such property for the TE mode and TM mode [214].

Therefore, we can use a combination of sinusoidal waves to express the power distribution inside the cavity resonator by applying the Poynting's theorem on the TE-mode and TM-mode. The Poynting vector is defined as $\mathbf{S} = \mathbf{E} \times \mathbf{H}$, and we can compute power density, $S = |\mathbf{S}|$, as follows:

$$S(x,y,z) = \left|\mathbf{E}(x,y,z) \times \mathbf{H}(x,y,z)\right| = \frac{1}{2}\epsilon\left|\mathbf{E}(x,y,z)\right|^2 = C^2\left[\sin^2\left(\frac{\pi x}{a}\right)\sin^2\left(\frac{\pi y}{b}\right)\sin^2\left(\frac{\pi z}{l}\right)\right] \tag{4.5}$$

where C, a, b, l are all constants. Upon acquiring the wireless channel $\mathbf{h}$ at different reference points (x_p, y_p, z_p), we solve the following optimization problem at each antenna i to perform interpolation at the

remaining points:

$$\min_{C_n,a_n,b_n,l_n} \quad \sum_p \left| S(x_p,y_p,z_p) - \left| \mathbf{h}_n(x_p,y_p,z_p) \right| \right|^2 \tag{4.6}$$

which is locally convex and can be solved by gradient descent. The interpolated channel response of antenna i at any desired position (x,y,z), denoted by $\hat{\mathbf{h}}_n(x,y,z)$, would be:

$$\left| \hat{\mathbf{h}}_n(x,y,z) \right| = S(x_p,y_p,z_p) \tag{4.7}$$

Algorithm 3 summarizes our overall approach when M LED harvesters are used. We can similarly estimate the phase of channel response $\angle\hat{\mathbf{h}}_n(x,y,z)$ of the channel by linearly interpolating the angles of probed channel response $\mathbf{h}_n(x_p,y_p,z_p)$ at the reference points. For simplicity, in the simulation and experiments, we only focus on estimating channel on the $x-y$ plane and thus the z term is omitted as a constant phase shift, although our approach can be extended to 3-D. The design of our soft robot described in Sec. 4.4 has a planar shape so that the robot initially lies on one plane ($x-y$ plane). Also, we assume to have prior knowledge of its initial location, and accurate models on how its shape changes when actuated. The robot is designed to only move its parts of body along a 2-dimensional plane, and this allows us to make this simplifying assumption across varied applications.

Control Mechanism of a Continuous Shaping Morphing: The algorithm above interpolates and infers the optimal beamforming vectors for each point within the cavity regardless of the location of the soft robot. Note that WASER assumes that the location of the soft robot itself is known and that its basic movements are predictable (designed during the fabrication of the robot). Thus, WASER can apply a series of beamforming vectors consecutively to enable a series of shape changes.

Algorithm 3 WASER Algorithm

$\mathbf{W}_1, \mathbf{W}_2, \cdots, \mathbf{W}_K \leftarrow$ Randomly Initialized Beamforming Weights
for $k = 1$ to K **do**
    Beamform using $\mathbf{W}_k$, record the lightness of LEDs $L_1^{(k)}, L_2^{(k)}, \cdots, L_M^{(k)}$
end for
Use PhaseLift to solve $\mathbf{h}^{(1)}, \mathbf{h}^{(2)}, \cdots, \mathbf{h}^{(M)}$
for $n = 1$ to N **do**
    Use interpolation to find C_n, a_n, b_n, l_n according to $\mathbf{h}_n^{(1)}, \mathbf{h}_n^{(2)}, \cdots, \mathbf{h}_n^{(M)}$
end for
Loop:
Check the current position of soft robot (x,y,z)
Use C_n, a_n, b_n, l_n to find $\mathbf{h}(x,y,z)$
Beamform using $\mathbf{w_o} = P\frac{\bar{\mathbf{h}}}{|\mathbf{h}|}$
end

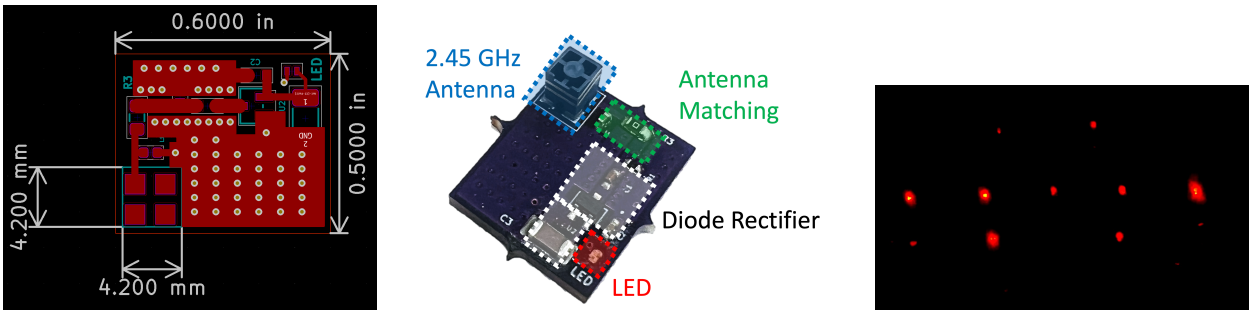

(a) PCB of WASER LED Harvester (b) WASER LED Harvester (c) Responses of the LED Array

Figure 4.2: WASER's Battery-free LED Harvesters.

4.3.4 Design of 2.4 GHz LED Harvesters

This section describes the design of the 2.4 GHz LED harvesters. We place an array of distributed LED harvesters in line-of-sight of WASER to infer the channel of the soft robot in non-line-of-sight. Fluctuations in energy at the harvesters allow us to respond in real-time to apply accurate high-power beamforming. Our objective is to design extremely low-power and low-cost LED harvesters that can respond to 2.4 GHz wireless signals.

Why LED Harvesters?: At first glance, one may choose to infer the channel distribution by deploying an array of RF/Wi-Fi backscatter instruments that modulate the channel information onto their reflected RF signals. While RF/Wi-Fi backscatter could dissipate less power than a wireless harvester equipped with an LED, the wireless signals reflected back from RF/Wi-Fi backscatters could be overwhelmed by the high-power signals transmitting from the beamforming actuation platform. This causes interference and requires an additional expensive RF receiving unit to handle high-power received signals so that it can process the signals from RF backscatter. Instead, WASER applies LED harvesters to indirectly capture information about the wireless channel at 2.4 GHz at multiple locations. The brightness of the LEDs can simply be read using a low-cost camera. By collecting our energy detection measurements through visible light signals rather than RF, our system functions even if the beamforming actuation platform is actively transmitting extremely high-power wireless signals.

Designing Wireless LED Harvesters at 2.4 GHz: To obtain visible light-based feedback for inferring the channel distribution, WASER uses a simple battery-free wireless energy harvester equipped with an ultra-compact low-power LED. Off-the-shelf components are chosen to keep cost and size down. Fig. 4.2 (a) shows the PCB layout of the WASER LED harvesters. Fig. 4.2 (b) describes the main components of our LED harvester: an antenna, the matching network, the diode rectifier, and the LED. WASER harvesters absorb wireless energy using 2.45 GHz omni-directional chip antennas [215] that only have a dimension

of 3 mm $\times$ 3 mm $\times$ 4 mm each. To ensure the antenna harvests at 2.4 GHz efficiently, we place an LC impedance matching circuit using datasheet-recommended component values for our antenna. RF power is rectified using a double-diode detector circuit with a BAT15-04 Schottky diode pair to charge a ceramic supercapacitor. The ultra-compact LED is connected to the supercapacitor, so that the brightness of the LED directly captures how much wireless power the harvester absorbs at a specific location. We choose to use red LEDs over other colors, since the power consumption of red LEDs is lower at the same luminous intensity [216]. Fig. 4.2 (c) shows the visible light responses from the LED harvester. According to the specification of the LED, the brightness is linear with the voltage harvested. Therefore, WASER uses the LED brightness to represent the magnitude of the channel at the distributed LED harvester locations as we described in Sec. 4.3.2. Note that our LED harvesters have a threshold of received signal power > 11 dBm to respond.

4.4 Design of Wirelessly Actuated Electronics-free Soft Robots

In this section, we seek to optimize the material of the soft robot to improve the speed of the actuation. While our accurate high-power beamforming algorithm in Sec. 4.3 is one approach to improve heat transfer to the robot, its actual material architecture also plays an equally pivotal role. To see why this is important, note that prior work [180] takes around 40 seconds to actuate an LCE-based soft uniaxial actuator using an 800 Watt maximum output microwave power. In contrast, to ensure safety and minimize microwave leakage, WASER limits itself to 20 Watt maximum total output wireless power. This section describes how: (1) WASER introduces novel material structures to significantly enhance heating efficiency; (2) WASER consists of materials and fabrication methods that ensure the soft robot's mechanical compliance and deformability.

4.4.1 Structural Design of Soft Robots

In this section, we explore a structural design that maximizes heating efficiency of the robot. Recall from Sec. 4.1.5 that the heating efficiency of dielectric heating is much lower than induction heating. However, today's material structure of soft robots is liquid-crystal composite – a polymer-like dielectric material. Thus, the energy conversion of the current soft robots is largely attributed to the dielectric losses.

Combining Shape-Morphing Polymer with Liquid Metal: WASER utilizes a novel material architecture that combines LCE materials with a gallium-based LM alloy that is highly conductive and can be patterned into a thin film. The key idea is to leverage both dielectric heating and induction heating mechanisms when the 2.4 GHz wireless signals are applied to the soft robot. Specifically, a thin layer of liquid metal

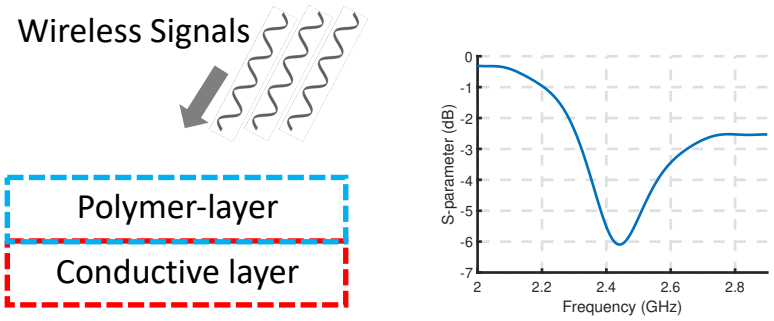

(a) Layered Structure (b) Frequency Response

Figure 4.3: (a) The cross-section of WASER's layered material structure. WASER uses a layered material structure to build the soft robot including a polymer layer and a conductive layer (see details about the materials in Sec. 4.4.2). 2.4 GHz wireless signals can induce heat using dielectric and induction heating for the polymer and conductive layer respectively. The conductive layer will generate more heat than the polymer layer due to the strong induced eddy current. Thus, the heat induced on the conductive layer will also further transfer to the polymer layer to facilitate the shape change of the polymer layer (LCE). Note that wireless signals at 2.4 GHz can penetrate through dielectric materials such as LCE. (b) The S11 frequency response of our soft robot design simulated by CST Studio.

composite is applied to the lower layer of the soft robot shown in Fig. 4.3 (a). As described below in Sec. 4.4.2, the composite is composed of eutectic gallium-indium LM alloy fused with a thin coating of silver ink. This layer induces induction heating that also transfers to LCE materials via conduction – allowing for rapid shape changes to the soft robot. The conductive layer is built in a facing-down manner to reduce reflection of the wireless signals and thus improve the energy absorption efficiency.

Optimizing the Geometry of the Soft Robot: The geometry of the liquid metal layer directly influences heat transfer efficiency. To see why, consider for example various frequencies of commercial antennas such as Wi-Fi, mmWave, and FM antennas that have different length and width. Likewise, WASER must carefully design the geometry of the liquid metal composites so that they resonate well with 2.4 GHz wireless signals like Wi-Fi antennas do. In addition, WASER must optimize the geometry within the constraints of a limited volume since adding excess weight to the soft robot limits its mobility.

To achieve an optimized geometry, WASER uses an empirical approach. Specifically, we simulate various geometries of the soft robot that resonate well with the 2.45 GHz ISM band. Fig. 4.3 (b) plots the frequency response of a design simulated by CST Studio. Then we build the prototypes to test the performance (e.g., temperature distribution) using our beamforming platform. From our optimization, we show a planar design with a thickness of 50 microns for the polymer layer and 5 microns for the liquid metal layer enabling improved heating efficiency.

Combating Edge Effects of the Conductive Layer: Microwave energy can raise the temperature of the conductive layer at ultra-low latency due to induced eddy currents. However, a flip side of the eddy

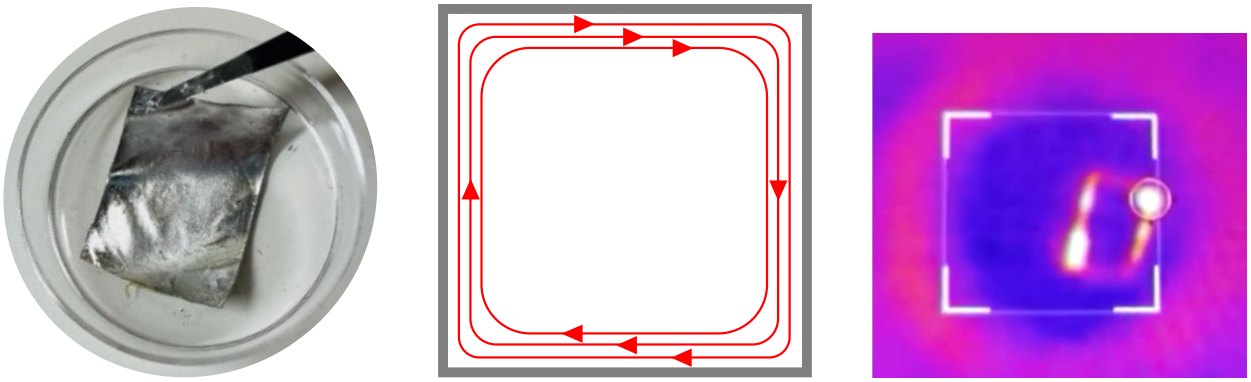

(a) Uniform Liquid Metal (b) Eddy Current Path (c) Edge Effects

Figure 4.4: (a) A rectangular sample with uniformly distributed liquid metal. (b) Red lines show the path of induced eddy currents on the rectangular sample. (c) Heat distribution captured by a thermal camera: the figure shows the heat distribution of the sample with uniformly distributed liquid metal when the 2.4 GHz wireless signals are applied. The brighter color shows a higher temperature. We note that edge effects occur at the corners and edges of the uniform liquid metal conductive layer.

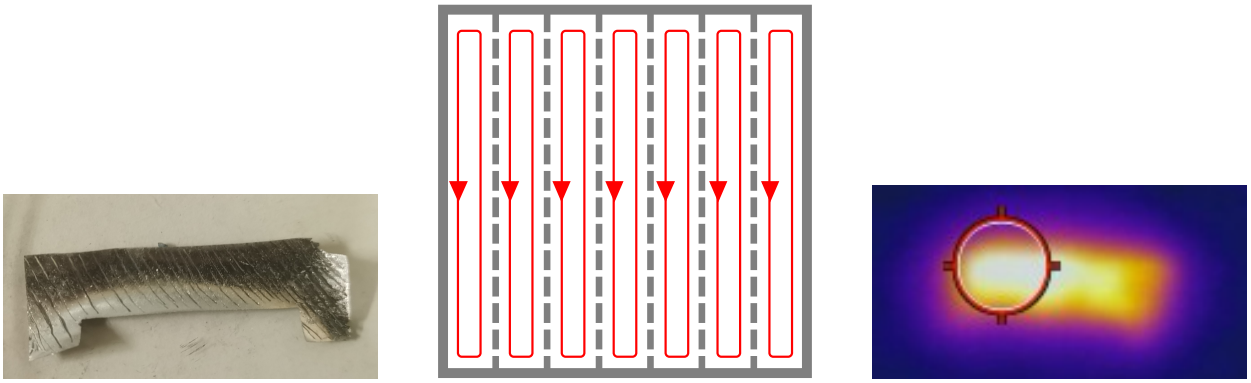

(a) Patterned Liquid Metal (b) Eddy Current Path (c) Mitigated Edge Effects

Figure 4.5: (a) A sample with patterned/textural liquid metal layer. (b) The path of induced eddy currents on a patterned liquid metal layer. (c) The heat distribution of the patterned liquid metal layer shows that the edge effect is significantly mitigated, and the high temperature is uniformly distributed through the entire sample.

current is that it tends to drive the electrons to corners and edges of the conductor (e.g. in Fig. 4.4 (b)) [217], leading to these regions being disproportionately heated. Fig. 4.4 (c) shows a heat distribution of a *uniformly distributed* liquid metal composite (shown in Fig. 4.4 (a)) when 2.4 GHz wireless energy is applied. We notice that only the edges and corners are heated to a high temperature. This phenomenon significantly distorts the heat distribution on the soft robot which degrades the actuation performance of the whole soft robot.

To combat the edge effect, WASER designs a fine-grained non-uniform texture on the liquid metal layer instead of using a uniformly distributed composite. The key idea is to design thin laminates of liquid metal conductors on the conductive layer as shown in Fig. 4.5 (a). Using this approach, the eddy

current will be induced within each lamination to mitigate the overall edge effects. Fig. 4.5 (b) shows the illustration of induced eddy current in our soft robot design and Fig. 4.5 (c) shows a much more uniform heat distribution, compared to the naïve design in Fig.. 4.4 (a). We implement multiple textures of the conductive layer and evaluate their heating efficiency in Sec. 4.6.2.

4.4.2 Soft Robot Material and Fabrication

This section describes the material architecture and method for fabricating the WASER soft robot.

Material Selection: To address the needs of a material architecture that can both be rapidly heated by wireless signals and exhibit repeatable actuation, we combine both liquid crystal elastomers (LCE) with a thin film of mechanically deformable conductive material. LCE was selected for its ability to be a soft actuator that can contract and expand based on thermal stimuli. LCEs are a class of shape memory polymers in which loosely crosslinked polymer networks of monomer mesogens exhibit controllable and reversible actuation [218]. These networks of mesogens reorient themselves directionally in a monodomain (ordered) anisotropic state in which polymer chains are aligned along one direction, or in a polydomain (disordered) isotropic state in which chains are aligned in all directions. These states are controlled by heating and cooling this material above and below its nematic to isotropic transition temperature. Below this temperature, in the nematic state, LCEs exhibit liquid like properties more similar to elastomers as the monodomain polymer networks slide past each other freely. Upon heating to the isotropic phase, these mesogens align themselves in a disordered state similar to a disordered crystalline structure, thus exhibiting properties closer to crystalline solids. At the macroscopic level this realignment induces controllable contractions and expansions making LCEs excellent for soft robotic applications. For this application, LCEs are fabricated in thin films (50 μm) that upon heating contract along one director and induce bending. To induce this uniform contraction polarized light is used to photopolymerize and align the mesogens along one director. For the conductive layer, a 3-6 μm coating of silver (Ag) and eutectic gallium-indium (EGaIn) liquid metal is applied to the surface of the thin film LCE. This thin layer does not impact mechanical performance due to its compliance and deformability. In this way, it allows for wireless heating of the LCE without interfering with its shape change and motion.

Fabrication Method: Fabrication of these thin film soft actuators is divided into three parts: (1) LCE synthesis; (2) LCE photoalignment; (3) deposition of soft conductive films used for wireless heating. Steps 1 and 2 are based on [175], while Step 3 is adapted from [219, 220]. The LCE is fabricated by mixing 1.500 grams of 1,4-bis-[4-(6-acryloyloxyhexyloxy)-benzoyloxy]-2-methylbenzene monomer (RM82, Wilshire Chemicals), 0.116 grams of n-butylamine chain extender molecule (Sigma Aldrich), and 0.025

grams of Irgacure I-369 photoinitiator. This was heated using a heat gun until the materials had dissolved. This is followed by three cycles of vortex mixing (Scilogex MX-F) for 40 seconds and heating. For the photo alignment setup, two 75×50 mm glass slides were cleaned using IPA and Acetone. To guarantee good bonding between the photoalignment dye and glass surface, the slides were treated with oxygen plasma at at 50 mW power and 100 mTorr pressure for 2 minutes. The photoalignment dye, brilliant yellow (Sigma Aldrich), was dissolved by 1% weight in DMF (Sigma Aldirch). Using a 0.45 μm filter, this solution was coated onto the surface of the slides and then spin coated for 750 rpm for 10 seconds and 1500 rpm for 30 s. Using a polarizing lens (Thorlabs LCRM2/M) placed in front of a projector (Vivitek) six inches apart, each slide was placed in front of the polarizer to photo-align along one direction. Once the two slides were photoaligned along the same director, they were bonded together with 50 μm spacer separating them. The bonded slides were then placed on a hot plate at 80 °C with the LCE material inserted between the slides using capillary action. After the cell has been filled, the setup is moved to a 70 °C oven for 16 hours for oligomerization. In order for polymer crosslinking to occur, and induce bending, the setup was placed in a UV chamber (UVP CL-1000 ultraviolet crosslinker) at 254 nm for 20 minutes on each side. The slides were then broken open and the LCE was removed. The LCE was then placed on a slide, coated with polyvinyl alcohol (PVA) (5 wt % aqueous solution), and spin coated for 1 minute at 2000 rpm. PVA was used to guarantee a smooth and uniform conductive layer. Next silver ink (Metalon JS-A102A) was airbrushed (Master airbursh G22) for 40 passes at 6 inches to place a uniform coat of silver ink into the LCE. The solvent was then evaporated for 12 hours at 60 °C. In order to make this surface conductive for WASER heating, eutectic Gallium Indium liquid metal (EGaIn) is wetted onto the deposited Ag surface. This combination of Ag and LM creates percolating pathways between Ag flakes and produces a continuous Ag-In-Ga film [219, 220].

4.5 Implementation

We build the first high-power programmable beamforming platform for wireless heating applications that will be available to researchers in wireless, robotics, material science, and HCI. The main components of the platform are software-defined radios, RF LDMOS integrated power amplifiers, an enclosed wireless heating stage that minimizes the microwave leakage, and the WASER software including an in-house safety monitoring software.

RF Chain: The platform shown in Fig. 4.6 consists of multiple RF chains. Each RF chain has one Ettus USRP X310 software-defined radio (SDR) as the signal generator. The SDR transmits signals at 2.45 GHz with a maximum output power of 20.26 dBm (0.106 W) that we verified using a spectrum analyzer. We

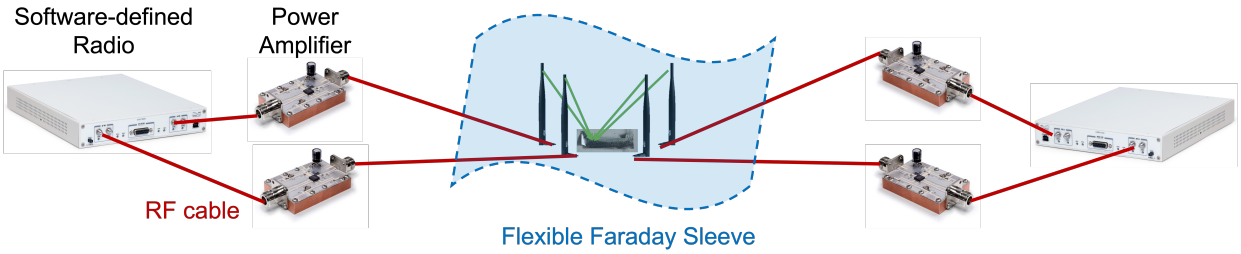

Figure 4.6: The WASER Platform.

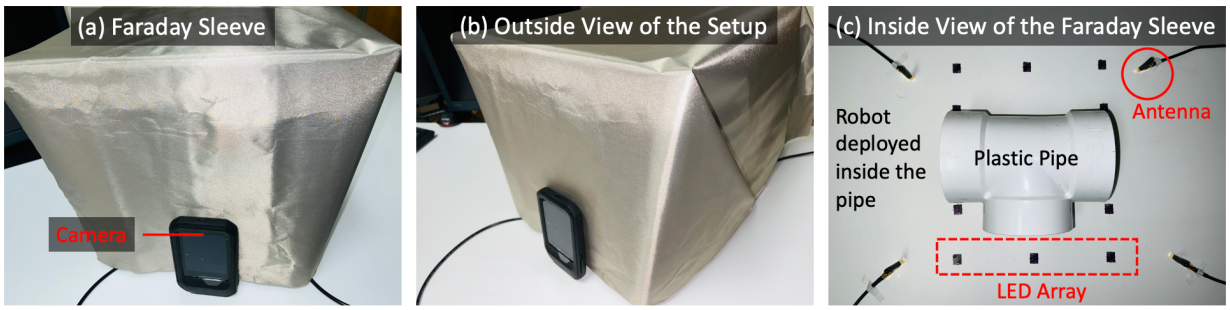

Figure 4.7: If the robot is deployed in **non-metallic spaces**, e.g., a plastic pipe, the system can be placed within a flexible Faraday sleeve. (a) & (b) The flexible Faraday enclosure. The Faraday sleeve can be scaled to volumes as large as 100 cubic-feet. The camera views the LED harvesters inside of the enclosure via a small opening (6 mm wide) of the Faraday sleeve that results in minimal leakage. (c) Top view of the system setup inside the Faraday enclosure. The robot is deployed inside the plastic pipe. Four antennas and ten LED harvesters are placed outside the plastic pipe.

then feed one RF LDMOS integrated power amplifier to the output of the SDR. While the SDR is capable of sending signals with a power of 20.26 dBm, we limit its output power below the maximum input power of the amplifier which is 20 dBm. The power amplifier provides a 29.7 dB linear power amplification to the 2.45 GHz RF spectrum when the input power is in the range of 1 dBm and 12 dBm [221]. At the end of each RF chain, we have a 2.4 GHz omni-directional antenna that supports a 50 W maximum output [222]. During the experiments, we have four RF chains in total, but the WASER platform can support six chains in maximum. The total transmitted power of WASER is limited to 20 W, but each chain is theoretically capable of providing 30dBm (1W) - 41.5dBm (14W) of wireless power.

WASER System Setup: Fig. 4.7 shows an example of the WASER system setup including the Faraday shield, a soft robot deployed inside a plastic pipe, and an array of LED-equipped wireless energy harvesters. Both antennas and the harvesters are placed within the Faraday shield. A camera views the harvesters via a small opening (6 mm wide) of the Faraday shield that results in minimal leakage, owing to the larger wavelength of microwaves [223].

As shown in Fig. 4.7 (c), should the pipe be made of non-metallic materials such as Polyvinyl chloride (PVC) plastic, wireless signals can easily pass through without severe penetration loss, and the LED

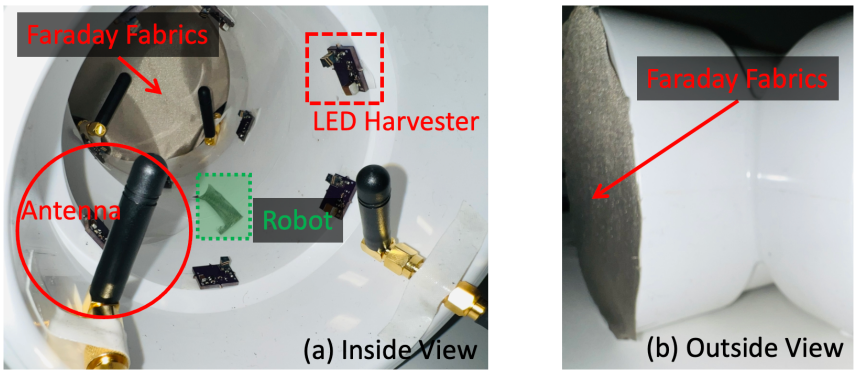

Figure 4.8: If the robot is deployed in **metallic spaces**, e.g., a metallic pipe, WASER reuses the existing metallic structures as RF shields. For example, a soft robot deployed to inspect metallic pipes does not require additional shields around the pipes themselves – only with minor modifications to the ends of the pipes covered by Faraday fabrics. The LED harvesters would need to be placed within the pipe with the camera viewing the lights from one end.

harvesters could be placed outside the pipe but still within Faraday shielding. In such a scenario, the entire system (the pipe, RF antennas, and LED-equipped harvesters) is enclosed by the Faraday shield to ensure minimum microwave leakage. Should a pipe (the enclosed space that the robot deployed) be metallic, the LED harvesters would need to be placed within the pipe with the camera viewing the lights from any one end (see Fig. 4.8 (a)). In this case, WASER re-uses the metallic pipe as the Faraday shield and only requires minor modifications to the setup – adding shields to the ends of the pipes (see Fig. 4.8 (b)).

Flexible Faraday Sleeves: Should the pipe be non-metallic, WASER uses a flexible, mobile, and low-cost Faraday sleeve to operate the actuation (see Fig. 4.7 (a) & (b)). The shielding sleeve is composed of double-layered EMF conductive fabrics [224]. It encloses the entire system including the robot, the RF antennas, and LED-equipped harvesters. The volume of the entire RF shielded space can be scaled to 100 ft^3. Note that this flexible Faraday fabric provides 40 dB attenuation to the wireless signals at 2.45 GHz and prevents the microwave leakage in compliance with FCC regulations. Indeed, our RF system can also integrate with other commercial portable EMI/RFI-shielded Faraday tents, such as the double-layered Holland Shielding Shielded Faraday Fabric [225] which provides 80 dB attenuation at 2.45 GHz.

WASER Software: WASER runs its real-time channel probing and beamforming algorithm in C++ on a Linux laptop. WASER captures the LED responses using a low-cost camera outside the microwave cavity. WASER implements the channel interpolation on MATLAB infused with electromagnetic and thermal modeling on CST studio. In addition to WASER's algorithm, we develop a safety monitoring software that monitors the temperature of power amplifiers, the thermal distribution inside the cavity, and the

microwave leakage outside the cavity.

Ground Truth and Baseline: We use a FLIR C5 Thermal Camera to measure the heat distribution, and use it as the ground truth to evaluate the heating resolution performance of WASER. We compare WASER against a baseline system, that performs *phase sweeping*. The baseline setup uses 4 antennas and 16 probes. During the probing process, each antenna keeps the precoder amplitude as the same value according to the power constraint, and shifts their phases among $0, \pi/2, \pi, 3\pi/2$. Therefore the total probing overhead would be 16 for each antenna to sweep through 4 phases. For each probe we record the signal strength according to the brightness of the LED harvester. After 16 probes, we select the phase combination of the 4 antennas that provides the maximum signal strength as the beamforming weight for the beamforming process. No interpolation is involved since such channel estimation is simply not accurate enough in practice. The baseline algorithm is denoted as *Sweeping* in our experimental results and graphs.

4.6 Results

4.6.1 Channel Probing Accuracy

We evaluate the accuracy of our channel probing and channel interpolation algorithm described in Sec. 4.3. We use an array of LED harvesters for probing and interpolation. The interpolated channel information is then used for beamforming towards the soft robot in the later sections.

4.6.1.1 Effect of Impairments

We evaluate the various sources of error in measurements from LED harvesters. These include: (1) Observation and quantization errors that can occur when we quantize the brightness of LED harvesters into real values; (2) Cameras may not respond linearly to the brightness of LED harvesters, and APIs providing the LED brightness can introduce quantization errors as well; (3) The positions of LED harvesters and soft robots may not be perfectly aligned leading to errors. We evaluate the effect of these impairments systematically through simulation.

Simulation Study: First, we use CST studio to simulate the EM wave inside the microwave oven. We then evaluate the effect of (1) positioning errors; (2) thermal noise; (3) quantization; and (4) camera non-linearity.

Fig. 4.9 (b) shows the results with a 4-antenna simulation. Our results show that: (1) *Pure Estimation:* The error of PhaseLift estimation drops from -1.5 dB to -4.5 dB from a single probe to a 16-probe experiment relative to the no impairment context. (2) *Thermal Noise:* We further note that the presence of

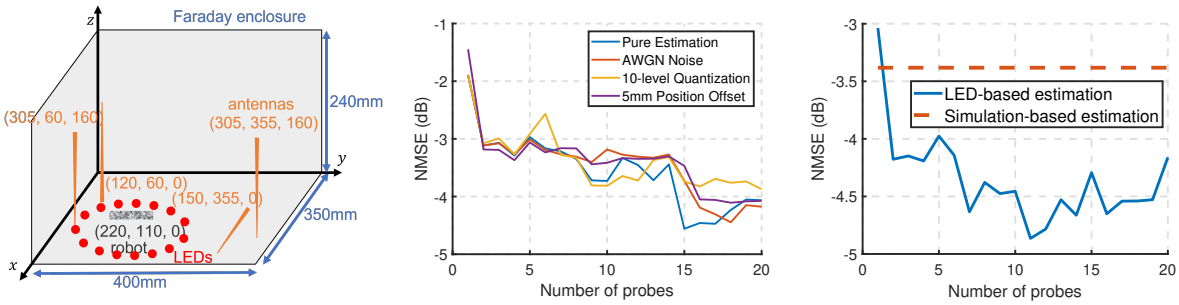

(a) Relative Positions (b) Error in Simulation (c) Error in Testbed

Figure 4.9: (a) Relative Positions: the antennas (orange) are placed with different orientations to provide diverse spatial diversity. The soft robot (silver) is placed at $(220, 110, 0)$ and LEDs (red) surround the soft robot with a spacing of 25 mm. (b & c) The Normal Mean Squared Error of WASER's channel probing. WASER's channel probing algorithm is resilient to various types of impairments, and the error converges to -4.5 dB after seven probes. Simulation-based estimation gives 1dB higher error than LED-based estimation.

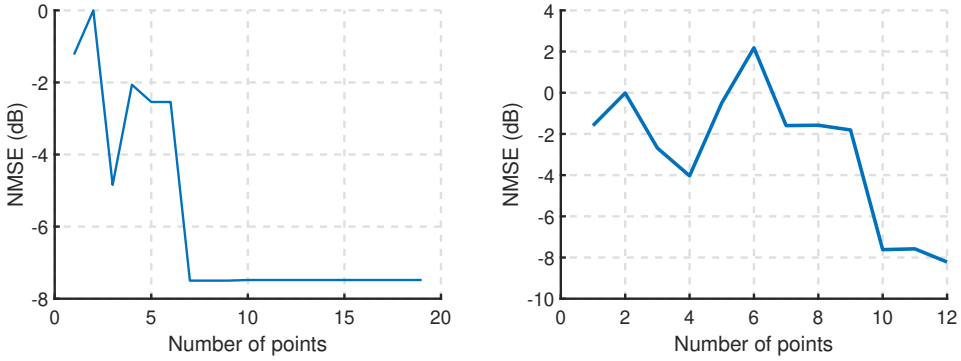

(a) Error in Simulation (b) Error in Testbed

Figure 4.10: Empirical NMSE of channel interpolation. (a) shows that the error converges to -7.7 dB for more than 7 points in simulation. (b) shows that the error converges to -7.9 dB for more than 10 points on our testbed.

additive white Gaussian noise (to around -30 dB as is typical in our experiments) does not significantly impact estimation. (3) *Quantization:* A 10-level quantization, where we uniformly quantize the values in $\mathbf{b}$ into 10 discrete values, causes the eventual error to degrade only slightly by about $0.3 - 0.7$ dB. (4) *Positioning:* We model a 5 mm error in positioning the LED harvester, which causes a slower convergence of the phase retrieval algorithm and results in a $0 - 0.9$ dB error increment to channel estimation. This is because a 5 mm shift in locations leads to 14.4 degrees of phase shifts.

Testbed Study: Next, we evaluate the accuracy of channel estimation using our testbed with 4-antennas and a single LED. We use 16 probes for channel estimation and reserve another 16 probes to evaluate the accuracy. As shown in Fig. 4.9 (b & c), our system converges to -4.5 dB error within 7 probes, and maintains around the error value with slight fluctuation caused by the impairments mentioned above.

In addition, we compare the WASER's channel estimation performance with a baseline that infers the channel without using LED harvesters. Specifically, we measure the channel by simulating the radio environment inside a Faraday shield. Results in Fig. 4.9 (c) show that this simulation-based channel estimation approach gives a 1 dB higher error. We note that this is because the radio environment is multipath-rich – we cannot infer the accurate wireless channel of the soft robot by only relying on the known location of the robot and RF antennas. Thus, WASER needs to deploy an array of LED harvesters to learn the multipath components in the environment. We further note that the WASER's LED-based channel estimation approach is generalizable for a more complex and dynamic Faraday shield (e.g., a flexible Faraday sleeve), given that using the simulation-based approach is time-consuming.

4.6.1.2 Interpolation Accuracy

We validate the interpolation accuracy both in simulation and on our testbed.

Simulation Study: We study the channel interpolation approach using CST Studio. We use p points for solving Eqn. 4.6 and infer the channels at an arbitrary point inside the closed surface bounded by p. Note that both in our simulations and experiments, we do not have the ground truth phase of $\mathbf{h}$. We therefore only present the estimation accuracy of the amplitude $|\mathbf{h}|$. Fig. 4.10 (a) shows the accuracy of interpolation with the changing number of reference points p, where the error reaches -7.7 dB for over 7 points of interpolation along a 2-D surface.

Testbed Study: We also evaluate the accuracy of interpolation in our testbed using a single antenna and 13 LEDs with a 50 mm spacing, among which 12 LEDs are used for providing interpolation data and another LED is used for evaluation. We use the result given by the simulation to initialize the optimization in the testbed. As shown in Fig. 4.10, the channel interpolation error in the testbed falls to -7.9 dB with 10 points. The difference might be caused by the setting difference in the simulation and in real testbed. In the testbed, the microwave oven has several bulges and gaps on the edges, causing extra load in the channel interpolation. Moreover, placement difference by several millimeters of the antenna in simulation and testbed can also cause a big difference in the optimized parameters.

4.6.2 Heating Efficiency on Various Textural Liquid Metal

Method: This section evaluates the heating efficiency on different textural designs of the thin liquid metal layer. The goal of designing a textural liquid metal layer instead of a uniform one is to combat edge effects (see Fig. 4.5 (b)) and improve the heating distribution over the soft robot. Specifically, WASER implements two patterns of the textural liquid metal layer: random texture and uniformly thin laminates.

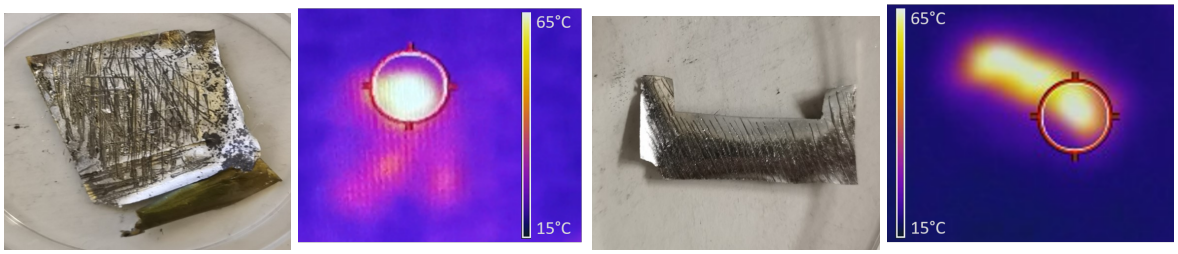

(a) Random Texture (b) Temperature Profile (c) Thin Laminates (d) Temperature Profile

Figure 4.11: WASER implements different textures on the conductive layer and evaluates their heat distributions. We show that the thin lamination design significantly mitigates the edge effects and achieves a uniform heat distribution.

Different textures are applied to the liquid metal during the material fabrication process. We deploy the two samples shown in Fig 4.11 (a) and (c) at the same location inside the heating stage. We then compare the heat distribution of these two textures when applying a single-antenna 2.4 GHz energy source with a 12.5 W transmitted power.

Results: Fig. 4.11 (b) and (d) show the heat distribution of the random texture and the thinly laminated texture respectively. The samples are heated at a room temperature ($\sim$ 20 °C). We observe the thinly laminated texture achieves a uniform heat distribution across the entire sample. The average temperature of entire sample increases to 75 °C within 3.5 seconds using a single-antenna 12.5 W wireless power. However, the heat distribution of the random textural liquid metal sample is random and one of the corners is overheated than other parts. We note that this is because the induced Eddy currents across the random textural liquid metal is uncontrollable. Therefore, we use the thinly laminated texture as our final design for the soft robot.

4.6.3 Heating Efficiency vs Number of Antennas

In this experiment, we evaluate WASER's heating efficiency as we vary the number of distributed antennas when the total transmitted power remains the same (2.5 Watts).

Method: We deploy up to four distributed antennas in our enclosed heating stage. The locations of the antennas is carefully calibrated using a Vector Network Analyzer so that these commercial antennas provide reasonable antenna gain ($\sim$ 3.0 dBi) at 2.45 GHz even if they are placed inside the enclosure. We consider various locations and orientations of the soft robot in both line-of-sight (LoS) and non-line-of-sight (NLoS) across experiments. We compare our algorithm with two baselines: (1) fixed phase *sweeping* as we described in Sec. 4.5; (2) a single-antenna setup that blasts the wireless energy blindly without any energy focusing. Across the experiments, we keep the total transmitted power the same which is 2.5 W

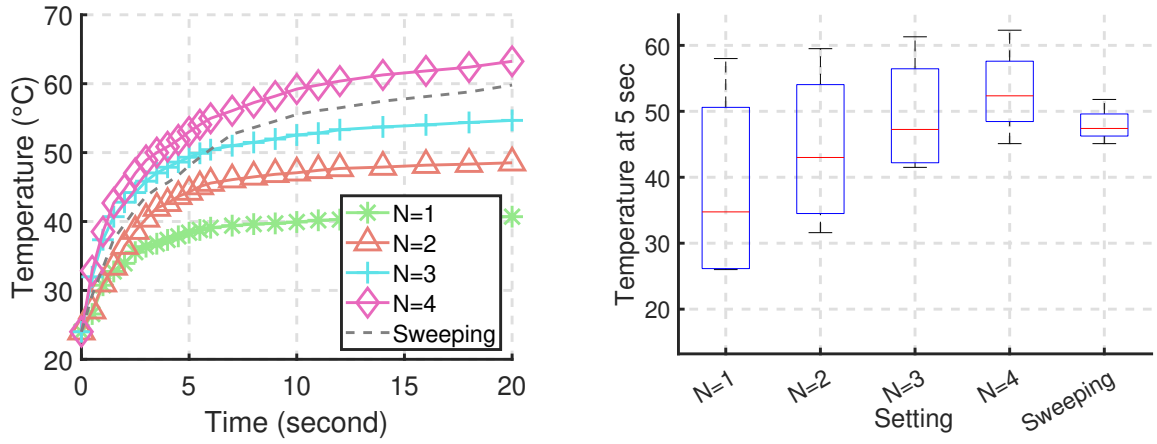

(a) Temperature vs. Time (b) Box Plot

Figure 4.12: Heating efficiency with different number of antennas N under total transmitted power 20W. We show that WASER's multi-antenna beamforming approach achieves a higher heating efficiency in the speed of heating soft robots as the number of antennas increases, and outperforms the baseline considering various locations of the soft robot in both LoS and NLoS.

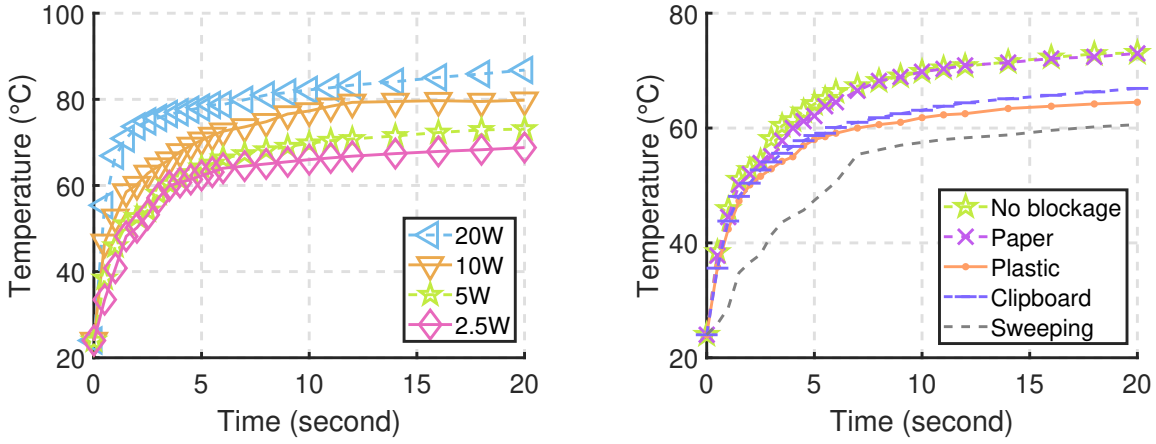

(a) Power (b) Blockage

Figure 4.13: (a) WASER achieves heating a soft robot to 60 °C at 1.5 seconds when applying a 20 W total transmitted power. (b) WASER is resilient to various materials of blockage.

or 34 dBm, and each antenna transmits the same wireless power. We note that our experiments consider distances between the soft robot and the closest antennas at least longer than 12 cm, but shorter than 22 cm which is limited by the size of the microwave cavity.

Results: Fig. 4.12 (a) shows the soft robot's average temperature across various locations with respect to increasing the time of actuation when using various number of antennas. We observe that for the single-antenna baseline, as expected, the temperature increases with the slowest speed compared to others and the temperature barely increases after 10 seconds as the heating effect and temperature dissipation reach a balance. WASER outperforms the single-antenna baseline by relying on its beamforming algorithm

which can precisely focus the wireless energy towards the soft robot. As WASER increases the number of antennas, it achieves a higher temperature at increasing speed. We note that WASER also outperforms the phase sweeping baseline that uses the same number of antennas. Specifically, WASER heats the robot to 60 °C taking 11 seconds on average considering various locations, but the phase sweeping baseline takes 20 seconds on average to reach 60 °C. This is because simply applying phase shifts to antennas without inferring the optimal beamforming vector cannot efficiently use wireless energy to heat the soft robot. But for all cases, the speed of temperature increase decreases with time until equilibrium.

Fig. 4.12 (b) depicts the bar plot of the soft robot's temperature at 5 seconds of heating. Again, we consider various locations of the soft robot in both LoS and NLoS. We find that the median of the temperature increases and the variance of the temperature decreases as WASER uses more distributed antennas. Overall, WASER's performance in heating efficiency increases by 50%, and 8% when compared to the single-antenna and random phase sweep baseline, respectively.

4.6.4 Temperature vs. Total Transmitted Power

Method: We evaluate the impact of total transmitted power on the temperature using the four-antenna setup. Specifically, we apply wireless energy with a total transmitted power of 2.5 W, 5 W, 10 W, and 20 W, and each antenna transmits the same wireless power. Here, we keep the soft robot at the same location across experiments.

Results: Fig. 4.13 (a) shows the temperature vs. the increase in time under different total transmitted power. We observe that as expected, a higher transmitted power significantly increases the temperature and the speed of temperature increase. WASER achieves heating the soft robot to 60 °C at 1.5 seconds when applying a 20 W total transmitted power from four antennas – a 27 times improvement compared to prior work [180]. We note that even though we apply a high transmitted power at the transmitter, the microwave leakage from our enclosed heating stage is well below the FCC regulation.

4.6.5 Impact of Blockage

Method: We evaluate the impact of blockage on the heating efficiency when the soft robot is in the non-line-of-sight from the beamforming platform. We infer the optimal beamforming vector to actuate the soft robot using our algorithm described in Algorithm. 3. We deploy different types of blockage such as paper, plastic and wood to emulate a non-line-of-sight environment. Across the experiments, we use our four-antenna setup with a 5 W (37 dBm) total transmitted power.

Results: Fig. 4.13 (b) shows the temperature of the soft robot being heated under various types of

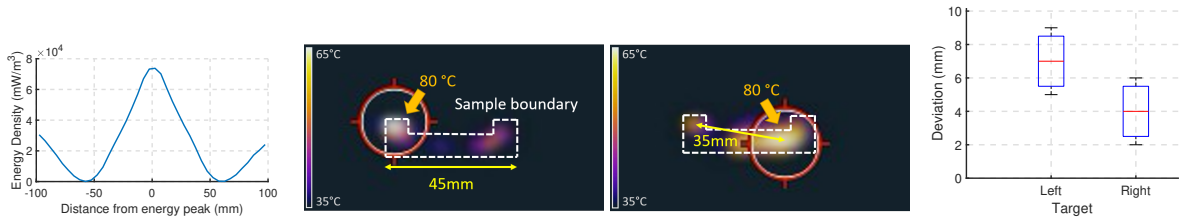

(a) Energy density at Y axis (b) Towards the left (c) Towards the right (d) Control accuracy

Figure 4.14: (a) The theoretical heating resolution that WASER can achieve. We define the heating resolution as the distance that the energy density attenuates 3 dB from the peak. (b)&(c) We evaluate the heating resolution of WASER by focusing energy towards the two ends of a 45 mm long soft robot. From the experiments, we show a 17.8 mm heating resolution on average. (d) We evaluate the robot movement accuracy. Specifically, we use WASER to apply beamformed energy towards the three individual parts (leftmost, rightmost) of the soft robot. We show that the most heated point deviates from the target point by an average of 7mm, 4mm, respectively.

blockage. As seen in the figure, the paper occlusion has a minimal impact in our system performance, which only increases the time to reach 60 °C by about 0.5 seconds compared to the situation where no blockage exists. Further, the performance under plastic and wooden blockage appears to be slightly lower which increases the time to reach 60 °C by 2 and 3 seconds respectively compared to the LoS setting. Overall, the WASER beamforming algorithm outperforms the random phase sweeping baseline, improving the time to reach 60 °C by at least 9 seconds, even if the soft robot is in non-line-of-sight.

4.6.6 Heating Resolution

Fig. 4.14 (a) shows the energy density at different parts in the Faraday shield along spatial axes. It can be seen that the energy density attenuates to half over a displacement of 20 mm. This demonstrates that our system is able to achieve a 20 mm heating resolution. We evaluate the heating resolution by beamforming at two ends of the sample shown in Fig. 4.5 (a). Fig. 4.14 shows the heat maps filmed by a thermal camera with the two beamforming weights. The sample is 45 mm long, with its boundaries depicted in the figure. In Fig. 4.14 (b), only the left end is heated to around 80 degrees, while the rest remains below 40 degrees. In Fig. 4.14 (c), the channel gradient at the right end of the sample is low, so the right end and its surrounding parts are heated to 80 degrees, while the left part is still under 40 degrees.

4.6.7 Accuracy of Actuation

Methods: The ability to accurately move the soft robot to exhibit desired shape changes is crucial. We evaluate the accuracy in making desired shape changes of WASER's soft robot by applying heat to a desired location. Specifically, we apply two different beamforming vectors to beam maximized energy at the leftmost and rightmost part of the soft robot respectively. By beamforming to these two parts,

the soft robot is expected to exhibit two different shape changes – the leftmost part of the soft robot flipping downwards and the rightmost part flipping downwards about a pre-defined point of maximal heat – the hotspot maximum. In the experiments, we measure the ground truth shape changes of the soft robot by using a camera. We then measure the deviation (Euclidean distance) between actual and desired maximum hotspot locations.

Results: Fig. 4.14 (d) shows the accuracy in actuation of a WASER's soft robot. We demonstrate that the soft robot exhibits an average of 7 mm and 4 mm deviation from the desired hotspot locations for the leftmost and rightmost movements respectively. While the deviations are low, we note that they can be further improved through more accurate channel estimation and reducing fabrication imperfections of the soft robot.

4.6.8 Shape Morphing and Motion Capabilities

We present two key proof-of-concept actuation examples that apply to potential soft robotic use-cases. These actuation examples: bending and folding, are both enabled by the programmable properties of LCE. These samples are cross-linked using UV light. By polarizing this light along one direction on the top side and in the perpendicular direction on the backside, the orient-able monomer mesogen backbone of the polymer chains are crosslinked in a twisted nematic pattern [175]. When heated above their transition temperature, these mesgons and their polymer chains contract along the direction of orientation inducing a bending. By placing the samples oriented face up or face down, bending and folding can be induced. Further information on the fabrication process of these samples is given in Section 6.

Bending: Bending can be induced by actuating a soft robot as described above. The robot is designed to flip downwards when particular parts are actuated. The robot is in a flat initial state, placed on a wooden surface for proper friction force, as shown in (a). We then beamform towards the left side of the robot, heating it, and induce a force downwards for the left part, as shown in (b). In (c), the left side is under medium actuation, and in (d), the left side is fully actuated, inducing a complete bend downwards to support its own weight.

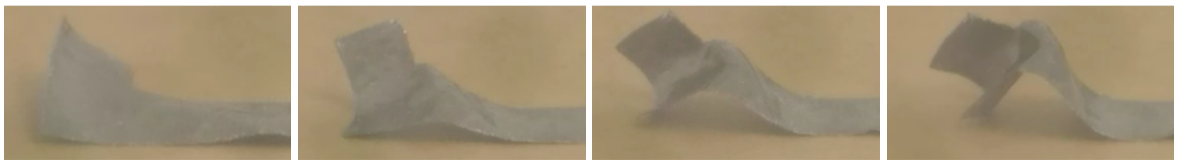

(a) Initial State (b) Start Actuation (c) Medium State (d) Full Actuation

Figure 4.15: Snapshots of WASER Bending.

(a) Initial State (b) Start Actuation (c) Medium State (d) Full Actuation

Figure 4.16: Snapshots of WASER Folding.

Folding: Next, we demonstrate a folding motion shown in Fig. 4.16. The initial state is flat on a low-friction surface, as can be seen in (a). After beamforming for a short time, in (b), the right part of the soft robot starts to actuate. (c) displays the medium state when the robot begins to fold, and (d) demonstrates the status when the soft robot is completely actuated, when the right part completely folds towards the center and clings to the left part of the robot.

4.7 Discussion and Limitations

To Eliminate Faraday Shields: WASER's use of a Faraday shield can be eliminated by reducing the total transmit power. We've demonstrated that WASER is able to actuate the soft robot with a very low transmit power (e.g., 2.5 W) which is the typical transmit power of a long-range Wi-Fi repeater. We believe a reduced transmit power (e.g., 1 Watt) without sacrificing the capability of actuating wireless electronics-free soft robots can be achieved by future improvements from multiple perspectives: (1) *Materials:* creating LCE materials be capable of performing shape changes at a lower temperature; using a higher conductive material such as gold to induce a stronger eddy current throughout the robot. (2) *Higher Frequency:* Moving to a much higher frequency such as 60 GHz millimeter wave can significantly reduce the side lobe energy leakage and enable better energy focus because of a much shorter wavelength ($\sim$ 5 mm) than 2.4 GHz — using a much less transmit power but beaming more focused RF energy to power the soft robot. We believe WASER is a starting point for a rich space for future exploration.

More Complex Movements: In this paper, we present relatively simple shape changes and movements of parts of the soft robot. However, a series of more complex moves can be performed based on the simple moves. Specifically, WASER applies a series of beamforming vectors to induce the continuous movements. For example with alternating bending and folding, a soft robot may be able to perform crawl-like movements. In addition, motion in reverse require the robot cooling down. This may take time for our current implementation, but researchers in soft robots propose innovations in reducing the cooling time [226]. More complex movements can be further improved by augmenting the shape morphing

capability of the soft robot itself [227, 176]. Doing so does require more advanced mechanical design of the soft robot, but we believe that the WASER platform is a natural complement to such research in soft robotics.

The Location of the Soft Robot: WASER assumes that the location of the soft robot is known, although its shape changes under heating can be inferred as shown in Sec. 4.6.8. Potential methods to localize and sense the actual shape change of the electronics-free soft robot could leverage wireless localization techniques based on acoustics or mmWave Radar – we leave this for future exploration.

Potential Use-cases: The actuation platform of WASER itself can serve as an important building block to enable future research in soft robotics, material science, and wireless systems. The electronics-free soft robot wirelessly actuated by WASERś actuation has wide-ranging potential use-cases. Here, we list some examples that could be built and evaluated in follow-up work. For example, the soft robot with an inner cavity can be remotely controlled to change their state of shape deformation and then perform on-demand release of chemicals (e.g., disinfectants) stored inside their cavity. In addition, with a more advanced design of the robot, WASER could remotely navigate the robot in confined spaces to carry and deliver payloads. Further, a swarm of soft robots can be deployed as a shape-changing surface whose roughness can be remotely controlled by 's actuation platform. We leave the evaluation of these use-cases to future work.

4.8 Summary

This chapter presents WASER, a solution that uses highly-controlled microwave radiation to control the movement and shape of elastomer-based soft robotic platforms. WASER achieves this through innovations in mobile systems and material science. On the mobile systems front, WASER develops a blind beamforming solution that rapidly focuses energy at desired points on the robot to heat and induce shape changes, without active electronics on the robot. On the materials front, we develop novel material designs that improve the heating efficiency of the platform without inducing high bulk and rigidity. WASER is fully implemented and evaluated on soft robotic platforms. We believe there is rich scope for future work that builds on WASER, particularly in exploring diverse use-cases and applications of wirelessly actuated soft robotics for industrial, aerospace, and medical applications.

Chapter 5

Speech Sensing Using Wirelessly Powered Tattoos

The dissertation so far describes the benefits of precise control of wireless energy delivery, such as powering battery-free tags at scale (Chapter 2), realizing fine-grained food heating (Chapter 3), and moving electronics-free soft robots (Chapter 4). The rest of the dissertation focuses on the opposite of the system—the energy receiving side—creating systems that can efficiently harvest wireless energy and be capable of advanced sensing capabilities. Specifically, this chapter presents RFTattoo [4] that can be powered wirelessly and capable of speech recognition.

In the US, more than 2 million people require digital Adaptive Alternative Communication (AAC) methods to help compensate for speech impairments [228]. While various classes of voice impairments exist, we target acquired conditions where users continue to be able to make facial and lip gestures associated with speech. We aim to learn these gestures over time to produce speech in real-time. Our approach applies to a wide range of temporary and permanent acquired dysphonia (voice disorders) ranging from hoarseness to complete loss of voice that occurs in about 1% of the global population [229].

While there is rich literature on assistive input-to-speech technologies for speech impairments, state-of-the-art solutions suffer important limitations. Camera-based [62, 63, 64, 65] visual solutions for real-time lip-reading require users to constantly be within line-of-sight of a camera, which may not be possible when the user is on the move. Audio-based assistive solutions [230] only apply to speech impairments where users are able to produce sounds and struggle in noisy environments. Past work has proposed a variety of face-worn sensors for speech sensing, particularly in clinical settings, such as magnets attached to the tongue [231], EEG helmets [232] and EMG electrodes on the face [233]. Assistive text-to-speech innovations require users to provide constant manual input to the system via keypads [234] or various

user interfaces [235] that require training and practice for proficiency. There remains a gap for an everyday intuitive assistive technology for voice impairments that does not require external infrastructure, bulky sensors on the face or manual hand input.

We present RFTattoo, the first wireless speech recognition platform for voice impairments through skin-friendly, wafer-thin, battery-free and stretchable RFID tattoos. We fabricate specialized light RFID tattoos attached to the skin surface of face at known locations. Each tag is fabricated to be stretchable, flexible, wafer-thin, extremely light and made with hypoallergenic materials. The tags are designed to be hidden under makeup and extremely skin-friendly. We track the strain of individual tags over time as they deform in response to motions generated by different intended sounds. However, it is often the case that certain distinct sounds produce similar facial movements. To this end, we build natural language processing models that combine identified facial gestures in context to construct meaningful words and sentences. A detailed user study with 10 users reveals 86% accuracy in recognizing the top-100 words in the English language.

RFTattoo's first challenge is to process signals from RFID tattoos to recognize distinct facial and lip gestures called *visemes*[1][236], that correspond to sounds the user intends to express. RFTattoo recognizes visemes by modeling the pure stretch of the flexible tag antenna. An intuitive approach to model tag stretch to infer its impact on the frequency at which it resonates. Specifically even a small change, say one millimeter, in the electrical length of an antenna lowers its resonant frequency by as much as 8 MHz in our experiments. Unfortunately, RFID tags in the U.S. operate in the FCC's unlicensed 900 MHz band with an effective bandwidth of 26 MHz. This makes it challenging to accurately capture the large frequency shifts induced by stretch. More importantly, requiring an RFID reader to hop through all frequencies even within the unlicensed band would be too time-consuming ($\sim$ few seconds) to recognize real-time speech.

RFTattoo addresses this challenge by probing multiple specially tuned RFID tags instead of probing multiple frequencies at the reader. In particular, we design an RFID tag that advertises the bits of its own current stretch value even if it is probed at one frequency (e.g. 915 MHz). Our approach to do so attaches multiple RFID chips to a common antenna, each tuned to multiple sets of specially chosen frequencies. We design the i^{th} chip in the RFID tag to respond at 915 MHz only if the i^{th} bit of the current stretch of the tag (expressed in millimeters) is one. In effect, this allows a tag to recover the bits of its own extent of stretch with a single frequency probe, with a small number of RFID chips per tag – logarithmic in the desired stretch resolution. We formulate a novel super-resolution optimization algorithm that improves this resolution even further by processing the power of the received signals across chips.

A second challenge RFTattoo must address is the dynamic radio environment – changing orientation

[1] A viseme is a set of phonemes that look the same, for example when lip reading.

of the RFID tags, multipath reflections as well as movement of the user's body. RFTattoo achieves this through a novel tag antenna design that isolates the impact of stretch from other aspects pertaining to the radio environment. Specifically we fabricate two co-located RFID antennas with two materials – one stretchable and one non-stretchable. We then compare the signals received across both RFID tags to isolate any effect from the tag location, orientation and radio environment.

Finally, RFTattoo builds a natural language processing framework to map stretch values of tattoos placed at different points in the face to recognize words and sentences the user intends to speak. A key challenge in this regard is the fact that some sounds produce identical facial and lip gestures (visemes) and therefore cause a high degree of ambiguity in the recognized phonemes. RFTattoo addresses this through two approaches. First, RFTattoo monitors subtle movements of the user's tongue through its impact on the magnitude and phase of the RFID tags on the skin's surface. We show how this allows for disambiguation of certain phonemes that produce identical facial movement. Second, RFTattoo leverages a useful property commonly exploited in natural language processing – the fact that adjacent phonemes are not completely independent but must follow the English dictionary and rules of grammar.

Limitations: We emphasize a few important limitations of RFTattoo: (1) RFTattoo achieves highest accuracy when the location of RFID tags on the face are known through a light-weight calibration a priori. This means that for optimal performance, one must re-calibrate should RFTattoo tags be peeled off and on, or with natural wear. (2) RFTattoo may miss visemes should specific tags be unresponsive owing to shadowing from the body relative to the reader. (3) RFTattoo's accuracy is poor in the face of unknown or untrained words (e.g. less common words and proper nouns). This is a common problem shared by voice recognition systems [237] (e.g. Siri, Alexa, etc.) as well as visual lip reading systems [62, 64, 65].

We implement RFTattoo by building custom tag antennas using stretchable Ag-PDMS conductors on PDMS substrates connected to three RFID chips. We use a meander-line antenna appropriately impedance tuned to respond at the 900 MHz ISM band. We use commodity Impinj RFID readers attached to the user's waist.[2] Our system is attached to the user's face using hypoallergenic stickers and covered with makeup. We conduct a detailed user study with 10 users including two users with temporary dysphonia (loss of voice). We also include results when all users are instructed to mouth words silently. Our results reveal that:

- RFTattoo achieves a median accuracy in stretch of 1.4 mm.
- RFTattoo distinguishes between eleven visemes of the English language [238] at an accuracy of 90%.

[2]Note while our implementation uses a relatively bulky 4-antenna RFID reader due to availability of channel state information, our system relies only on information from only one antenna. Our system will be compatible with much more portable and highly compact readers as future commercial products become more open.

- RFTattoo recognizes the most frequently used 100 words of the English language at an accuracy of 86%.

Contributions: Our main contribution is a novel system that recognizes intended speech of users with voice impairments using light-weight RFID tattoos attached to the face. Our contributions include:

- Algorithms that recognize subtle mm-accurate stretches of the tattoos as well as movement of the tongue by processing RF-backscatter signals at a handheld reader.
- A natural language processing framework that recognizes various facial gestures associated with speech to construct meaningful words and sentences.
- A detailed user study that reveals the promise of our approach in recognizing intended speech, even when users do not make any sounds.

5.1 Related Work

5.1.1 RFID-based on-body Sensing

Sensor-equipped RFID tags have been used to monitor temperature [239, 240], moisture [241], or even neural signals [242]. More recent work relies solely on the phase and RSSI of tags to accurately track their location [25, 26], including sensing the body skeleton [2], shape [1] and target imaging [243]. RFID tags have also been used for finger touch tracking by sensing the impedance mismatch between the tag chip and the antenna [244]. Recent work also utilizes coupling effect of the near-field antennas to distinguish different materials underlying RFID tags [245]. In contrast to these systems, RFTattoo seeks to infer the stretch of individual RFID tags at known locations on the user's face. It does so purely using signals received at a handheld RFID reader. We note that our system is designed to detect extremely subtle stretches (few mm) that are within the dimensions of an RFID tag from a handheld commodity single-antenna RFID reader. This is beyond the purview of state-of-the-art RFID location-tracking solutions and therefore necessitates new solutions.

5.1.2 Sensing Using Stretchable Electronics

There have been many recent advances in the use of soft materials for mechanically robust robots and electronics: crawling robots powered with compressed air [191], electrically-powered soft robots capable of locomotion through confined spaces [246], and self-healing robots that can continue to walk when their on-board circuitry is severed [247]. Advances in stretchable electronics have enabled soft circuits with

multiple sensing modalities [248, 249, 250, 195], but these typically rely on battery-powered sensors or an external power supply. Stretchable RFIDs has been explored for strain monitoring [251, 252] and gesture detection [253, 254] but cannot be applied solely in the 900 MHz ISM band alone (per FCC limits) and are primarily based on RSSI and therefore vulnerable to multipath in the on-body context. In the HCI community, skin-friendly tattoos or epidermal electronics [255, 256, 257, 258, 259, 260] have been used to provide on-skin user interfaces, which either operate only in the near-field or require active components on the skin. In this chapter, we focus on building passive, stretchable and flexible RFID tattoo that can sense speech with commercial RFID readers.

5.1.3 Automatic Speech Recognition

Recent research in automatic speech recognition for persons with voice impairments has explored visual and audio-based strategies [61]. Video-based systems such as lip reading by camera [64, 65] assume that the user is within the field-of-view of a camera. Audio-based solutions assume the user is able to make certain audible sounds and therefore do not apply to users with complete dysphonia (loss of voice). More recent work requires the use of specialized powered sensors to detect speech, for example: (1) detecting EMG signals when speaking by placing skin surface electrodes at the face [261, 262, 263]; (2) detecting tongue movements by tracking small magnets pasted on the tongue [231]; (3) detecting imagined speech by using EEG [66] or nerve signals [67]; (4) detecting air flow as a proxy for speech [264]; (5) detecting mouth movement of radio reflection using Wi-Fi [265]. These methods can work without any sound production, but either require intrusive equipment or are limited certain phonemes or words, precluding wide-scale deployment [266]. RFTattoo complements these approaches by providing a solution that detects a wide range of phonemes with a battery-free and light-weight RFID-based solution for speech recognition.

5.2 Overview of RFTattoo

RFTattoo's primary goal is to infer the speech of the user in real-time based on signals reflected off RFID tags. We assume the RFID tags are attached as light-weight tattoos on the skin at known locations on the face. We measure the signals reflected off these tags from a commodity portable RFID reader. We allow the reader to be portable and not in the line-of-sight of the tags. We measure the magnitude and phase of the signals reflected off multiple tags from the reader. We then process these signals to measure two physical properties: the stretch of individual RFID tags, and the position of the tongue. We monitor both these properties over time to classify between various visemes (facial gestures) the user makes. Finally, we process these visemes to infer the words and sentences spoken by the user. The rest of this chapter

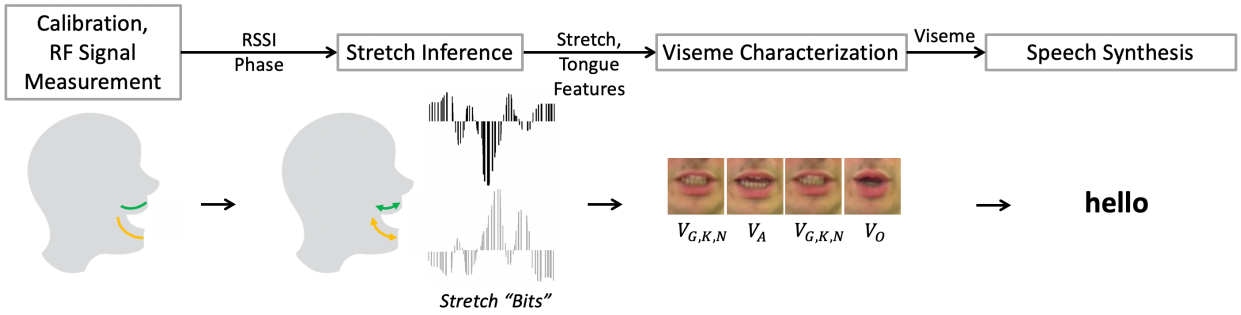

Figure 5.1: RFTattoo's Architecture: (1) measures the wireless channel of RFID tattoos; (2) infers stretch bits and tongue position based on the reflected signal power and phase of multiple three-chip RFID tattoos; (3) feeds the features into machine learning models to recognize corresponding facial gestures (viseme – see Table.5.1); images show the corresponding viseme from the GRID dataset [9]; (4) combines visemes to form meaningful words and sentences by natural language processing.

describes three key components of RFTattoo (Fig. 5.1).

(1) Inferring RFID Stretch and Tongue Position: RFTattoo actively measures the stretch of each RFID tag and the position of the tongue – two key aspects that help recognize speech. RFTattoo infers stretch by monitoring its effect on impedance due to the fact that a stretched tag is longer and thinner. RFTattoo specifically measures the frequency response, i.e. the change in magnitude of the reflected signal across frequencies to study this effect. We measure this property accurately and in real-time despite the limited bandwidth of commercial reader. We further show how RFTattoo can also infer the position of the tongue through its effect on RFID impedance. Sec. 5.3 details our solution.

(2) Designing the RFID tags: Next, we show how we can design RFID tags to specifically ease identification of stretch. We optimize both the material of the tag as well as the design of the antennas. First, the material of the tags must be optimized for facial skin – designing a material that is too thick will make it inconvenient to wear and designing one too thin will make it vulnerable to wear after repeated stretches. Our antenna design must also ensure that it resonates optimally with the RFID reader despite changing impedance when in contact to the skin and the limited area available on the skin. Sec. 5.4 describes our approach.

(3) Processing Speech: Finally, given the stretch of individual RFID tags, RFTattoo fuses these measurements to infer visemes, that are visual gestures of the face produced by different syllables pronounced by the user. We note that some visemes can be produced by multiple sounds, (e.g. "thee" and "tea" are indistinguishable visually). We show how we can disambiguate many such sounds using the position of the tongue. Sec. 5.5 describes our system that borrows from natural processing techniques to fuse the resulting phoneme measurements into meaningful words and sentences.

5.3 Processing RFTattoo Signals

In this section, we characterize two properties from signals reflected off RFID tags on the user's face: the stretch of RFID tags and the position of the tongue.

5.3.1 Inferring Tag Stretch

Our key approach to monitor tag stretch measures the change in impedance as a result of the tattoo elongating. Specifically, as tattoos are stretched, its effecting width decreases and length increases, both of which increases its resistance and reactance. In effect, this causes a change in the resonant frequency of the RFID tag. For ease of exposition, we first assume the absence of multipath and a fixed relative orientation of RFID tag to the RFID reader. We will explicitly deal with these challenges later in the section.

Why does resonant frequency shift with stretch?: The resonant frequency of an antenna is the frequency, where the amplitude is higher than at adjacent frequencies. Stretching an RFID tag changes its antenna's electrical length and therefore its resonant frequency. Specifically, as the antenna length increases, the wavelength at which it resonates also increases meaning that the resonant frequency will shift towards lower frequencies. Mathematically, the resonant frequency of a half-wave dipole antenna is written as [267]:

$$f_{res} = \frac{c}{2\sqrt{\varepsilon_e}} \frac{1}{L + L_e} \tag{5.1}$$

where L is the effective length of the half-wave dipole antenna and the L_e is the effective elongation of the antenna as stretching, ε_e is the effective relative permittivity of the antenna substrate and c is the speed of light in free-space. The ε_e can be estimated using the method mentioned from [268]. From this equation, we can see the resonant frequency is inversely proportional to its electric length leading to a very simple approach to infer stretch once electric length is found accurately.

Fig. 5.2 plots the frequency response of a half-wavelength dipole antenna when stretched to different lengths measured by the VNA and simulated by HFSS. Indeed it is surprising that a stretch of just 1 mm leads to a substantial resonant frequency shift of 8 MHz. On a positive note, this shows that resonant frequency allows for obtaining very fine grained values of stretch. However, on the flip side, detecting even an effecting stretch of a few millimeters will require an order-of-magnitude greater bandwidth when compared to the 26 MHz of bandwidth available in the 900 MHz ISM band. This motivates a challenging problem – how do identify resonant frequencies short of sweeping a wide range of frequencies.

How to find the resonant frequency?: A strawman approach to find the resonant frequency is to interpolate this by measuring RSSI across the available bandwidth at the 900 MHz ISM band (about 26 MHz).

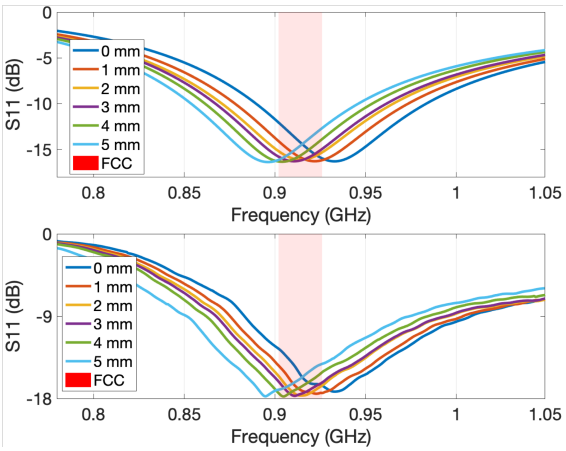

Figure 5.2: Simulated data from HFSS (upper) and real data collected by VNA (down) of the resonant frequency shift with stretch steps of 1mm. Red patch shows the ISM bandwidth from 902 to 928 MHz. An average of 8 MHz frequency shift is seen. After just 3.25 mm of stretch, the resonant frequency goes out of the ISM band. Even within the band, the gain drops up to 6 dB.

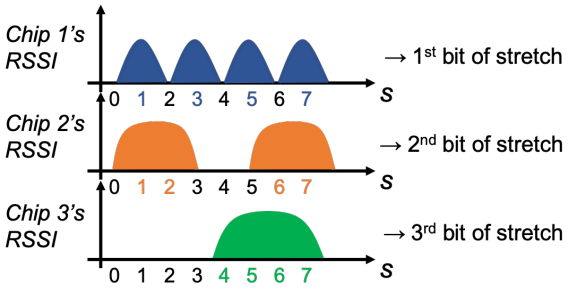

Figure 5.3: The i^{th} chip on an RFTattoo tag responds with a maximum signal power when the i^{th} bit of s is one, where s is tag stretch. For example, Chip 3 responds when the 3^{rd} bit of stretch is on which corresponds to the possible stretch of 4, 5, 6, 7 in the figure.

This constitutes a total of 50 hopping frequencies for Gen2 RFID protocol across 902 to 928 MHz (assuming the reactance is always conjugated with the RFID chip) accommodating a mere 3.25 mm of effective stretch in terms of resonant frequencies. Outside this range of stretch, we note that one can interpolate the RSSI drop across this small range of frequencies to infer the resonant frequency outside the FCC band. However, this approach remains vulnerable to multipath and noise. Further, it is time-consuming taking at least 10 seconds to sweep the 50 available frequencies with state-of-the-art readers [269].

RFTattoo finds the resonant frequency by relying on multiple co-located tags instead of multiple widely separated frequencies to obtain stretch. Specifically, we designing customized tags that all resonate at the 900 ISM band, but only when pulled to specific values of stretch. This means that by learning the distribution of signal power from the various tags, one can learn the precise distribution of signal power across frequencies.

Mathematically, let $t_1, \ldots, t_n$ be designed to resonate at a maximum power when stretched to values $s_1, \ldots, s_n$. Let $P(t_i)$ denote the power of the signal received from tag t_i. Then it is easy to see that the expected value of stretch s can be directly interpolated from the power of each tag as:

$$s = \sum_i s_i \frac{P(t_i)}{\sum_j P(t_j)}$$

The above approach has a unique advantage: it allows for finding the optimal value of stretch instantaneously without any frequency hopping beyond the FCC bands. However, it also has a notable disadvantage: it requires multiple co-located tags which can add to the bulk of the system. Indeed, using

a larger number of tags can ensure more accurate system performance, given that it allows resonant frequencies to be more finely sampled. Indeed, it appears that the number of tags required is linear in the number of discrete resonant frequencies (i.e. stretch values) that need to be sampled. This leads us to a fundamental question: Can we design a system that requires a sub-linear number of tags $k << O(n)$ to sample n discrete resonant frequencies?

RFTattoo develops a solution that requires $k \sim O(\log n)$ tags to sample n discrete resonant frequencies. RFTattoo formulates the question of finding the resonant frequency through a divide-and-conquer approach. Specifically, let us suppose the optimal resonant frequency f can be represented in k bits ($k = \lceil \log n \rceil$). Then we design k tags each designed to always resonate on a set of multiple frequencies. In particular, the j^{th} tag is designed to resonate at all frequencies where the jth bit of the resonant frequency index f is one. In effect, this ensures that by simply reading off the power of the $k = \lceil \log n \rceil$ tags, one can infer the true stretch of all tags (see Fig. 5.3 for an example). Mathematically, we can write the index of the resonant frequency among n possible values as:

$$f = \sum_{i=1}^{k} 2^{i-1} x_i \tag{5.2}$$

where x_i is one if and only if $P(t_i)$ is above a threshold and zero otherwise.

How do we map resonant frequency to stretch?: In practice, resonant frequencies do not lie on discrete values but vary continuously over a large range of a few hundred megahertz. RFTattoo captures this by noting that a tag that resonates at a particular discrete frequency will also strongly reflect neighboring frequencies. Now let use $P(t_i, f)$ denote the power of the signal received from tag t_i at f frequency, notice that f is within the range of 902 to 928 MHz following the Gen2 RFID protocol. This means that we can effectively model the expected power $EP(t_i, f|s)$ by manually measuring the power of each tag $P(t_i, f)$ at specific frequencies should it be stretched to a value s. We model this expected power by simulating the frequency spectrum for each chip across a range of stretch values using HFSS [270] that draws from well-known antenna stretch models based on Poisson's effect [271] where Poisson's ratio is 0.5. For example, as depicted in Fig. 5.2, the orange curve shows the expected power for the tag with 1 mm stretch value under the FCC frequency spectrum. We observe that the simulated expected power across stretch values closely aligns with empirical measurements from our prototype (in Sec. 5.7). We compute the expected stretch as:

$$\tilde{s} = \min_{s} \prod_{i} |EP(t_i, f|s) - P(t_i, f)| \tag{5.3}$$

We then solve this optimization problem using standard sequential least square programming with a zero stretch of the tattoo as the starting point. We also bound the bandwidth and stretch within the 26

MHz and 50 mm. To achieve the above design we need to fabricate tags that have the ability to resonate at multiple specific frequencies. In Sec. 5.4, we show how RFTattoo achieves this by adding multiple parallel impedances all coupling to the same chip.

Achieving Super-Resolution: We note that the above optimization problem deliberately models signal power to achieve high stretch resolution. Specifically, by simply treating whether a chip resonates or not at a specific stretch value as a binary 0 or 1, one can only resolve stretch at up to 2^k discrete values where k is the number of chips available on the tag. In contrast, as we model the continuous expected power quantity per chip with much more fined-grained stretch values instead of 2^k stretch, one can achieve significantly higher resolution than 2^k. Our results in Sec. 5.7 show median stretch accuracy of about a millimeter with only three chips per tag.

Impact of Tag Location, Orientation and Multipath: Our discussion thus far does not account for the unknown location of the RFID reader as well as multipath reflections from various surfaces between the reader and tag, including the user's body. In particular, the received signal power from an RFID tag depends upon four properties: (1) The location of the reader relative to the tag; (2) Signal multipath, including any attenuation and reflections of the user's body; (3) The orientation of the tag; and (4) The stretch of the tag. Our goal is to therefore isolate any received power change due to stretch from all remaining properties.

Our approach to do so attaches an additional RFID chip to the tattoo that is co-located but has a built-in antenna that does not undergo stretch. Given that this antenna is at the same location as the stretchable tag, it shares any attenuation across frequency owing to location, orientation and the effect of multipath. Indeed, any change in signal power across frequencies between this tag and the stretchable tag can be attributed purely to the stretch of the latter. Sec. 5.4 describes how we can use different antenna materials to create both stretchable and non-stretchable co-located tags.

5.3.2 Finding Tongue Position

RFTattoo uses the position of the tongue to disambiguate certain extremely similar sounds that produce virtually identical facial expressions, and therefore identical stretch values.

Finding tongue position: At a high level, our approach to find tongue position relies on the fact that the tongue changes the near-field radio environment of the tags that imposes additional impedance to the tags antenna. When the tongue moves from upper jaw front to lower jaw front, even though the tags shares the same surface of skin, the interior structure of the face changes so that the underlying material impedance changes. This would cause additionally resonant frequency shift apart from the stretch of

the tags. We note that since tongue position is effectively a component of multipath, we only use signal measurements from the stretchable tag and do not use measurements from the non-stretchable tag. We rely on machine learning models to decouple the effect of tongue position from other sources of signal multipath. Specifically, we classify tongue positions using Random Forest at normal resting position, upper jaw front, upper jaw back lower jaw front and lower jaw back, while keeping the same all other sources of facial movement (see Fig. 5.10(a)). In Sec. 5.7.2, we evaluate system performance with two other candidate classifiers.

Finding the collective spatial arrangements of the tags: Besides the position of the tongue itself, RF-Tattoo also processes information from the spatial arrangement of various tags on the face to gather information on facial expression. We rely on the rich literature on RFID tag localization to do so. Our work specifically relies on WiSh [1] which allows a single antenna RFID reader to track the spatial shape formed by multiple RFID tags as they vary in geometry. We note that while the spatial arrangement of tags are valuable, the value of stretch plays a much more crucial role in determining facial expressions during speech. This is because movements of the face during speech are subtle, meaning that one has to compute minute changes within the dimensions of a tag itself to accurately characterize them.

5.4 Tattoo Design and Fabrication

This section describes our RFID tattoo design and fabrication approach, specifically: (1) The design of the antenna pattern to fit within the dimensions of the user's face while providing high signal variations under stretch; (2) Materials and fabrication methods that maximize durability of the tags while keeping them light and thin.

5.4.1 Tattoo Antenna Design

Our objective is to design antennas that resonate at multiple frequencies and respond over well-defined bandwidths when stretched to a specific value. Doing so requires carefully designing the impedance of the RFID circuit. RFTattoo designs three components of the tag to achieve this: (1) the dipole, which is the physical far-field component that enables the communication between the tag and the RFID reader; (2) the inductor loop matches the impedance of the RFID chip; and (3) the coupling section, which transfers power between the dipole antenna and the inductor loop. We describe how each of these components are designed below:

Designing the dipole to be skin-compatible: RFTattoo uses a center-fed half-wavelength dipole antenna. The electrical length of a half-wave dipole is half-wavelength of the operating frequency of the antenna.

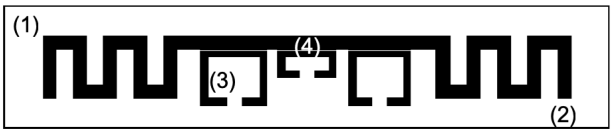

Figure 5.4: RFTattoo's Antenna Design: (1) Stretchable substrate. (2) Half-wavelength Dipole. (3) Impedance matching networks with a sets of high impedance modules. (4) Coupled spacing.

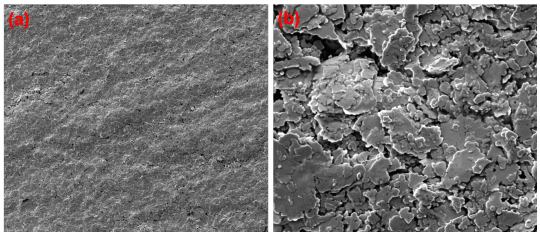

Figure 5.5: Under stretched RFTattoo sampled from scanning electron microscope. (a) 300× magnified. (b) 5000× magnified. Small cracks appear between the silver flakes embedded withing the elastomer composite when the tag has been stretched.

However, at the 900 MHz ISM band, the physical length of a half-wave dipole is approximately 16 cm, which is too long for a face tattoo. To reduce the size of the antenna, we adapt the technique of meandering design (see Fig. 5.4) that maintains a reduced size with a long reading range [251]. We tune the number of meandering bends of the antenna in a manner that minimizes antenna loss (i.e. increase trace width), subject to length constraints of the antenna to resonate at the appropriate frequency.

As we design the dipole antenna, we need to consider the impact on the resonant frequency of the material to which the tag is attached – the human skin, in our case. Skin has heavy dielectric properties which motivates the need to appropriately tune the number of meander bends [272]. In our antenna design process, we build an ANSYS simulation that models the impact of human skin through a multiple substrate layer – containing one layer of antenna substrate and one layer of a human model provided by ANSYS HFSS. Our simulation is designed to optimize the performance of the dipole to resonate at 915 MHz. Having said this, in practical experiments, we observe minor variations in the resonant frequency between different users due to differences in the electric properties of their skin. Fortunately, the impact of this shift remains consistent across stretch of an RFID tag and can therefore be calibrated for a priori when the user makes a neutral expression. In Sec. 5.7, we demonstrate how our approach generalizes to diverse types of human skin.

Impedance Matching to resonate at many frequencies: An important challenge we face in antenna design is to make it respond to multiple distinct frequencies. One option is to pack together multiple antennas of different lengths co-located in the vicinity of an RFID tag. Doing so would clearly be too bulky to attach as a wafer-thin tattoo. In contrast, RFTattoo tunes a single dipole antenna to multiple resonant frequencies through a multiple impedance matching units.

To elaborate, recall that to maximize the power transfer efficiency between the chip and the antenna at a given resonant frequency, any impedance mismatch should be resolved between these two components.

A tag's reflection coefficient (Γ) which accounts for the impedance mismatch between the chip and the antenna is given by $\Gamma = \frac{Z_C - Z_A}{Z_C + Z_A}$, where Z_A is the impedance of the antenna and Z_C is the impedance of the chip. To maximize the efficiency of wireless power transfer to an RFID tag, antenna impedance is tuned to be conjugate matched with the input impedance of the chip at its center frequency, i.e. $Z_A = Z_C^*$. We note that a matching circuit is necessary even if an RFID tag is tuned to a single resonant frequency, as it is inevitable that some impedance mismatch remains between a dipole antenna and chip, no matter how carefully designed.

Notice that while impedance may mismatch between chip and antenna at one particular frequency, it is simultaneously possible that the impedance could perfectly match at a different frequency. We rely on this notion to provide multiple resonant frequencies for one chip. In particular, we build a multiple impedance matching network for each co-located chip. Our method uses a set of parallel inductively coupled loops which can also be seen as multiple high impedance coupling units for one chip to source the dipole [273]. Each loop is composed by series LC components and a resistor. Each branch of the series LC components couples with the stretchable dipole antenna (Fig. 5.6(a)). Further, each branch of the series LC components also matches the impedance with the feed terminal at different frequencies shown in Fig. 5.6(b). The advantage of such an approach is that the input reactance of the dipole and each loop now is only depends on the loop inductance at the corresponding required resonant frequency. Thus, we can control the reactance of each loop by carefully designing it to directly match the reactance of the chip at multiple discrete frequencies so that $X_{loop}(f) + X_A(s, f) = -X_C(f)$ where X_{loop}, X_A and X_C are the reactance of the loop, antenna and the chip respectively. For example, when dipole antenna is stretched at s_1, the sum reactance of dipole antenna and the loop (L_1, C_1) will match with the reactance of the chip at frequency f_1, while the other loops mismatch at this particular frequency.

Inductively Coupled Spacing to tune bandwidth: Besides the center frequency, RFTattoo also must tune the bandwidth over which a tag resonates. This is important since RFTattoo may need tags respond to a range of consecutive closely spaced frequencies. RFTattoo achieves this by tuning the position of the inductor loop [274]. Recall that this directly impacts the power transfer between the chip and the dipole via mutual coupling.

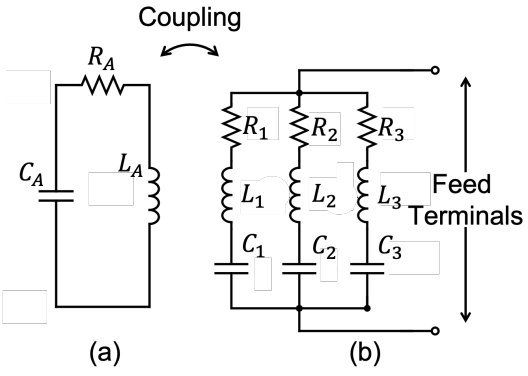

Figure 5.6: (a) Equivalent circuit of dipole antenna. (b) Equivalent circuit of multiple impedance matching networks composed of parallel inductively coupled loops. Overall, this figure shows that multiple parallel coupled loops intend to match with the impedance of feed terminals (chip) at different frequencies which in turn leads the

Specifically, a smaller spacing between the loop and dipole will increase power transfer efficiency while ensuring a narrower band tag. RFTattoo therefore achieves varied bandwidths by adapting the loop-dipole spacing. We further do so differently for each chip to reserve different desired bandwidths as shown in Fig.5.3. Our evaluation in Sec. 5.6 uses a loop-dipole spacing of 1 mm, 0.6 mm, 0.01 mm respectively.

5.4.2 Tattoo Material and Fabrication

This section describes our choice of materials in designing a stretchable RFID tag. A key trade-off that dictates our choice of materials is the balance between structural integrity and thickness. Thicker tags are more robust to repeated stretching and maintain integrity, yet are more uncomfortable to attach to skin. We therefore choose materials that provide the maximum structural integrity possible, while remaining as thin as 0.4 millimeters.

Material Selection of Stretchable Conductor Ink: The stretchable conductor is composed of a percolating network of electrically conductive filler embedded within a stretchable polymer matrix. Silver (Ag) flakes are selected as the filler due to their high conductivity (6.3×10^7 S/m) since this allows for stretchable traces with low electrical resistance. In addition, the platelet-like geometry of the flakes can reduce the electromechanical hysteresis of the composite during tensile loading and unloading cycles. We use poly(dimethylsiloxane) (PDMS) as the polymer matrix due to its low Young's modulus, high strain limit, and high elastic resilience (i.e. negligible mechanical hysteresis during loading and unloading). In order to achieve an adequate volumetric conductivity without sacrificing the elasticity of the matrix, we use a 30% volume fraction of Ag. This allows for a conductive elastomer composite with reasonable stretchability and conductivity (3000 S/cm). In addition, the printability of conductive ink is another key property for either fast prototyping or scalable manufacturing. To enable printability, we use methyl isobutyl ketone (MIBK) or order to reduce the viscosity of ink composite. The detailed fabrication process of the Ag-PDMS ink composite is presented in the Sec.5.6.

RFID Tags Fabrication: The RFID tags are compose of stretchable antennas, i.e. Ag-PDMS conductors, on a stretchable substrates. PDMS is used as the substrate to hold the unsolidified Ag-PDMS ink during curing and also to support the Ag-PDMS conductive traces during mechanical deformation. PDMS is

Table 5.1: Phoneme to viseme clustering mapping [8]

Viseme Class	Phonemes Set	Sound
V_A	/ɑː/ /aʊ/ /aɪ/ /ʌ/	car, house, fly
V_E	/e/ /eɪ/ /æ/	egg, same, cat
V_I	/iː/ /ɪ/	sheep, ship
V_O	/ɔː/ /ɔɪ/ /əʊ/	door, coin, nose
V_U	/ʊ/ /uː/	book, boot
$V_{J,CH}$	/ʤ/ /ʧ/ /ʃ/ /ʒ/	jeep, cheap, ...
$V_{P,M,B}$	/p/ /b/ /m/	pit, bit, map
$V_{F,V}$	/f/ /v/	fat, vat
$V_{D,T,S}$	/d/ /t/ /s/ /z/ /θ/ /ð/	din, tin, ...
$V_{R,W}$	/r/ /w/	run, we
$V_{G,K,N}$	/g/ /k/ /n/ /ŋ/ /l/ /y/ /h/	gut, cut, ...

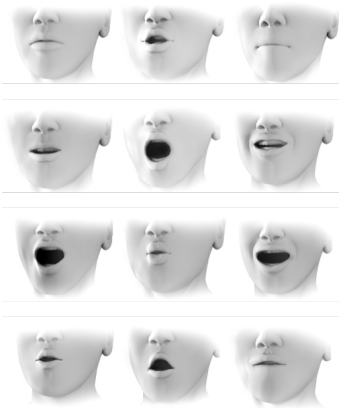

Figure 5.7: Samples of visemes [10]. 12 visemes are used to map to all possible phonemes including the rest pose.

particularly well-suited for this since it has an elastic modulus that is similar with human skin and is comfortable to wear. Moreover, the strong bond between PDMS substrates and Ag-PDMS composite facilitates the integrity of the RFID tags during mechanical deformation.

Stencil printing is used to pattern the Ag-PDMS and enables fast prototyping of the stretchable RFID tags. A commercially available Ag-filled conductive epoxy is used to solder the Ag-PDMS antennas with the packaged RFID chips. Typically, during stretching, there will be stress concentrations at the interface between the stretchable conductor and rigid RFID chip due to the compliance mismatch between the two. This stress concentration can lead to mechanical failure. To prevent this, we encapsulate the chip with an additional layer of PDMS in order to reduce the stress concentration and improve mechanical robustness.

Stretchable tag vs. Non-Stretchable tag: As mentioned in Sec. 5.3.1, RFTattoo uses a co-located stretchable and non-stretchable tag to account for location, orientation and signal multipath. While our approach above details the design of the stretchable tag, we also designed non-stretchable tags. We built the latter using copper-based antenna made from a commercially available flexible PCB (Pyralux FR8510R, DuPont) through solid ink (wax) based prototyping method [275]. In our experiments, we also use a miniature commercial RFID tag as the non-stretchable tag to be attached to the center of the stretchable tag.

5.5 From RF Signals to Speech

RFTattoo synthesizes speech by processing the stretch, location and tongue position of various points in the skin obtained from the RFID tag signals. It first uses this information to classify between various facial gestures called visemes that are unique to different sounds. RFTattoo then borrows from the rich literature on text-to-speech in natural language processing to synthesize speech in real-time.

5.5.1 Characterizing Visemes

A *viseme* is a unit of visual speech – more specifically, the visual equivalent of a phoneme (a unit of sound in speech recognition). Each viseme represents the shape of the face when the user attempts to speak a particular phoneme [236]. Past work on automated lip-reading have widely used visemes to recognize speech based on video input[276]. Recognizing shorter visemes as opposed to longer speech segments has several advantages such as needing less training effort and generalizing well for different speaker identities (speaking styles, accents, etc.).

Choice of Visemes: Phonemes map many-to-one to visemes, because many phonemes can not be distinguished using only visual cues (e.g. "p" vs. "m" sounds). Phoneme-to-viseme mappings have been constructed mainly by two approaches: linguistic and data-driven. RFTattoo uses the map from [8] obtained through a hybrid linguistic and data driven approach – a relatively sparse set which worked well experimentally (see Sec. 5.7). This map is composed by 38 phonemes and 11 classes (plus a silence class) in Table 5.1.

Viseme Classification: To classify the different viseme sets, we use the resonant frequency shift property of the RFID tattoo tag. RFID tattoos are attached to 4 different locations on the persons face: above the upper lip, below the lower lip, the left cheek and the right cheek. As the person utters the different phoneme sounds, these tags are stretched by different amounts resulting in diverse set of resonant frequencies. The response of all these tags are collected to obtain tag stretch as detailed in Sec. 5.3.

To train the classifier, we extract an extensive set of features mentioned in Table 5.2, based on the time and frequency properties of each tag's response. In particular, these features are extracted from tag stretch, tag and tongue location as well as normalized RSSI and phase values of each of the tags' response over time and frequency (within FCC bins). Notice that stretchable tags and reference tags have weak mutual coupling effect with each other. Small fluctuations in the RSSI value could occur, before extracting the features, we smooth the raw data over a 0.05 second time window. We select only the 125 most important features to train the classifier model after carrying out a set of relevance tests on the initial feature set. We train our classifier using 9 different classification models: k-Nearest Neighbors, RBF Support Vector Machine, Gaussian Process, Decision Tree, Random Forest, AdaBoost, Naive Bayes, Gaussian Mixture Model and QDA. We use Random Forest in classifying visemes which achieves the best accuracy (90%) among classifiers.

5.5.2 Speech Synthesis

Now we aim to synthesize the speech that the user intends to speak. At each unit of time, we predict a list of phoneme candidates, derived from the viseme mapping (see Table 5.1), as well as their likelihood scores. The likelihood scores are obtained by our machine learning model, which outputs the predicted viseme with corresponding probabilities for all possible visemes. Using these phoneme candidates, we can reconstruct words with ambiguities. We note that despite the 90% accuracy in viseme classification and even upon accounting for tongue position, the ambiguity of the phonemes could significantly impact speech reconstruction accuracy.

To address this, RFTattoo draws from a salient advantage of natural language processing – adjacent phonemes and words are not independent – they are limited by the English dictionary and rules of grammar. We leverage this fact to disambiguate the recognition results and recognized the transcript of what the user intends to speak. Finally, we synthesize the speech using a public text-to-speech API[3].

During the operation, our recognition algorithm will produce a prediction stream that contains the recognition result and the corresponding time windows. Each recognition result r_k consists a list of phoneme candidates $r_k c_1, r_k c_2, ..$ and associated likelihoods $r_k l_1, r_k l_2, ...$

Word & Sentence Segmentation: To perform speech recognition, we first organize the recognized phonemes into words and sentences, based on their recorded time stamps. Here, a word is comprised of phonemes, and words compose a sentence. The lengths of pauses between in-word syllables, words, and sentences often vary. We run a pilot study with four participants and empirically determine the pause thresholds for both word-level separation and sentence level separation. If the pauses between multiple adjacent phonemes are smaller than the word/sentence threshold, we group these phonemes into the same word/sentence.

On-the-fly Word Disambiguation: Next, for each set of phonemes constituting a word, we derive a set of most likely word candidates using a pronouncing dictionary [277]. A pronouncing dictionary defines the mapping between sequences of phonemes and words. We then need to select one word from each group to assemble the final sentence. Choosing the words randomly, or even the most likely word per phoneme sequence often results in gibberish, since the words may not form meaningful sentences in combination. Leveraging this fact, we build a Bayesian model to evaluate the naturalness of the sentence formed by different word sequences. Let $N(w_k|w_1, w_2, ..w_{k-1})$ denote the naturalness score of choosing the word w_n among a group G_k, given that the prior sequence $w_1, w_2, ..w_{k-1}$ is determined. The selection

[3]https://cloud.google.com/text-to-speech/

of the incoming word W_k^* is equivalent to finding word that can maximize:

$$w_k^* = \arg\max_{w_k} N(w_k|w_1, w_2, ..w_{k-1}) * l(w_k) \tag{5.4}$$

where $l(w_k)$ is the word likelihood score from the earlier viseme recognition.

We use a co-occurrence relation to measure the naturalness. We count the frequencies of m consecutive words appear in a large document collection, and use these frequencies to indicate the naturalness. In other words, the more common the word sequence is, the more natural the sequence would be. We measure the co-occurrence in a sliding window of m consecutive words.

$$N(w_k|w_1, w_2, ..w_{k-1}) = \prod_{1}^{m} p(w_k, w_{k-1}, .., w_{k-m+1}) \tag{5.5}$$

where $p(w_k, w_{k-1}, .., w_{k-m+1})$ is the non-zero frequency of these consecutive words in a large document collection. If we cannot find a specific consecutive word sequence in the corpus, We set $p = 1e^{-3}$ to avoid multiplication by zero. Our implementation sets $m = 3$ and measures the frequency in Cornell Movie Dialog Corpus [278]. While our approach is simple and easy to reproduce, a more specific and contextual corpus, such as including sentences used most commonly in daily conversation can improve performance.

5.6 Implementation

Hardware of RFID Tattoo: We build the dipole and inductor loop of RFID tattoo using an Ag-PDMS composite to make the antenna of tattoo maintain high conductivity as well as stretchablity. We use PDMS as the substrate of the tattoo, which provides up to 45 mm elongation for the selected tattoo dimensions. We use three LXMS31ACNA [279] chips to embed in each stretchable RFID tattoo (Fig. 5.8(b)). In our experiments, we use one Xerafy [280] chip as the non-stretchable reference tag which is placed on the center of each RFID tattoo. After careful fabrication and multiple interactions of adjustment, our wafer-thin RFID tattoos have a dimension of 80×10×0.5 mm.

Participants in the IRB Study: Consent forms[4] for participation in the research study were obtained at the time of the study, and all participants received a comprehensive description of the experimental procedures. We hold an active IRB protocol, which allows for attaching RFID tags to monitor bio-signals from human participants. We test our system with 10 participants from different genders, ages from 24 to 50 and include two participants who have temporary dysphonia. During the experiments, participants with healthy speech ability are also asked to speak without sound.

Evaluation and System Setup: We fully implement RFTattoo on an Impinj RFID reader attached to the user's waist. We note that we use a relatively bulky 4-antenna RFID reader as it provides wireless channel

[4]This work does not raise any ethical issues

information and is readily programmable [1]. Our system however relies on information from only one antenna, allowing for much more portable and compact commercial readers in future commercial deployments. Our reader uses frequency-hopping spread spectrum to hop fifty frequencies across 902 MHz to 928 MHz (ISM band) using a pseudo-random sequence every 200 milliseconds. RFTattoo's algorithms are implemented in Python in real-time. The RFID reader is connected with only one single antenna and is carried by the participants. We test the participants in static and mobile indoor multipath-rich office space with walls, cubicles and furniture. We chose to attach four customized RFID tattoos on the upper and lower jaw and two sides of mouth which captures the most significance features of speech [281]. We mount the antenna with different relative distances to the users face resulting in three different levels of signal-to-noise ratio. Unless specified otherwise, all our experiments below use a multipath-rich indoor office space with wall and furniture, with the reader along the user's waist and often in non-line-of-sight relative to tattoos on the face.

Calibration Procedure: Unless specified otherwise, prior to the experiments, the participants are involved in doing a light-weight calibration where the relative location of RFID tags on the face are calculated and calibrated a priori. This means that the system is re-calibrated should RFTattoo tags be peeled off and on, or with natural wear. Also, we ensure that antennas of adjacent tattoos do not overlap with each other have a gap of at least 3 mm between them. Once participant attaches the four tags to their face, they are asked to rest in place and speak three words to test the signal level of the chip responses. Since the tags are specifically designed to operate on human skin, if the tags are not properly attached to the skin, the tag response level will be weakened and limit the communication distance of our system. We also calibrate for a priori when the user makes a neutral expression to mitigate minor variations in the resonant frequency between different users due to differences in the electric properties of their skin. The whole process requires less than 2 minutes. We report our accuracy for cases where per-user prior calibration or training was not performed in Sec. 5.7.4.

Data Management: We split all collected data into separate random training and testing data with the size ratio to be 4:1 [282, 283]. We evaluate a score of the training dataset using cross-validation over five folds for overfitting analysis. Then we train our model using the training dataset. Once we have our trained model, we predict and evaluate the accuracy of the testing dataset without using cross-validation. In the evaluation section, all accuracy is testing accuracy.

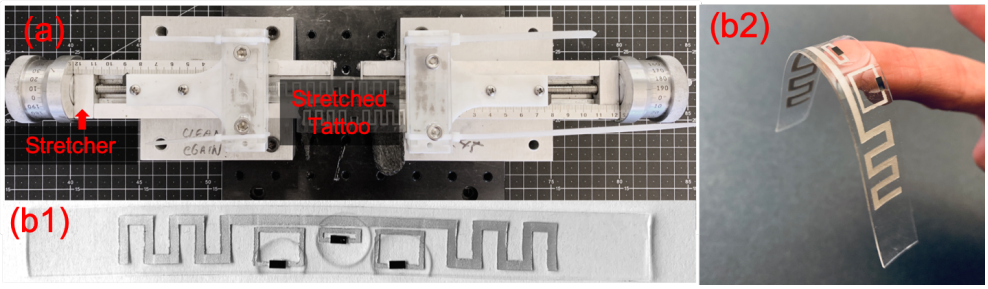

Figure 5.8: (a) Rotary system for stretch evaluation of RFTattoo. (b) RFID Tattoo.

5.7 Results

We first present our results from two microbenchmarks evaluating stretch and tongue position. We then evaluate the performance of our speech recognition system at three different levels: viseme, word and sentence.

5.7.1 Accuracy of Stretch

Method: We evaluate the accuracy of stretch on RFID Tattoos via a rotary system [284] which can control the level of stretch with a resolution of 5 μm. We manually clamp the two ends of the tattoo ensuring that it experiences exactly zero stretch force (just in the rest) when the stretcher shows zero on the scale. We then use the rotary system to stretch the RFID tattoo from 0 (in the resting state) to 30 mm in steps of 1mm. We repeat the experiments many times in different settings by placing the RFID reader in varying orientations with respect to the tattoo to emulate multipath and non-line-of-sight. As we stretch the stretchable RFID tattoos, we also co-locate our non-stretchable tattoos with the stretchable one, in order to isolate the stretch effect from the radio environment impact. We model the expected power by monitoring the response from each chip, where each RFID chip transmits its own unique pre-defined identification number following the standard RFID protocol. Further, we map the power to stretch by leveraging the ability of the tag to resonate at multiple discrete frequencies (see Sec. 5.3).

Results: We stretch the RFID tattoo from 0 to 30 mm. Fig. 5.10(a) shows the accuracy of stretch inference. We notice that RFTattoo achieves a 1.2 mm accuracy in stretch at the RSSI range of -40 to -45 dBm at around 30 cm distance between tattoo and reader. Further, our system drops in accuracy to 1.9 mm accuracy at lower RSSI (below -50 dBm) at around 1.2 m. We evaluate the tattoo's stretchability no further than 30 mm, since common facial movement produces an average stretch of 25 mm. Our RFID reader can successfully decode the response from the RFTattoo tag at a minimum of -72 dBm of RSSI which supports RFTattoo tags to be read from up to 2.5 meters away. During the experiment, we notice that as

Table 5.2: Feature Set used by RFTattoo.

Minimal Features	Specialized Features
Mean, Variance, SD	Sample Entropy
Length	FFT Coefficients
Max, Min	Aggregated FFT
Longest strike above/below mean	Wavelet Transform coefficients
Variance > SD	Energy ratio by chunks
Mean absolute change	Linear trend (absolute value, slope)
Skewness	Time series complexity
% of reoccurring values	Quantile (5%,15%,85%,95%)

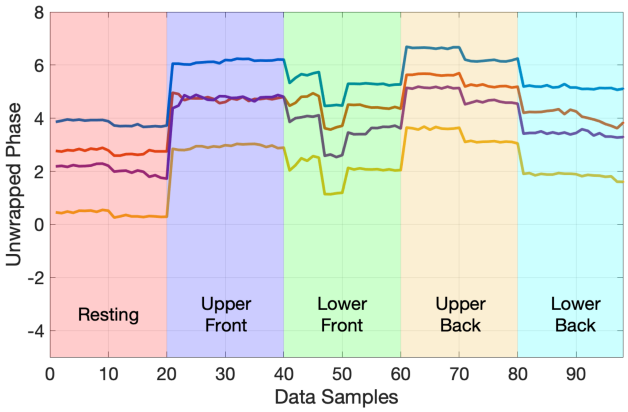

Figure 5.9: Unwrapped phase of the reflected signal from four RFID chips in different RFID tattoos. The phase relationship changes with different tongue positions due to the impedance variation. We feed this as one of the raw data into feature extraction. See details in Sec. 5.3.2

the RSSI reduces owing to the range or poor orientation, the inference accuracy drops. Since our inference model depends highly on power monitoring, we notice that a further-away tag may not respond at some resonant frequencies. One could achieve higher range performance by optimizing the material used for the RFID tags or using multiple RFID readers.

5.7.2 Tongue Positioning Accuracy

Method: RFTattoo uses tongue positions to distinguish between certain phonemes that are produced with the same facial gesture. We classify tongue positions using the resonant frequency shift of the RFID tattoo due to the near-field effects induced by the changes in tattoo's proximate environment, i.e., the tongue, as described in Sec. 5.3.2. We test this by classifying five tongue positions: (1) normal resting position, (2) upper jaw front: the tip of the tongue touching the base of the teeth on the upper jaw, (3) upper jaw back: the tip of the tongue touching the upper jaw ceiling, (4) lower jaw front: the tip of the tongue touching the base of the teeth on the lower jaw, and (5) lower jaw back: the tongue rolled inwards with the tip touching the lower jaw floor. Fig. 5.10(b) visualizes the different tongue positions. Our experiments were performed by requesting the users to keep the mouth slightly open allowing enough room for the tongue movement.

Results: Fig. 5.10(b) shows the confusion matrix for the Random Forest Classifier achieving an accuracy of 92%. Adaboost and Naive Bayes classifier also achieved similar accuracy on average. We have seen that our system is less sensitive to distinguish the position between the upper and lower jaw front, only one

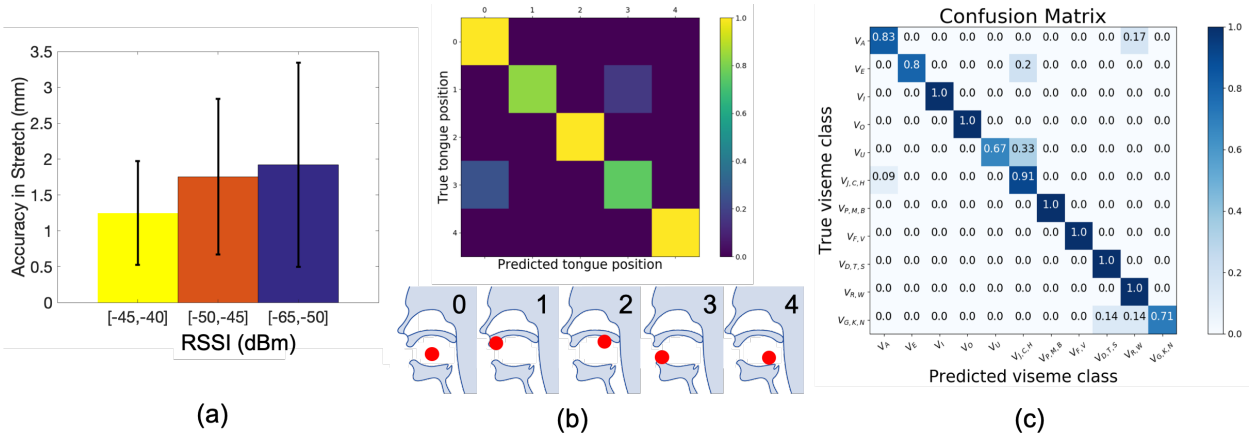

Figure 5.10: (a) Stretch Accuracy: RFTattoo achieves a 1.4 mm median accuracy in stretch inference (Facial muscle movement has an average stretch of 25 mm.) (b) Tongue positioning: RFTattoo achieves 92% mean accuracy in classification of five tongue positioning. Index 0 to 4 represent tongue positions at: normal resting position, upper jaw front, upper jaw back, lower jaw front, lower jaw back. (c) Confusion matrix of 11 visemes: RFTattoo achieves an average test accuracy of 90%.

tag on both upper and lower jaw may be the reason to have lower accuracy. Fig. 5.9 shows the phase of the signal from different RFID chips while the tongue is moving among five tongue positions. We observe that the tongue can significantly influence the signal due to the impedance variation introduced to the RFID tattoos.

5.7.3 Accuracy of Viseme Classification

Method: To classify visemes, we collected an extensive set of RFID tattoo responses corresponding to the 38 phoneme classes from our test subjects. Remember that multiple phoneme classes belong to a single viseme class. In particular, the 38 phonemes can be grouped into 11 viseme classes. Our initial classification model was built based on the phoneme classes, but we realized that the phonemes belonging to the same viseme class were difficult to differentiate for reasons discussed in Sec. 5.5.1. We then decided to use the best representative phoneme from for each viseme class and build our classifier based on them. This is because once we have a viseme class, we can rely on our natural language processing framework to produce the word and sentence by feeding the likelihood of the candidate visemes to it. The details of the classifier are described in the Sec. 5.5. We ensured that our data set contained a mixture of responses corresponding to both facial gestures accompanied by sound and just the facial gestures by having the test subjects to utter the viseme without generating any sound.

Result: We trained our system on 9 classification models (see Sec.5.5.1) and select the best five classifier (shown in Fig.5.11(a)). Our Random Forests model gives the best performance with an average accuracy

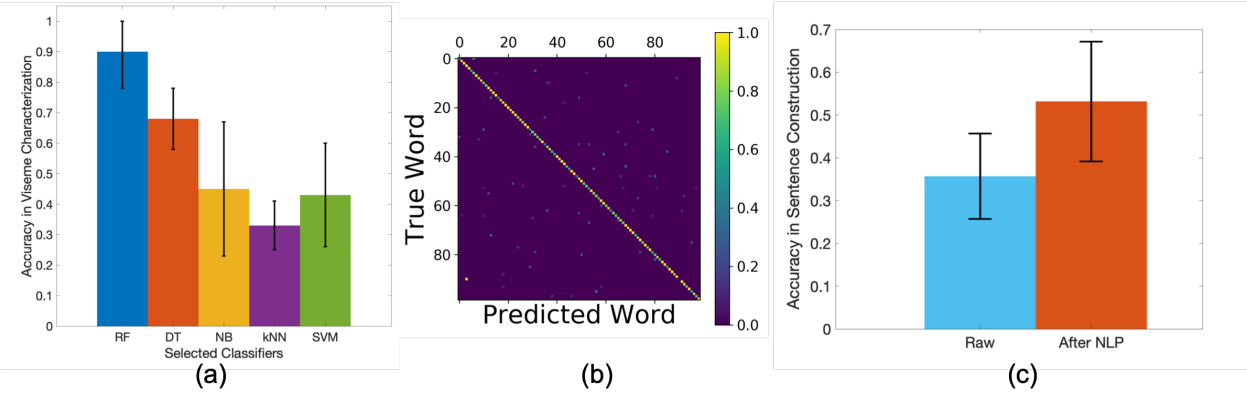

Figure 5.11: (a) Accuracy in predicting viseme classes for selected classifiers: Random Forest, Decision Tree, Naive Bayes, kNN and SVM. (b) Confusion matrix of 100 words classification. RFTattoo achieves an average test accuracy of 86%. (c) Accuracy in sentence construction.

of 90% followed by Naive Bayes (84%) and Decision Tree (72%). Initial results based on phoneme classification revealed that it is difficult to differentiate between phonemes belonging to the same viseme solely based in the time and frequency response of the RFID tattoos. This is due to the inherent similarity of the phonemes belonging to the particular viseme class, that is, these phoneme sounds are generated with the same facial gesture and tongue position, there by producing almost the similar amount of stretch and impedance change in the RFID tattoos. This led to a low classification accuracy when classifying phonemes mapping to the same viseme for each of the models, albeit a good accuracy classifying phonemes mapping to different visemes. We then chose the best representative phoneme for each of the viseme class and trained the classifier based on them. Fig. 5.10(c) shows the confusion matrix for the Random Forest Classifier based on the dataset collected across users. While almost all the viseme classes can be classified with sufficiently large accuracy, class V_U (corresponding to the "u" sound) has a low accuracy due to the nature of the stretch it produces on different face structures. While at first glance, it would appear that the Viseme class V_U will produce the most distinct response as it generates the maximum facial distortion, upon taking a closer look we find that this distortion is highly dependent on the user's specific face structure. This is not an issue with other viseme classes as the facial distortion is generally consistent across users with diverse face structures. Another observation from our results is that the classifier performance is virtually unchanged by the presence or absence of sound.

5.7.4 Accuracy in the Word Classification

Method: Although the viseme classification produced a substantially high accuracy, we observed that mapping a combinations of visemes to a word purely by stitching together the responses of different viseme candidates proved to be relatively unreliable. This is partly because the probability of an error

in the word increases exponentially with the number of phonemes it is composed of. However, such an approach would ignore the structure of the English language, where certain words are more common than others and certain combinations of phonemes are valid while others are not. This led us to building classifier system at the word-level to feed the candidate word likelihood to the NLP layer. Our vocabulary dataset is composed of 100 most commonly used words in English. To evaluate the impact of the size of vocabulary on the accuracy, we randomly select 10, 20, 30, 40, 50, 60, 70, 80, 90, 100 words in our vocabulary dataset and follow our dataset management to split the words into non-overlapping training and testing datasets. We use the same model feature configuration to train all the models for different sizes of vocabulary. We notice that different words may have varying accuracy that impacts our system, so we repeat the experiment over 100 iterations for each size of the vocabulary. Apart from the most commonly used words, we also incorporate a good mix of monosyllabic, disyllabic and tri-syllabic words to train the classification model. As in the tests for viseme classification, we ensured that the dataset contained the responses corresponding to facial expressions accompanied with sound and without sound.

Result: We trained our classifier based on the 10 classifiers listed in the previous section. We achieve an average test accuracy of 90% based on Random Forest classifier followed by 67% on Decision trees with a size of 50 words vocabulary dataset. We achieve an average accuracy of 86% with a size of 100 words. Fig. 5.11(b) shows the confusion matrix obtained on one of the test dataset containing a good mix of all the 100 words uttered in different ways (with/without generating sound). We observe that one word is nearly recognized to be a different word, that is because of the lack of speech context, which led to our NLP solution. It should also be noted that a good proportion of words can be accurately classified since the utterance of a word generates a much richer time and frequency response owing to the longer time series containing instances of multiple stretches. Fig. 5.12(a) shows our system performance versus the size of vocabulary from 10 to 100 words with a step size of 10. We observe that our system accuracy drops when the vocabulary size gets larger. One possible solution to that is to use more RFID chips that resonate at more fine-grained discrete frequencies to model the stretch.

The Effect of Misalignment: We have tested our system on previously untrained speakers, where no pre-calibration is performed before the experiments. We observe a 72% accuracy over the above 50 words using our trained model based on Random Forest classifier.

Different Types of Users: We achieve an average accuracy of 86% among 10 users. Participants who have temporary dysphonia have a very similar accuracy of 85%. During the experiments, we also observe an average accuracy of 87.1% and 86.9% for participants speaking with sound and without sound.

Number of Reader Antennas: We conduct an experiment to evaluate the impact of the number of RFID

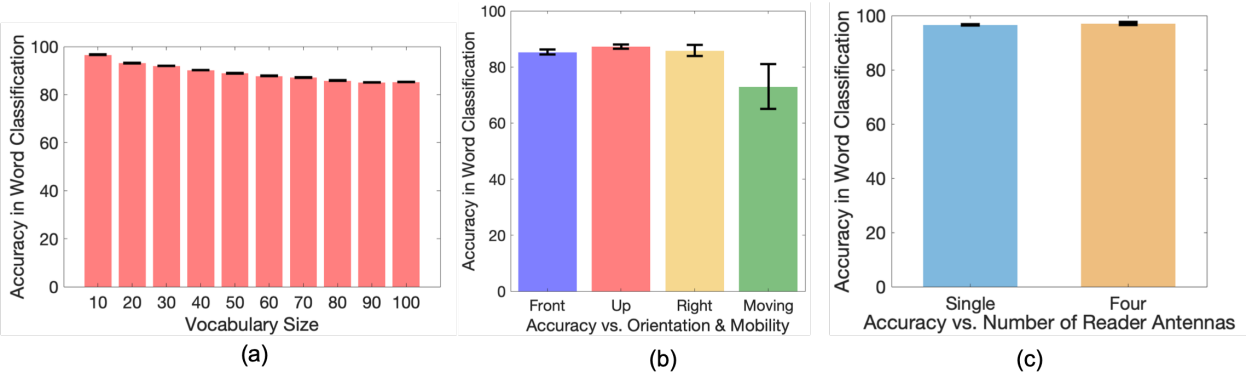

Figure 5.12: Accuracy (%) vs. (a) Vocabulary Size; (b) user's orientation; (c) number of reader antennas.

reader antennas. Specifically, we collect 20 words using a single antenna and four antennas connected to the RFID reader. The four antennas are arranged in line with 5 cm spacing. Again, we follow our dataset management mechanism to obtain the training and testing dataset (described in Sec. 5.6). In Fig. 5.12(c), we show the prediction accuracy versus the number of reader antennas, we got 93.1% and 97.2% accuracy when using a single antenna and four antennas. To process the data from four antennas, we individually feed raw data from different antennas into the prediction model and average the accuracy across all the antennas. We observe a higher accuracy when using four antennas since multiple antennas give us richer spatial diversity in the RFID tags reading (higher overall response rate of the tags and better spatial resolution). In the rest of our evaluation, we obtain results using a single reader antenna, given that it allows for more portable and cost-effective platforms.

The Effect of User Orientation: We ask users to face the reader antenna with three different orientations: front, top and right. We collect 20 words for each orientation and use them as testing data. In Fig. 5.12(b), we show that our system is robust at most of the orientations. While we get 85.3%, 87.2% and 85.8% towards different orientations, we also observe that at some orientations, the accuracy dips when multiple RFID tattoos experience extreme shadowing from the human body.

Mobility / Body Position: We conduct an experiment where users are moving with walking speed (*sim* 1 m/s) and the relative position between the tattoos and the reader antenna changes over time. We collect 20 words and use them as testing data and predict the word accuracy using our trained model of 100 words. We observe a 76.5% accuracy while the users are moving at walking speed versus accuracy of 86% when the user is static.

5.7.5 Accuracy in Sentence Construction

Method: We conducted a pilot experiment for the sentence construction in a regular office space. We first request participants to wear four RFID tags in a way they feel comfortable. We then asked participants to speak three warm-up words to get familiar with the system. We also perform a calibration phase per-user (see Sec. 5.6). During the experiment, we presented each participant a paper sheet and asked them to read the sentences on the sheet without sound, 10 repetitions for each sentences (to measure system variance). We performed two types of experiments: (1) First, the sentences on the sheet were among twenty candidate sentences often used in daily conversations with the pool of sentences known to RFTattoo a priori. (2) Second, the sentences are arbitrary chosen grammatically correct sentences that span 4 words to 10 words, the typical range of sentence lengths in English speech [285], with the pool of sentences spoken by the user not known a priori to RFTattoo. At the end of the experiment, we helped the participants remove the RFID tags and provided alcohol-soaked cotton swab to help participants clean the glue residue.

Result: Figure 5.11(c) illustrates the sentence level recognition accuracy of RFTattoo. We first observe that for the 20 commonly used sentences in day-to-day use known to our system, RFTattoo shows an average of 91% accuracy. In contrast, the raw recognition accuracy for sentences unseen by RFTattoo is 35.7%. Integrating natural language processing based correction further boosts the average accuracy to 53.2%. We note that this is within the performance range for unknown sentences of state-of-the-art vision-based lip-reading software that requires line-of-sight (e.g. 46.8% in [286]). We also find that RFTattoo works better for longer sentences, which contains more contextual information. Our results reveal that RFTattoo holds promise in reconstructing sentences for users with voice impairments. Our accuracy can be further improved over time with more data to tune to the user's particular speaking habits.

5.8 Limitations

Calibration: RFTattoo requires a light-weight calibration step to be performed a priori. To achieve optimal system performance, one must re-calibrate should RFTattoo tags be peeled off and on, or with natural wear. We describe the details of the procedure in Sec. 5.6

Extreme shadowing: Due to the 2-D dipole antenna used by RFID tags, RFTattoo tags may experience shadowing due to the human body at certain orientations of the tag relative to the reader. This can in part

be circumvented by attaching more RFID tattoos to the user's face.

Words out of the vocabulary dataset: RFTattoo has poor accuracy in predicting the untrained words. This is a common problem shared by voice recognition systems as well as lip-reading systems.

Sensing and Computational Latency: Our current RFTattoo prototype has an end-to-end refresh rate of around 0.8 Hz which is 48 words per minute, running on MacBook Pro. The window size for RFTattoo feature extraction is set to be 1 second. On average, the stretch inference and feature extraction takes 0.24 second in total, the output of our prediction model for a 4-word sentences takes 5.2 milliseconds. The Impinj RFID reader can process 3000 readings per second. We note that our end-to-end latency can be improved with better compute infrastructure at the edge.

Communication Range: RFTattoo has a limited communication range up to 2.5 meters. The reason is two-fold. First, human skin has lossy dielectric material properties, which means the range performance of the RFTattoo tags are highly associated with the permittivity and conductivity of the skin. Further, we note that the electrical properties of the human body can vary from person to person. RFTattoo shows excellent performance when the RFID reader is around the waist. Second, RFTattoo infers the tag stretch based on the response of the different resonant chips, and its range is a function of the choice of material of the tag. One could achieve a higher range performance by optimizing the material used for the RFID tags or using blind beamforming techniques [3].

5.9 Discussion

Facial Expressions: Many other facial movements other than talking can also affect the face, such as expressions (e.g., a smile), eating and drinking. This is also a well-known challenge for visual lip-reading techniques. Indeed, recognizing the type of facial expression itself remains a challenge for lip-reading systems [287]. Past work has relied on specialized features such as the movement of lip corners to distinguish between similar expressions, e.g., posed versus genuine smiles [288]. While the scope of RFTattoo is restricted to speech recognition, designing RFID tattoos which can perform advanced expression sensing by sensing the stretch of lip corners and beyond remains an open challenge for the future work.

Interference in the 915 MHz Band: We test our system where only one RFID reader is present in the environment. Our system is at least as resilient as standard RFID systems since we use commercial RFID chips within our RFID tattoos. As with any wireless system, the presence of interfering sources in the shared 915 MHz band can cause a reduction in the performance of RFTattoo. We believe future

implementations of RFTattoo can benefit from the rich literature on RFID interference cancellation [289].

RFID Tattoo Signal Collisions: RFTattoo uses commercial RFID chips which fully run the standard commercial RFID protocol (EPC Gen2 UHF RFID protocol). This protocol inherently mitigates collisions due to the presence of a large number of RFID tags based on Slotted Aloha. Indeed, the signal collision gets worse if the environment presents massive RFID tags, like the warehouses. One potential method is to pose SELECT queries [290] on RFID readers to which only a user-defined group of RFID tags respond to the RFID reader. For now, our RFID reader can process 100 RFID tags simultaneously with 3000 tag readings per second.

RFTattoo System Cost: The cost of each RFID tattoo is similar to the commercial RFID tag. While the RFID reader used in our experiment cost around $1500, there are cheaper RFID readers ($200) available in the market. We envision that RFID reader chips will be integrated with personal devices such as smartphones to preserve user privacy.

5.10 Summary

This chapter presents the first system that recognizes intended speech of users with voice impairments using light-weight RFID tattoos attached to the face. We present algorithms that recognize subtle stretches of the tattoos as well as movement of the tongue by processing signals reflected off the tags received at a handheld RFID reader. Our system then builds a natural language processing framework that recognizes various facial gestures associated with speech to construct meaningful words and sentences. We present results from a detailed user study that reveals the promise of our approach in recognizing intended speech, even when users do not make any sounds.

Chapter 6

Conclusion and Future Directions

6.1 Conclusion

This dissertation addresses and builds systems to overcome fundamental limitations of today's wireless energy delivery systems. Instead of blasting wireless energy blindly, this dissertation presents a series of blind beamforming systems to deliver precisely controllable wireless energy to battery-free devices and objects without any active feedback from them. Specifically, it enhances the connectivity of commercial battery-free tags, enables software-defined cooking and soft robot actuation, and enriches advanced sensing capabilities.

Further, this dissertation opens up several new research domains for the community of wireless system researchers. First, it explores the potential of wireless actuation. For example, this dissertation re-visits a 50-years-old invention and re-designs the microwave oven to enable software-defined cooking. In addition, this dissertation builds on the same heating concept and introduces the first distributed multi-antenna high-power RF system for general-purpose heating and demonstrates actuating battery-free robots. Second, this dissertation explores the intersection of soft materials and wireless systems. For example, it presents soft robots and stretchable tags that can harvest wireless energy efficiently and be capable of various sensing tasks. We detail the contribution as follows:

6.1.1 Connecting Commercial Battery-free Tags at Scale

PushID is the first system that pushes the communication range of commercial battery-free RFID systems by eightfold. PushID achieves this without modifying the hardware of the RFID tags. The key enabler is a novel distributed blind beamforming algorithm to efficiently discover battery-free tags without a response from tags. Other than PushID, we also developed TextileSense [7] which extends the range of commercial

NFC tags by fourfold. More importantly, we demonstrate that such a blind beamforming system can be used as an accurate human-computer interface such as hand gesture tracking. We briefly summarize the contributions of TextileSense below:

- TextileSense pushes the range limits of near-field communication (NFC). NFC is based on magnetic coupling, the essential difference from RFIDs. The magnetic coupling makes the signal power transmitted by the NFC reader drop dramatically as the signal goes. Its main goal is to extend NFC's short communication range of 5 centimeters maximum. While this dissertation presents novel blind beamforming techniques in the far field (e.g., the RFID context), doing so for near-field communication is more complicated. The energized pattern is not only determined by the set of phase shifts applied across the array of NFC readers but is also impacted by the location, orientation, and impedance of the unknown objects. To address this challenge, TextileSense develops a blind *near-field* beamforming algorithm to sense objects in the near-field with unknown locations, orientation, and impedance. TextileSense achieves a four-times improvement in range compared to commercial NFC systems. We further demonstrate that such a system unlocks many applications: high-accuracy localization, body posture sensing, and accurate user interface.

6.1.2 Software-defined Cooking & Beyond

Beyond powering battery-free devices at scale, this dissertation presents wireless systems for fine-grained heating. Specifically, SDC redesigns the microwave oven system to heat food (e.g., pizza) in a software-defined and fine-grained thermal trajectory. In addition, WASER develops a distributed multi-antenna high-power RF system for general-purpose heating, and demonstrates actuating thermally driven electronics-free soft robots.

6.1.3 Co-optimizing Soft Materials and Energy-harvesting Systems

We have collaborated with soft robotic researchers to create battery-free soft materials and machines capable of efficiently harvesting wireless energy for various sensing tasks. RFTattoo develops skin-friendly temporary tattoos that recognize the intended speech of users with speech impairments. Further, we created battery-free soft machines for shape sensing [1, 291] and flexible user interfaces [292].

6.2 Future Directions

6.2.1 Wireless Heating in Cyber-Physical Systems

My work so far has demonstrated how to harness precisely controllable wireless energy to provide long-range communication for battery-free IoT, programmable microwave cooking, etc. Indeed, applying precisely controllable wireless energy for heating purposes can also be extended to other domains: additive manufacturing and medical treatments. In the short term, I plan to develop a novel 3-D printing mechanism that improves the efficiency of additive manufacturing. Today's laser-based sintering technology has limited efficiency when fabricating plastics and metal since existing approaches require sophisticated mechanical motor systems and constructing prototypes layer by layer in a 2-D manner. A promising solution incorporating my previous work can enable a new type of 3-D printer that fabricates materials in a 3-D manner with a much lower cost, fine-grained resolution, and high efficiency. In the long term, I'm eager to explore applying precise and programmatic control of wireless energy in medical treatments, such as enabling non-invasive, safe tumor ablation and nerve stimulation. Building on my previous work to design a non-invasive approach for medical treatments, super-resolution heat distribution controls, and exploring how wireless signals interact with human bodies could have a widespread impact in healthcare, radiology, and in-body wireless system research.

6.2.2 Sustainable Wireless Infrastructure

Hundreds of cellular base stations have been deployed in urban areas, especially in the era of 5G. As a wireless system researcher working in the field of battery-free computing, I ask, "Can we build wideband, battery-free cellular tags utilizing the proliferation of wireless energy from those cellular towers to enable long-range, battery-free communication and energy-efficient sensing at city-scales?" Imagine billions of these low-cost, battery-free tags being deployed in the city. These tags can leverage the ambient energy in the air and enable city-scale, real-time sensing (such as structural health and weather monitoring) for free without installing added power-hungry infrastructure. Moreover, these tags can improve the network throughput and connectivity of the existing network infrastructure by acting as passive relays and monitoring wireless propagation characteristics (e.g., multipath effects) in the spectrum of 5G networks throughout the city. To enable these applications, it's also critical to develop an energy-efficient communication protocol to power and communicate with these tags across the city. My research has hinted at how we could minimize the transmitted wireless energy to connect passive RFID tags at scale by exploiting the collaboration between existing distributed RFID base stations (readers). Building on my previous work,

I believe developing such a city-scale, battery-free technology combined with the precise control of delivering wireless energy using the existing cellular infrastructure will lead us work toward a **sustainable** future.

6.2.3 Smart Materials for Biomedical Devices & Robotics:

Tethered and battery-powered devices can restrict mobility and flexibility and make maintenance more expensive, thereby limiting their utility in various domains. Instead, my research emphasizes that wireless and battery-free technology benefits systems and ordinary objects by miniaturizing the form factor of devices with unlimited operational lifetimes and enabling low-cost, maintenance-free sensing capabilities. I envision this allowing unique research opportunities in soft robotics and biomedical devices. For example: Can we turn ordinary manual intravenous (IV) sets into a wireless, battery-free, and low-cost sensing system that can detect air bubbles inside the IV tube, reducing heart attack risks? Can we build bio-compatible battery-free materials that monitor vital signs continuously? Can we create battery-free micro/nanorobots that can be moved by ambient wireless energy? Unique system challenges and real-world constraints must be considered for these new research domains. I want to leverage my expertise in developing low-cost, soft, battery-free, and wireless systems to explore these areas.

6.2.4 Security for Battery-free IoT

My research has demonstrated multiple sensing services and enhanced connectivity for battery-free IoT systems, exposing their vulnerabilities and raising alarming security challenges. For example, how can we hold the short-range security guarantee for battery-free IoT devices against eavesdroppers? How can we prevent adversarial sensing of our body and living environment in the full spectrum from Hz to THz? I believe solutions to these challenges rely on a fundamental understanding of physical-layer wireless networks and advanced signal processing. Possible solutions such as beamforming destructive interference at eavesdroppers could prevent adversarial sensing. I am excited to leverage my expertise in state-of-the-art wireless systems to resolve potential security challenges.

Bibliography

[1] H. Jin, J. Wang, Z. Yang, S. Kumar, and J. Hong, "Wish: Towards a wireless shape-aware world using passive rfids," in *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '18. New York, NY, USA: ACM, 2018, pp. 428–441. [Online]. Available: http://doi.acm.org/10.1145/3210240.3210328 iii, 6, 90, 97, 105, 116

[2] ——, "Rf-wear: Towards wearable everyday skeleton tracking using passive rfids," in *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*. ACM, 2018, pp. 369–372. iii, 90

[3] J. Wang, J. Zhang, R. Saha, H. Jin, and S. Kumar, "Pushing the range limits of commercial passive rfids," in *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. Boston, MA: USENIX Association, 2019, pp. 301–316. [Online]. Available: https://www.usenix.org/conference/nsdi19/presentation/wangjingxian iii, 3, 8, 58, 113

[4] J. Wang, C. Pan, H. Jin, V. Singh, Y. Jain, J. I. Hong, C. Majidi, and S. Kumar, "Rfid tattoo: A wireless platform for speech recognition," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 3, no. 4, pp. 1–24, 2019. iii, 57, 87

[5] H. Jin, J. Wang, S. Kumar, and J. Hong, "Software-defined cooking using a microwave oven," in *The 25th Annual International Conference on Mobile Computing and Networking*, 2019, pp. 1–16. iii, 3

[6] J. Wang, V. Ranganathan, J. Lester, and S. Kumar, "Ultra low-latency backscatter for fast-moving location tracking," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 6, no. 1, pp. 1–22, 2022. iii

[7] J. Wang, J. Zhang, K. Li, C. Pan, C. Majidi, and S. Kumar, "Locating everyday objects using nfc textiles," in *Proceedings of the 20th International Conference on Information Processing in Sensor Networks (co-located with CPS-IoT Week 2021)*, 2021, pp. 15–30. iii, 115

[8] S. Lee and D. Yook, "Audio-to-visual conversion using hidden markov models," in *Pacific Rim International Conference on Artificial Intelligence*. Springer, 2002, pp. 563–570. vii, 101, 102

[9] M. Cooke, J. Barker, S. Cunningham, and X. Shao, "An audio-visual corpus for speech perception and automatic speech recognition," *The Journal of the Acoustical Society of America*, vol. 120, no. 5, pp. 2421–2424, 2006. xii, 92

[10] "Wolf paulus' journal: Viseme model with 12 mouth shapes," https://wolfpaulus.com/jounal/software/lipsynchronization/, 01 2019. xiii, 101

[11] D. Vasisht, S. Kumar, H. Rahul, and D. Katabi, "Eliminating channel feedback in next-generation cellular networks," in *Proceedings of the 2016 ACM SIGCOMM Conference*, 2016, pp. 398–411. 1

[12] H. S. Rahul, S. Kumar, and D. Katabi, "Jmb: scaling wireless capacity with user demands," *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 235–246, 2012. 1

[13] S. Kumar, E. Hamed, D. Katabi, and L. Erran Li, "Lte radio analytics made easy and accessible," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 4, pp. 211–222, 2014. 1

[14] S. Kumar, D. Cifuentes, S. Gollakota, and D. Katabi, "Bringing cross-layer mimo to today's wireless lans," in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, 2013, pp. 387–398. 1

[15] E. Hamed, H. Rahul, M. A. Abdelghany, and D. Katabi, "Real-time distributed mimo systems," in *Proceedings of the 2016 ACM SIGCOMM Conference*, 2016, pp. 412–425. 1, 2, 7, 19

[16] S. Gollakota, S. D. Perli, and D. Katabi, "Interference alignment and cancellation," in *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, 2009, pp. 159–170. 1

[17] C. M. Yetis, T. Gou, S. A. Jafar, and A. H. Kayran, "On feasibility of interference alignment in mimo interference networks," *IEEE Transactions on Signal Processing*, vol. 58, no. 9, pp. 4771–4782, 2010. 1

[18] X. Xie, X. Zhang, S. Kumar, and L. E. Li, "pistream: Physical layer informed adaptive video streaming over lte," in *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, 2015, pp. 413–425. 1

[19] L. M. Ni, Y. Liu, Y. C. Lau, and A. P. Patil, "Landmarc: indoor location sensing using active rfid," *Wireless networks*, vol. 10, no. 6, pp. 701–710, 2004. 1

[20] A. Agrawal, G. J. Anderson, M. Shi, and R. Chierichetti, "Tangible play surface using passive rfid sensor array," in *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing*

Systems, ser. CHI EA '18. New York, NY, USA: ACM, 2018, pp. D101:1–D101:4. [Online]. Available: http://doi.acm.org/10.1145/3170427.3186514 1

[21] H. Jin, C. Xu, and K. Lyons, "Corona: Positioning adjacent device with asymmetric bluetooth low energy rssi distributions," in *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, ser. UIST '15. New York, NY, USA: ACM, 2015, pp. 175–179. [Online]. Available: http://doi.acm.org/10.1145/2807442.2807485 1

[22] L. Shangguan, Z. Zhou, and K. Jamieson, "Enabling gesture-based interactions with objects," in *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 2017, pp. 239–251. 1

[23] J. Wang, D. Vasisht, and D. Katabi, "Rf-idraw: virtual touch screen in the air using rf signals," in *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 4. ACM, 2014, pp. 235–246. 1, 9

[24] H. Li, P. Zhang, S. Al Moubayed, S. N. Patel, and A. P. Sample, "Id-match: A hybrid computer vision and rfid system for recognizing individuals in groups," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ser. CHI '16. New York, NY, USA: ACM, 2016, pp. 4933–4944. [Online]. Available: http://doi.acm.org/10.1145/2858036.2858209 1

[25] L. Yang, Y. Chen, X.-Y. Li, C. Xiao, M. Li, and Y. Liu, "Tagoram: Real-time tracking of mobile rfid tags to high precision using cots devices," in *Proceedings of the 20th annual international conference on Mobile computing and networking*. ACM, 2014, pp. 237–248. 1, 6, 90

[26] Y. Ma, N. Selby, and F. Adib, "Minding the billions: Ultra-wideband localization for deployed rfid tags," in *ACM MobiCom*, 2017. 1, 6, 90

[27] Wikipedia, "Airtag - wikipedia," https://en.wikipedia.org/wiki/AirTag, 02 2019, (Accessed on 02/13/2019). 1

[28] Google, "Soli - google atap," https://atap.google.com/soli/, 02 2019, (Accessed on 02/13/2019). 1

[29] M. Vollmer, "Physics of the microwave oven," *Physics Education*, vol. 39, no. 1, p. 74, 2004. 1, 34, 35, 54

[30] J. M. Osepchuk, "The history of the microwave oven: A critical review," in *2009 IEEE MTT-S International Microwave Symposium Digest*. IEEE, 2009, pp. 1397–1400. 1

[31] E. Howell, "Things you shouldn't cook in a microwave | microwave safety," https://www.livescience.com/34406-microwave-safety-experiments.html, 01 2013, (Accessed on 03/17/2019). 1, 31, 34

[32] BabyGreenThumb.com, "Safe plastic numbers guide," http://www.babygreenthumb.com/p-122-safe-plastic-numbers-guide.aspx, 06 2011, (Accessed on 02/24/2019). 1, 34

[33] K. Pitchai, "Electromagnetic and heat transfer modeling of microwave heating in domestic ovens," 2011. 1, 30

[34] H.-S. Yu, "Microwave oven with a turntable and mode stirrers," 2 1999, uS Patent 5,877,479. 1, 35

[35] M. G. Co., "Cn102635881a - microwave oven - google patents," https://patents.google.com/patent/CN102635881A/en, 08 2012, (Accessed on 03/15/2019). 1, 35

[36] J. Wang and D. Katabi, "Dude, where's my card?: Rfid positioning that works with multipath and non-line of sight," in *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4. ACM, 2013, pp. 51–62. 1, 2, 3, 9

[37] Y. Ma, X. Hui, and E. C. Kan, "3d real-time indoor localization via broadband nonlinear backscatter in passive devices with centimeter precision," in *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*. ACM, 2016, pp. 216–229. 1, 2, 3, 9

[38] I. Ahmed, H. Khammari, A. Shahid, A. Musa, K. S. Kim, E. De Poorter, and I. Moerman, "A survey on hybrid beamforming techniques in 5g: Architecture and system model perspectives," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 4, pp. 3060–3097, 2018. 1

[39] W. Roh, J.-Y. Seol, J. Park, B. Lee, J. Lee, Y. Kim, J. Cho, K. Cheun, and F. Aryanfar, "Millimeter-wave beamforming as an enabling technology for 5g cellular communications: Theoretical feasibility and prototype results," *IEEE communications magazine*, vol. 52, no. 2, pp. 106–113, 2014. 1

[40] N. D. Sidiropoulos, T. N. Davidson, and Z.-Q. Luo, "Transmit beamforming for physical-layer multicasting," *IEEE transactions on signal processing*, vol. 54, no. 6, pp. 2239–2251, 2006. 2

[41] R. Mudumbai, G. Barriac, and U. Madhow, "On the feasibility of distributed beamforming in wireless networks," *IEEE Transactions on Wireless communications*, vol. 6, no. 5, pp. 1754–1763, 2007. 2

[42] F. Rashid-Farrokhi, K. R. Liu, and L. Tassiulas, "Transmit beamforming and power control for cellular wireless systems," *IEEE Journal on selected areas in communications*, vol. 16, no. 8, pp. 1437–1450, 1998. 2

[43] R. Mudumbai, D. R. B. Iii, U. Madhow, and H. V. Poor, "Distributed transmit beamforming: challenges and recent progress," *IEEE Communications Magazine*, vol. 47, no. 2, pp. 102–110, 2009. 2

[44] V. Venkateswaran and A.-J. van der Veen, "Analog beamforming in mimo communications with phase shift networks and online channel estimation," *IEEE Transactions on Signal Processing*, vol. 58, no. 8, pp. 4131–4143, 2010. 2

[45] V. Chawla and D. S. Ha, "An overview of passive rfid," *IEEE Communications Magazine*, vol. 45, no. 9, pp. 11–17, 2007. 2

[46] M. Liukkonen and T.-N. Tsai, "Toward decentralized intelligence in manufacturing: recent trends in automatic identification of things," *The International Journal of Advanced Manufacturing Technology*, vol. 87, no. 9-12, pp. 2509–2531, 2016. 2, 6

[47] Y.-C. Chen, C. Chu, R.-S. Chen, H. M. Sun, and P. Ju, "Rfid-based bonded warehouse for science park," *International Journal of Radio Frequency Identification Technology and Applications*, vol. 5, no. 1, pp. 1–23, 2018. 2, 6

[48] A. K. Biswal, M. Jenamani, and S. K. Kumar, "Warehouse efficiency improvement using rfid in a humanitarian supply chain: Implications for indian food security system," *Transportation Research Part E: Logistics and Transportation Review*, vol. 109, pp. 205–224, 2018. 2, 6

[49] R. Vadivambal and D. Jayas, "Non-uniform temperature distribution during microwave heating of food materials - a review," *Food and Bioprocess Technology*, vol. 3, no. 2, pp. 161–171, 2010. 3, 30, 48

[50] M. J. Ford, C. P. Ambulo, T. A. Kent, E. J. Markvicka, C. Pan, J. Malen, T. H. Ware, and C. Majidi, "A multifunctional shape-morphing elastomer with liquid metal inclusions," *Proceedings of the National Academy of Sciences*, vol. 116, no. 43, pp. 21 438–21 444, 2019. 4, 57, 61

[51] Q. He, Z. Wang, Y. Wang, A. Minori, M. T. Tolley, and S. Cai, "Electrically controlled liquid crystal elastomer–based soft tubular actuator with multimodal actuation," *Science advances*, vol. 5, no. 10, p. eaax5746, 2019. 4, 57, 61

[52] T. A. Kent, M. J. Ford, E. J. Markvicka, and C. Majidi, "Soft actuators using liquid crystal elastomers with encapsulated liquid metal joule heaters," *Multifunctional Materials*, vol. 3, no. 2, p. 025003, jun 2020. [Online]. Available: https://doi.org/10.1088/2399-7532/ab835c 4, 57

[53] J. M. Boothby, J. C. Gagnon, E. McDowell, T. Van Volkenburg, L. Currano, and Z. Xia, "An Untethered Soft Robot Based on Liquid Crystal Elastomers," *Soft Robotics*, vol. 00, no. 00, pp. 1–9, 2021. 4, 57, 61

[54] Amazon, "Korad ka3010d -precision variable adjustable 30v," https://www.amazon.com/KA3010D-Precision-Variable-Adjustable-Regulated/dp/B00G0VRFUG, 01 2022, (Accessed on 01/15/2022). 4, 57

[55] S. Ma, X. Li, S. Huang, J. Hu, and H. Yu, "A light-activated polymer composite enables on-demand photocontrolled motion: Transportation at the liquid/air interface," *Angewandte Chemie*, vol. 131, no. 9, pp. 2681–2685, 2019. 4, 58

[56] A. H. Gelebart, D. Jan Mulder, M. Varga, A. Konya, G. Vantomme, E. Meijer, R. L. Selinger, and D. J. Broer, "Making waves in a photoactive polymer film," *Nature*, vol. 546, no. 7660, pp. 632–636, 2017. 4, 58, 62

[57] X. Liu, S.-K. Kim, and X. Wang, "Thermomechanical liquid crystalline elastomer capillaries with biomimetic peristaltic crawling function," *Journal of Materials Chemistry B*, vol. 4, no. 45, pp. 7293–7302, 2016. 4, 58

[58] M. Rogóż, K. Dradrach, C. Xuan, and P. Wasylczyk, "A millimeter-scale snail robot based on a light-powered liquid crystal elastomer continuous actuator," *Macromolecular rapid communications*, vol. 40, no. 16, p. 1900279, 2019. 4, 58, 61, 62

[59] H. Tian, Z. Wang, Y. Chen, J. Shao, T. Gao, and S. Cai, "Polydopamine-coated main-chain liquid crystal elastomer as optically driven artificial muscle," *ACS applied materials & interfaces*, vol. 10, no. 9, pp. 8307–8316, 2018. 4, 58, 61

[60] J. Hu, Z.-Y. Kuang, L. Tao, Y.-F. Huang, Q. Wang, H.-L. Xie, J.-R. Yin, and E.-Q. Chen, "Programmable 3d shape-change liquid crystalline elastomer based on a vertically aligned monodomain with cross-link gradient," *ACS Applied Materials & Interfaces*, vol. 11, no. 51, pp. 48 393–48 401, 2019. 4, 58, 61

[61] G. Sterpu, C. Saam, and N. Harte, "Attention-based audio-visual fusion for robust automatic speech recognition," 2018. 5, 91

[62] A. B. Hassanat, "Visual words for automatic lip-reading," *arXiv preprint arXiv:1409.6689*, 2014. 5, 87, 89

[63] N. EVENO, P. DELMAS, and P.-Y. COULON, "Vers l'extraction automatique des lèvres d'un visage parlant." 5, 87

[64] S. Agrawal, V. R. Omprakash *et al.*, "Lip reading techniques: A survey," in *Applied and Theoretical Computing and Communication Technology (iCATccT), 2016 2nd International Conference on*. IEEE, 2016, pp. 753–757. 5, 87, 89, 91

[65] Y. M. Assael, B. Shillingford, S. Whiteson, and N. de Freitas, "Lipnet: Sentence-level lipreading," *CoRR*, vol. abs/1611.01599, 2016. [Online]. Available: http://arxiv.org/abs/1611.01599 5, 87, 89, 91

[66] K. Brigham and B. V. Kumar, "Imagined speech classification with eeg signals for silent communication: a preliminary investigation into synthetic telepathy," in *Bioinformatics and Biomedical Engineering (iCBBE), 2010 4th International Conference on*. IEEE, 2010, pp. 1–4. 5, 91

[67] NASA, "Subvocal Speech," https://www.nasa.gov/centers/ames/news/releases/2004/subvocal/subvocal.html, 2004. 5, 91

[68] C. H. Loo, A. Z. Elsherbeni, F. Yang, and D. Kajfez, "Experimental and simulation investigation of rfid blind spots," *Journal of Electromagnetic Waves and Applications*, vol. 23, no. 5-6, pp. 747–760, 2009. [Online]. Available: https://doi.org/10.1163/156939309788019750 6

[69] K.-H. Hsia, M.-G. Wu, J.-N. Lin, H.-J. Zhong, and Z.-Y. Zhuang, "Development of auto-stacking warehouse truck." *JRNAL*, vol. 4, no. 4, pp. 334–337, 2018. 6

[70] H. Jin, J. Wang, Z. Yang, S. Kumar, and J. Hong, "Rf-wear: Towards wearable everyday skeleton tracking using passive rfids," in *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*, ser. UbiComp '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 369–372. [Online]. Available: https://doi.org/10.1145/3267305.3267567 6

[71] J. Lee, H. Kwon, and B. Lee, "Design consideration of uhf rfid tag for increased reading range," in *2006 IEEE MTT-S International Microwave Symposium Digest*, June 2006, pp. 1588–1591. 6

[72] N. Tran, B. Lee, and J.-W. Lee, "Development of long-range uhf-band rfid tag chip using schottky diodes in standard cmos technology," in *Radio Frequency Integrated Circuits (RFIC) Symposium, 2007 IEEE*. IEEE, 2007, pp. 281–284. 6

[73] D. Kim and J. Yeo, "Dual-band long-range passive rfid tag antenna using an amc ground plane," *IEEE Transactions on Antennas and Propagation*, vol. 60, no. 6, pp. 2620–2626, 2012. 6

[74] Y. Ma, N. Selby, and F. Adib, "Drone relays for battery-free networks," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. ACM, 2017, pp. 335–347. 6

[75] V. Pillai, H. Heinrich, D. Dieska, P. V. Nikitin, R. Martinez, and K. S. Rao, "An ultra-low-power long range battery/passive rfid tag for uhf and microwave bands with a current consumption of 700 na at 1.5 v," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 54, no. 7, pp. 1500–1512, 2007. 6

[76] Y. Ma, Z. Luo, C. Steiger, G. Traverso, and F. Adib, "Enabling deep-tissue networking for miniature medical devices," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*. ACM, 2018, pp. 417–431. 6, 9

[77] S. Chen, S. Zhong, S. Yang, and X. Wang, "A multiantenna rfid reader with blind adaptive beamforming," *IEEE Internet of Things Journal*, vol. 3, no. 6, pp. 986–996, 2016. 6, 9, 24

[78] "Regulatory status for using rfid in the epc gen2 (860 to 960 mhz) band of the uhf spectrum," https://www.gs1.org/docs/epc/uhf_regulations.pdf, 2018, (Accessed on 08/27/2018). 6

[79] C. Jiang, Y. He, X. Zheng, and Y. Liu, "Orientation-aware rfid tracking with centimeter-level accuracy," in *Proceedings of the 17th ACM/IEEE International Conference on Information Processing in Sensor Networks*. IEEE Press, 2018, pp. 290–301. 6

[80] AtlasRFIDstore, "Impinj rhcp far field rfid antenna (fcc/etsi)," https://www.atlasrfidstore.com/impinj-rhcp-far-field-rfid-antenna-fcc-etsi/, (Accessed on 04/30/2018). 7

[81] R. Bar-Yehuda and S. Even, "A linear-time approximation algorithm for the weighted vertex cover problem," *Journal of Algorithms*, vol. 2, no. 2, pp. 198–203, 1981. 7

[82] J. Jadidian and D. Katabi, "Magnetic mimo: How to charge your phone in your pocket," in *Proceedings of the 20th annual international conference on Mobile computing and networking*. ACM, 2014, pp. 495–506. 8

[83] L. Shi, Z. Kabelac, D. Katabi, and D. Perreault, "Wireless power hotspot that charges all of your devices," in *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*. ACM, 2015, pp. 2–13. 8

[84] J. R. Smith, A. P. Sample, P. S. Powledge, S. Roy, and A. Mamishev, "A wirelessly-powered platform for sensing and computation," in *International Conference on Ubiquitous Computing*. Springer, 2006, pp. 495–506. 8

[85] V. Talla, M. Hessar, B. Kellogg, A. Najafi, J. R. Smith, and S. Gollakota, "Lora backscatter: Enabling the vision of ubiquitous connectivity," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 1, no. 3, p. 105, 2017. 8

[86] A. Varshney, O. Harms, C. Pérez-Penichet, C. Rohner, F. Hermans, and T. Voigt, "Lorea: A backscatter architecture that achieves a long communication range," in *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*. ACM, 2017, p. 18. 8

[87] A. Dongare, R. Narayanan, A. Gadre, A. Luong, A. Balanuta, S. Kumar, B. Iannucci, and A. Rowe, "Charm: exploiting geographical diversity through coherent combining in low-power wide-area networks," in *Proceedings of the 17th ACM/IEEE International Conference on Information Processing in Sensor Networks*. IEEE Press, 2018, pp. 60–71. 8

[88] R. Eletreby, D. Zhang, S. Kumar, and O. Yağan, "Empowering low-power wide area networks in urban settings," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. ACM, 2017, pp. 309–321. 8

[89] N. Sornin, M. Luis, T. Eirich, T. Kramp, and O. Hersent, "Lorawan specification," *LoRa alliance*, 2015. 8

[90] J. C. Zuniga and B. Ponsard, "Sigfox system description," *LPWAN@ IETF97, Nov. 14th*, 2016. 8

[91] J.-F. Cardoso and A. Souloumiac, "Blind beamforming for non-gaussian signals," in *IEE proceedings F (radar and signal processing)*, vol. 140, no. 6. IET, 1993, pp. 362–370. 9

[92] D. T. M. Slock, "Blind fractionally-spaced equalization, perfect-reconstruction filter banks and multichannel linear prediction," in *Proceedings of ICASSP '94. IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. iv, April 1994, pp. IV/585–IV/588 vol.4. 9

[93] B. G. Agee, S. V. Schell, and W. A. Gardner, "Spectral self-coherence restoral: a new approach to blind adaptive signal extraction using antenna arrays," *Proceedings of the IEEE*, vol. 78, no. 4, pp. 753–767, April 1990. 9

[94] A. Bletsas, A. Lippman, and J. N. Sahalos, "Simple, zero-feedback, distributed beamforming with unsynchronized carriers," *IEEE journal on selected areas in communications*, vol. 28, no. 7, 2010. 9

[95] G. Sklivanitis, K. Alexandris, and A. Bletsas, "Testbed for non-coherent zero-feedback distributed beamforming," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, 2013, pp. 2563–2567. 9

[96] S. M. Birari and S. Iyer, "Mitigating the reader collision problem in rfid networks with mobile readers," in *Networks, 2005. Jointly held with the 2005 IEEE 7th Malaysia International Conference on Communication., 2005 13th IEEE International Conference on*, vol. 1. IEEE, 2005, pp. 6–pp. 9

[97] V. Sarangan, M. R. Devarapalli, and S. Radhakrishnan, "A framework for fast rfid tag reading in static and mobile environments," *Computer Networks*, vol. 52, no. 5, pp. 1058–1073, 2008. 9

[98] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong, "Argos: Practical many-antenna base stations," in *Proceedings of the 18th annual international conference on Mobile computing and networking*. ACM, 2012, pp. 53–64. 9

[99] H. S. Rahul, S. Kumar, and D. Katabi, "Jmb: Scaling wireless capacity with user demands," in *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, ser. SIGCOMM '12. New York, NY, USA: ACM, 2012, pp. 235–246. [Online]. Available: http://doi.acm.org/10.1145/2342356.2342401 11, 19, 20

[100] D. Kuester and Z. Popovic, "How good is your tag?: Rfid backscatter metrics and measurements," *IEEE Microwave Magazine*, vol. 14, no. 5, pp. 47–55, 2013. 11

[101] Q.-U.-A. Nadeem, A. Kammoun, and M.-S. Alouini, "Elevation beamforming with full dimension mimo architectures in 5g systems: A tutorial," *arXiv preprint arXiv:1805.00225*, 2018. 12

[102] D. Tse, "Fundamentals of wireless communication," 2004. 12, 16, 18

[103] J. Yang and J. Y.-T. Leung, "A generalization of the weighted set covering problem," *Naval Research Logistics (NRL)*, vol. 52, no. 2, pp. 142–149, 2005. 13, 14

[104] R. M. Karp, "Reducibility among combinatorial problems," in *Complexity of computer computations*. Springer, 1972, pp. 85–103. 14

[105] D. Alevras, M. Padberg, and M. W. Padberg, *Linear optimization and extensions: problems and solutions*. Springer Science & Business Media, 2001. 14

[106] P. Raghavan, *Randomized Algorithms*. 14

[107] U. Feige, "A threshold of ln n for approximating set cover," *Journal of the ACM (JACM)*, vol. 45, no. 4, pp. 634–652, 1998. 14

[108] S. H. Zak, V. Upatising, and S. Hui, "Solving linear programming problems with neural networks: a comparative study," *IEEE Transactions on Neural Networks*, vol. 6, no. 1, pp. 94–104, 1995. 15

[109] M. D. Zeiler, "Adadelta: an adaptive learning rate method," *arXiv preprint arXiv:1212.5701*, 2012. 15

[110] S. V. Crowder, "A simple method for studying run–length distributions of exponentially weighted moving average charts," *Technometrics*, vol. 29, no. 4, pp. 401–407, 1987. 19

[111] H. Rahul, H. Hassanieh, and D. Katabi, "Sourcesync: a distributed wireless architecture for exploiting sender diversity," *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, pp. 171–182, 2011. 20

[112] S. Wold, A. Ruhe, H. Wold, and W. Dunn, III, "The collinearity problem in linear regression. the partial least squares (pls) approach to generalized inverses," *SIAM Journal on Scientific and Statistical Computing*, vol. 5, no. 3, pp. 735–743, 1984. 22

[113] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Predictable 802.11 packet delivery from wireless channel measurements," *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, pp. 159–170, 2011. 26

[114] T. Harrold and A. Nix, "Intelligent relaying for future personal communication systems," 2000. 29

[115] A. E. Sloan, "Demographic redirection," *FOOD TECHNOLOGY*, vol. 67, no. 7, pp. 38–49, 2013. 30

[116] E. Ackerman, "A brief history of the microwave oven - ieee spectrum," https://spectrum.ieee.org/tech-history/space-age/a-brief-history-of-the-microwave-oven, 09 2016, (Accessed on 02/10/2019). 30

[117] J. M. Osepchuk, "A history of microwave heating applications," *IEEE Transactions on Microwave theory and Techniques*, vol. 32, no. 9, pp. 1200–1224, 1984. 30, 34

[118] A. Kooser, "Heat map microwave: Nuke your food with thermal predator-vision - cnet," https://www.cnet.com/news/heat-map-microwave-gives-thermal-view-while-nuking-your-food/, 02 2015, (Accessed on 11/21/2018). 30, 35

[119] N. Uhas, "How to film inside a microwave - youtube," https://www.youtube.com/watch?v=-3IpXgMPKFA, 09 2017, (Accessed on 11/21/2018). 30, 35

[120] FLIR Systems, "Infrared camera accuracy and uncertainty in plain language | flir systems," https://www.flir.com/discover/rd-science/infrared-camera-accuracy-and-uncertainty-in-plain-language/, 04 2016, (Accessed on 03/09/2019). 30, 35, 38

[121] Wikipedia, "Maillard reaction - wikipedia," https://en.wikipedia.org/wiki/Maillard_reaction, 02 2019, (Accessed on 02/13/2019). 31

[122] United States Department of Agriculture, "Microwave ovens and food safety," https://www.fsis.usda.gov/wps/portal/fsis/topics/food-safety-education/get-answers/food-safety-fact-sheets/

appliances-and-thermometers/microwave-ovens-and-food-safety/ct_index, 08 2013, (Accessed on 02/14/2019). 32, 44

[123] H. This, *Molecular gastronomy: exploring the science of flavor*. Columbia University Press, 2006. 33

[124] V. Talla, B. Kellogg, S. Gollakota, and J. R. Smith, "Battery-free cellphone," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 1, no. 2, p. 25, 2017. 33

[125] H. Jin, J. Wang, Z. Yang, S. Kumar, and J. Hong, "Wish: Towards a wireless shape-aware world using passive rfids," in *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 2018, pp. 428–441. 33

[126] H. Jin, Z. Yang, S. Kumar, and J. I. Hong, "Towards wearable everyday body-frame tracking using passive rfids," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 1, no. 4, p. 145, 2018. 33

[127] J. Wang, J. Zhang, R. Saha, H. Jin, and S. Kumar, "Pushing the range limits of commercial passive rfids," in *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. Boston, MA: USENIX Association, 2019, pp. 301–316. [Online]. Available: https://www.usenix.org/conference/nsdi19/presentation/wangjingxian 33

[128] Wikipedia, "Dielectric heating - wikipedia," https://en.wikipedia.org/wiki/Dielectric_heating, 09 2018, (Accessed on 09/05/2018). 34

[129] U.S. Food & Drug Administration, "Resources for you (radiation-emitting products) > microwave oven radiation," https://www.fda.gov/radiation-emittingproducts/resourcesforyouradiationemittingproducts/ucm252762.htm, 2017, (Accessed on 02/10/2019). 34

[130] B. Manz, "Rf energy is finally cooking," https://www.mwrf.com/community/rf-energy-finally-cooking, 12 2017, (Accessed on 03/15/2019). 35

[131] M. T. C. Smith, "Corningware browner instructions," https://www.microwavecookingforone.com/Charts/Browner.html, 01 1986, (Accessed on 03/15/2019). 35

[132] Nordic Ware, "Egg boiler | microwave | nordic ware," https://www.nordicware.com/egg-boiler, 03 2019, (Accessed on 03/15/2019). 35

[133] Wikipedia, "Susceptor - wikipedia," https://en.wikipedia.org/wiki/Susceptor, 07 2019, (Accessed on 07/04/2019). 35

[134] FISO Technologies Inc., "Mws microwave work station - product datasheet," http://tech-quality.com/images/stories/pdf/OT/mws.pdf, 12 2018, (Accessed on 11/21/2018). 35

[135] Wiciot OY, "°cmicro - °cmicro," http://www.cmicro.io/, 07 2019, (Accessed on 07/04/2019). 35

[136] M. Salakari, "Cmicro | f6s," https://www.f6s.com/cmicro, 07 2019, (Accessed on 07/04/2019). 35

[137] Y. Kawahara, X. Bian, R. Shigeta, R. Vyas, M. M. Tentzeris, and T. Asami, "Power harvesting from microwave oven electromagnetic leakage," in *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp '13. New York, NY, USA: ACM, 2013, pp. 373–382. [Online]. Available: http://doi.acm.org/10.1145/2493432.2493500 35, 54

[138] W. Wei, A. Nakamata, Y. Kawahara, and T. Asami, "Food recognition via monitoring power leakage from a microwave oven," *Journal of Information Processing*, vol. 23, no. 6, pp. 835–844, 2015. 35

[139] C. Torres, J. Chang, A. Patel, and E. Paulos, "Phosphenes: Crafting resistive heaters within thermoreactive composites," in *Proceedings of the 2019 on Designing Interactive Systems Conference*, ser. DIS '19. New York, NY, USA: ACM, 2019, pp. 907–919. [Online]. Available: http://doi.acm.org/10.1145/3322276.3322375 36

[140] M. Mizrahi, A. Golan, A. B. Mizrahi, R. Gruber, A. Z. Lachnise, and A. Zoran, "Digital gastronomy: Methods & recipes for hybrid cooking," in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, 2016, pp. 541–552. 36

[141] A. Zoran and D. Cohen, "Digital konditorei: Programmable taste structures using a modular mold," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 2018, p. 400. 36

[142] K. Fukuchi, K. Jo, A. Tomiyama, and S. Takao, "Laser cooking: a novel culinary technique for dry heating using a laser cutter and vision technology," in *Proceedings of the ACM multimedia 2012 workshop on Multimedia for cooking and eating activities*. ACM, 2012, pp. 55–58. 36

[143] J. Potter, *Cooking for geeks: real science, great hacks, and good food*. OReilly, 2016. 36

[144] Food Safety and Inspection Service, United States Department of Agriculture, "Appendix a to compliance guidelines," https://www.fsis.usda.gov/Oa/fr/95033f-a.htm?redirecthttp=true, 06 1999, (Accessed on 02/21/2019). 36

[145] P. Muehlschlegel, H.-J. Eisler, O. J. Martin, B. Hecht, and D. Pohl, "Resonant optical antennas," *science*, vol. 308, no. 5728, pp. 1607–1609, 2005. 40

[146] H. K. Khattak, P. Bianucci, and A. D. Slepkov, "Linking plasma formation in grapes to microwave resonances of aqueous dimers," *Proceedings of the National Academy of Sciences*, vol. 116, no. 10, pp. 4000–4005, 2019. 40

[147] P. Rattanadecho, "The simulation of microwave heating of wood using a rectangular wave guide: Influence of frequency and sample size," *Chemical Engineering Science*, vol. 61, no. 14, pp. 4798–4811, 2006. 41

[148] M. Azpurua, K. D. Ramos *et al.*, "A comparison of spatial interpolation methods for estimation of average electromagnetic field magnitude," *Progress In Electromagnetics Research M*, vol. 14, no. 2010, pp. 135–145, 2010. 41

[149] D. Fleet and Y. Weiss, "Optical flow estimation," in *Handbook of mathematical models in computer vision*. Springer, 2006, pp. 237–257. 42

[150] S. Haykin, *Kalman filtering and neural networks*. John Wiley & Sons, 2004, vol. 47. 43

[151] Y.-C. Fu, B. B. Yang, C.-J. Chen, and S.-S. Sheen, "An apparatus for on-line measuring electric field intensity during microwave heating: a theoretical consideration," *International journal of food science & technology*, vol. 40, no. 7, pp. 743–757, 2005. 43

[152] C. Pike-Burke and S. Grunewalder, "Optimistic planning for the stochastic knapsack problem," in *Artificial Intelligence and Statistics*, 2017, pp. 1114–1122. 46

[153] G. B. Dantzig, "Discrete-variable extremum problems," *Operations research*, vol. 5, no. 2, pp. 266–288, 1957. 46

[154] B. C. Dean, M. X. Goemans, and J. Vondrák, "Approximating the stochastic knapsack problem: The benefit of adaptivity," *Mathematics of Operations Research*, vol. 33, no. 4, pp. 945–964, 2008. 46

[155] United States Department of Agriculture, "Microwave ovens and food safety," https://www.fsis.usda.gov/wps/portal/fsis/topics/food-safety-education/get-answers/food-safety-fact-sheets/appliances-and-thermometers/microwave-ovens-and-food-safety/ct_index, 12 2017, (Accessed on 02/10/2019). 52

[156] A. Spiegel, "Pasteurized vs. homogenized milk: What's the difference? | huffpost life," https://www.huffpost.com/entry/pasteurized-homogenized-milk_n_5606168, 07 2014, (Accessed on 07/04/2019). 52

[157] T. Ludvigsen, "Hangprinter - reprap," https://reprap.org/wiki/Hangprinter, 01 2019, (Accessed on 07/03/2019). 54

[158] P. N. Shrivastava, "A practical guide to the determination of the human exposure to radio frequency fields. ncrp report no. 119. recommendations of the national council on radiation protection and measurements," *Medical Physics*, vol. 22, no. 5, pp. 658–659, 1995. [Online]. Available: https://aapm.onlinelibrary.wiley.com/doi/abs/10.1118/1.597580 54

[159] Wikipedia, "Cavity magnetron - wikipedia," https://en.wikipedia.org/wiki/Cavity_magnetron, 07 2019, (Accessed on 07/03/2019). 55

[160] M. R. Hossan, D. Byun, and P. Dutta, "Analysis of microwave heating for cylindrical shaped objects," *International Journal of Heat and Mass Transfer*, vol. 53, no. 23-24, pp. 5129–5138, 2010. 55, 63

[161] L. Herranz, W. Min, and S. Jiang, "Food recognition and recipe analysis: integrating visual content, context and external knowledge," *arXiv preprint arXiv:1801.07239*, 2018. 55

[162] CooksRecipes.com, "Internal temperature reference chart for cooking meats and poultry," https://www.cooksrecipes.com/tips/meat-cooking-temperature-chart.html, 07 2019, (Accessed on 07/05/2019). 55

[163] J. T. Gonzalez and S. E. Hudson, "Layer by layer, patterned valves enable programmable soft surfaces," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 6, no. 1, mar 2022. [Online]. Available: https://doi.org/10.1145/3517251 57

[164] S. H. Yoon, L. Paredes, K. Huo, and K. Ramani, "Multisoft: Soft sensor enabling real-time multimodal sensing with contact localization and deformation classification," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 2, no. 3, pp. 1–21, 2018. 57

[165] M. T. Tolley, R. F. Shepherd, B. Mosadegh, K. C. Galloway, M. Wehner, M. Karpelson, R. J. Wood, and G. M. Whitesides, "A resilient, untethered soft robot." 57

[166] H. Kim, S.-k. Ahn, D. M. Mackie, J. Kwon, S. H. Kim, C. Choi, Y. H. Moon, H. B. Lee, and S. H. Ko, "Shape morphing smart 3d actuator materials for micro soft robot," *Materials Today*, vol. 41, pp. 243–269, 2020. 57

[167] B. A. Trimmer, H.-T. Lin, A. Baryshyan, G. G. Leisk, and D. L. Kaplan, "Towards a biomorphic soft robot: design constraints and solutions," in *2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*. IEEE, 2012, pp. 599–605. 57

[168] W. Hu, G. Z. Lum, M. Mastrangeli, and M. Sitti, "Small-scale soft-bodied robot with multimodal locomotion," *Nature*, vol. 554, no. 7690, pp. 81–85, 2018. 57

[169] M. Wehner, R. L. Truby, D. J. Fitzgerald, B. Mosadegh, G. M. Whitesides, J. A. Lewis, and R. J. Wood, "An integrated design and fabrication strategy for entirely soft, autonomous robots," *nature*, vol. 536, no. 7617, pp. 451–455, 2016. 57

[170] L. Hines, K. Petersen, G. Z. Lum, and M. Sitti, "Soft actuators for small-scale robotics," *Advanced materials*, vol. 29, no. 13, p. 1603483, 2017. 57

[171] S. I. Rich, R. J. Wood, and C. Majidi, "Untethered soft robotics," *Nature Electronics*, vol. 1, no. 2, pp. 102–112, 2018. 57

[172] E. M. Terentjev, "Liquid-crystalline elastomers," *Journal of Physics: Condensed Matter*, vol. 11, no. 24, p. R239, 1999. 57

[173] C. Ohm, M. Brehmer, and R. Zentel, "Liquid crystalline elastomers as actuators and sensors," *Advanced materials*, vol. 22, no. 31, pp. 3366–3387, 2010. 57

[174] X. Huang, M. Ford, Z. J. Patterson, M. Zarepoor, C. Pan, and C. Majidi, "Shape memory materials for electrically-powered soft machines," *Journal of Materials Chemistry B*, vol. 8, no. 21, pp. 4539–4551, 2020. 57, 61

[175] H. Kim, J. Gibson, J. Maeng, M. O. Saed, K. Pimentel, R. T. Rihani, J. J. Pancrazio, S. V. Georgakopoulos, and T. H. Ware, "Responsive, 3d electronics enabled by liquid crystal elastomer substrates," *ACS applied materials & interfaces*, vol. 11, no. 21, pp. 19 506–19 513, 2019. 57, 73, 84

[176] T. H. Ware, M. E. McConney, J. J. Wie, V. P. Tondiglia, and T. J. White, "Voxelated liquid crystal elastomers," *Science*, vol. 347, no. 6225, pp. 982–984, feb 2015. [Online]. Available: www.sciencemag.org/content/347/6225/978/suppl/DC1 57, 86

[177] R. S. Kularatne, H. Kim, J. M. Boothby, and T. H. Ware, "Liquid crystal elastomer actuators: Synthesis, alignment, and applications," *Journal of Polymer Science Part B: Polymer Physics*, vol. 55, no. 5, pp. 395–411, mar 2017. [Online]. Available: https://onlinelibrary.wiley.com/doi/full/10.1002/polb.24287https://onlinelibrary.wiley.com/doi/abs/10.1002/polb.24287https://onlinelibrary.wiley.com/doi/10.1002/polb.24287 57

[178] C. Yakacki, M. Saed, D. Nair, T. Gong, S. Reed, and C. Bowman, "Tailorable and programmable liquid-crystalline elastomers using a two-stage thiol–acrylate reaction," *Rsc Advances*, vol. 5, no. 25, pp. 18 997–19 001, 2015. 57

[179] M. Zadan, D. K. Patel, A. P. Sabelhaus, J. Liao, A. Wertz, L. Yao, and C. Majidi, "Liquid crystal elastomer with integrated soft thermoelectrics for shape memory actuation and energy harvesting," *Advanced Materials*, p. 2200857, 2022. 58

[180] X. Wang, Y. Wang, X. Wang, H. Niu, B. Ridi, J. Shu, X. Fang, C. Li, B. Wang, Y. Gao *et al.*, "A study of the microwave actuation of a liquid crystalline elastomer," *Soft Matter*, vol. 16, no. 31, pp. 7332–7341, 2020. 60, 62, 70, 82

[181] K. Behnke, "Is this anything to worry about? 5g health issues explained," https://www.grandmetric.com/2019/03/26/5g-health-issues-explained/, 01 2022, (Accessed on 01/15/2022). 60

[182] Metageek, "Wi-fi signal strength basics," https://www.metageek.com/training/resources/wifi-signal-strength-basics/, 01 2022, (Accessed on 01/15/2022). 60

[183] D. Rus and M. T. Tolley, "Design, fabrication and control of soft robots," *Nature*, vol. 521, no. 7553, pp. 467–475, 2015. 61

[184] M. Rogóż, H. Zeng, C. Xuan, D. S. Wiersma, and P. Wasylczyk, "Light-driven soft robot mimics caterpillar locomotion in natural scale," *Advanced Optical Materials*, vol. 4, no. 11, pp. 1689–1694, 2016. 61

[185] C. Huang, J.-a. Lv, X. Tian, Y. Wang, Y. Yu, and J. Liu, "Miniaturized swimming soft robot with complex movement actuated and controlled by remote light signals," *Scientific reports*, vol. 5, no. 1, pp. 1–8, 2015. 61

[186] T. Wallin, J. Pikul, and R. Shepherd, "3d printing of soft robotic systems," *Nature Reviews Materials*, vol. 3, no. 6, pp. 84–100, 2018. 61

[187] J. Huber, N. Fleck, and M. Ashby, "The selection of mechanical actuators based on performance indices," *Proceedings of the Royal Society of London. Series A: Mathematical, physical and engineering sciences*, vol. 453, no. 1965, pp. 2185–2205, 1997. 61

[188] C. Yuan, D. J. Roach, C. K. Dunn, Q. Mu, X. Kuang, C. M. Yakacki, T. Wang, K. Yu, and H. J. Qi, "3d printed reversible shape changing soft actuators assisted by liquid crystal elastomers," *Soft Matter*, vol. 13, no. 33, pp. 5558–5568, 2017. 61

[189] M. Wang, Z.-W. Cheng, B. Zuo, X.-M. Chen, S. Huang, and H. Yang, "Liquid crystal elastomer electric locomotives," *ACS Macro Letters*, vol. 9, no. 6, pp. 860–865, 2020. 61

[190] A. D. Marchese, R. Tedrake, and D. Rus, "Dynamics and trajectory optimization for a soft spatial fluidic elastomer manipulator," *The International Journal of Robotics Research*, vol. 35, no. 8, pp. 1000–1019, 2016. 61

[191] R. F. Shepherd, F. Ilievski, W. Choi, S. A. Morin, A. A. Stokes, A. D. Mazzeo, X. Chen, M. Wang, and G. M. Whitesides, "Multigait soft robot," *Proceedings of the national academy of sciences*, vol. 108, no. 51, pp. 20 400–20 403, 2011. 61, 90

[192] C. Wang, K. Sim, J. Chen, H. Kim, Z. Rao, Y. Li, W. Chen, J. Song, R. Verduzco, and C. Yu, "Soft ultrathin electronics innervated adaptive fully soft robots," *Advanced Materials*, vol. 30, no. 13, p. 1706695, 2018. 61

[193] S. Labware, "Digital hotplate/stirrer, 115v | southern labware," https://www.southernlabware.com/digital-hotplate-stirrer-115v.html?gclid=Cj0KCQiAr5iQBhCsARIsAPcwRONpHh-5qj3IzA89_DuiWzwY9Dr5fRnJ2s6lzhDL18w8BDwBdlSZJWAaAv8IEALw_wcB, 08 2012, (Accessed on 03/15/2019). 61

[194] Y. Nishikawa and M. Matsumoto, "A design of fully soft robot actuated by gas–liquid phase change," *Advanced Robotics*, vol. 33, no. 12, pp. 567–575, 2019. 61

[195] T. Hellebrekers, K. B. Ozutemiz, J. Yin, and C. Majidi, "Liquid metal-microelectronics integration for a sensorized soft robot skin," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 5924–5929. 61, 91

[196] M. D. Bartlett, N. Kazem, M. J. Powell-Palm, X. Huang, W. Sun, J. A. Malen, and C. Majidi, "High thermal conductivity in soft elastomers with elongated liquid metal inclusions," *Proceedings of the National Academy of Sciences*, vol. 114, no. 9, pp. 2143–2148, 2017. 61

[197] J. Morrow, S. Hemleben, and Y. Menguc, "Directly fabricating soft robotic actuators with an open-source 3-d printer," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 277–281, 2016. 61

[198] B. Zuo, M. Wang, B.-P. Lin, and H. Yang, "Visible and infrared three-wavelength modulated multi-directional actuators," *Nature communications*, vol. 10, no. 1, pp. 1–11, 2019. 62

[199] M. Abu-Saleem, Y. Zhuge, R. Hassanli, M. Ellis, M. M. Rahman, and P. Levett, "Microwave radiation treatment to improve the strength of recycled plastic aggregate concrete," *Case Studies in Construction Materials*, vol. 15, p. e00728, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2214509521002436 62

[200] I. A. Ozkan, B. Akbudak, and N. Akbudak, "Microwave drying characteristics of spinach," *Journal of food engineering*, vol. 78, no. 2, pp. 577–583, 2007. 62

[201] T. K. Palazoğlu and W. Miran, "Experimental comparison of microwave and radio frequency tempering of frozen block of shrimp," *Innovative Food Science & Emerging Technologies*, vol. 41, pp. 292–300, 2017. 62

[202] U. Rosenberg and W. Bogl, "Microwave pasteurization, sterilization, blanching, and pest control in the food industry," *Food technology (USA)*, 1987. 62

[203] S. Young, D. Graves, M. Rohrer, and R. Bulard, "Microwave sterilization of nitrous oxide nasal hoods contaminated with virus," *Oral Surgery, Oral Medicine, Oral Pathology*, vol. 60, no. 6, pp. 581–585, 1985. [Online]. Available: https://www.sciencedirect.com/science/article/pii/003042208590355X 62

[204] S. Gertsman, A. Agarwal, K. O'Hearn, R. Webster, A. Tsampalieros, N. Barrowman, M. Sampson, L. Sikora, E. Staykov, R. Ng, J. Gibson, T. Dinh, K. Agyei, G. Chamberlain, and J. McNally, "Microwave- and heat-based decontamination of n95 filtering facepiece respirators: a systematic review," *Journal of Hospital Infection*, vol. 106, no. 3, pp. 536–553, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0195670120304035 62

[205] A. de la Hoz, A. Díaz-Ortiz, and P. Prieto, "Microwave-assisted green organic synthesis," 2016. 62

[206] M. Gupta and E. W. W. Leong, *Microwaves and metals*. John Wiley & Sons, 2008. 63

[207] Y. M. Lu and M. Vetterli, "Sparse spectral factorization: Unicity and reconstruction algorithms," in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2011, pp. 5976–5979. 66

[208] S. Nayer, P. Narayanamurthy, and N. Vaswani, "Provable low rank phase retrieval," *IEEE Transactions on Information Theory*, vol. 66, no. 9, pp. 5875–5903, 2020. 66

[209] S. Nayer and N. Vaswani, "Sample-efficient low rank phase retrieval," *IEEE Transactions on Information Theory*, vol. 67, no. 12, pp. 8190–8206, 2021. 66

[210] E. J. Candes, T. Strohmer, and V. Voroninski, "Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming," *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1241–1274, 2013. 66

[211] H. Sahinoglou and S. D. Cabrera, "On phase retrieval of finite-length sequences using the initial time sample," *IEEE transactions on circuits and systems*, vol. 38, no. 8, pp. 954–958, 1991. 66

[212] M. Fickus, D. G. Mixon, A. A. Nelson, and Y. Wang, "Phase retrieval from very few measurements," *Linear Algebra and its Applications*, vol. 449, pp. 475–499, 2014. 66

[213] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004. 66

[214] B. S. Guru and H. R. Hiziroglu, *Electromagnetic field theory fundamentals*. Cambridge university press, 2009. 67

[215] Molex, "Antenna omni-directional 3.3dbi gain 2483.5mhz," https://www.arrow.com/en/products/0479480001/molex, 01 2022, (Accessed on 01/15/2022). 69

[216] R. Semiconductor, "Sml-p11utt86r - rohm semiconductor," https://fscdn.rohm.com/en/products/databook/datasheet/opto/led/chip_mono/sml-p11-e.pdf, 01 2022, (Accessed on 01/15/2022). 70

[217] C. Design, "Successful induction heating of rcs billets," 2008. 72

[218] S. W. Ula, N. A. Traugutt, R. H. Volpe, R. R. Patel, K. Yu, and C. M. Yakacki, "Liquid crystal elastomers: An introduction and review of emerging technologies," *Liquid Crystals Reviews*, vol. 6, no. 1, pp. 78–107, 2018. 73

[219] A. F. Silva, H. Paisana, T. Fernandes, J. Góis, A. Serra, J. F. Coelho, A. T. de Almeida, C. Majidi, and M. Tavakoli, "High resolution soft and stretchable circuits with pva/liquid-metal mediated printing," *Advanced Materials Technologies*, vol. 5, no. 9, p. 2000343, 2020. 73, 74

[220] M. Tavakoli, M. H. Malakooti, H. Paisana, Y. Ohm, D. Green Marques, P. Alhais Lopes, A. P. Piedade, A. T. de Almeida, and C. Majidi, "Egain-assisted room-temperature sintering of silver nanoparticles for stretchable, inkjet-printed, thin-film electronics," *Advanced Materials*, vol. 30, no. 29, p. 1801852, 2018. 73, 74

[221] NXP, "Rf ldmos integrated power amplifier," https://www.nxp.com/docs/en/data-sheet/MHT2012N.pdf, 07 2018. 75

[222] L. C. Inc., "Rf ant 2.4ghz whip str sma male," https://www.digikey.com/en/products/detail/laird-connectivity-inc/WXE2400SM/2392260, 01 2022, (Accessed on 01/15/2022). 75

[223] J. T. Murugan, T. S. Kumar, P. Salil, and C. Venkatesh, "Dual frequency selective transparent front doors for microwave oven with different opening areas," *progress in electromagnetics research letters*, vol. 52, pp. 11–16, 2015. 75

[224] Amazon, "Faraday fabric," https://www.amazon.com/dp/B0875V2F7M?ref_=cm_sw_r_cp_ud_dp_VBFCZDPK98VWAM2W3S1Z, 03 2022, (Accessed on 03/15/2022). 76

[225] H. Shielding, "Emi/rfi-shielded faraday tent," https://hollandshielding.com/EMI-RFI-shielded-Faraday-tent, 03 2022, (Accessed on 03/15/2022). 76

[226] X. Huang, Z. Ren, and C. Majidi, "Soft thermal actuators with embedded liquid metal microdroplets for improved heat management," in *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, 2020, pp. 367–372. 85

[227] E. Hawkes, B. An, N. M. Benbernou, H. Tanaka, S. Kim, E. D. Demaine, D. Rus, and R. J. Wood, "Programmable matter by folding," *Proceedings of the National Academy of Sciences*, vol. 107, no. 28, pp. 12 441–12 445, 2010. 86

[228] "How a new technology is changing the lives of people who cannot speak," https://www.theguardian.com/news/2018/jan/23/voice-replacement-technology-adaptive-alternative-communication-vocalid, 07 2019. 87

[229] "Voice disorder - an overview | sicencedirect topics," https://www.sciencedirect.com/topics/medicine-and-dentistry/voice-disorder, 12 2018. 87

[230] L. Muda, M. Begam, and I. Elamvazuthi, "Voice recognition algorithms using mel frequency cepstral coefficient (mfcc) and dynamic time warping (dtw) techniques," *arXiv preprint arXiv:1003.4083*, 2010. 87

[231] A. Bedri, H. Sahni, P. Thukral, T. Starner, D. Byrd, P. Presti, G. Reyes, M. Ghovanloo, and Z. Guo, "Toward silent-speech control of consumer wearables," *Computer*, vol. 48, no. 10, pp. 54–62, 2015. 87, 91

[232] P. Suppes, Z.-L. Lu, and B. Han, "Brain wave recognition of words," *Proceedings of the National Academy of Sciences*, vol. 94, no. 26, pp. 14 965–14 969, 1997. 87

[233] M. Janke and L. Diener, "Emg-to-speech: Direct generation of speech from facial electromyographic signals," *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 25, no. 12, pp. 2375–2385, 2017. 87

[234] "Textspeak," https://www.alimed.com/textspeak.html, 12 2018. 87

[235] "Nu-vois iii artificial larynx," http://www.weitbrecht.com/NuVois-III-speech-aid.html, 01 2019. 88

[236] C. G. Fisher, "Confusions among visually perceived consonants," *Journal of Speech and Hearing Research*, vol. 11, no. 4, pp. 796–804, 1968. 88, 102

[237] "How siri works," https://electronics.howstuffworks.com/gadgets/high-tech-gadgets/siri4.html, 01 2019. 89

[238] "The oec: Facts about the language," https://web.archive.org/web/20111226085859/http://oxforddictionaries.com/words/the-oec-facts-about-the-language, 12 2018. 89

[239] J. Wang, O. Abari, and S. Keshav, "Challenge: Rfid hacking for fun and profit," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '18. New York, NY, USA: ACM, 2018, pp. 461–470. [Online]. Available: http://doi.acm.org/10.1145/3241539.3241561 90

[240] J. Yin, J. Yi, M. K. Law, Y. Ling, M. C. Lee, K. P. Ng, B. Gao, H. C. Luong, A. Bermak, M. Chan *et al.*, "A system-on-chip epc gen-2 passive uhf rfid tag with embedded temperature sensor," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 11, pp. 2404–2420, 2010. 90

[241] J. Gao, J. Siden, and H.-E. Nilsson, "Printed electromagnetic coupler with an embedded moisture sensor for ordinary passive rfid tags," *IEEE Electron Device Letters*, vol. 32, no. 12, pp. 1767–1769, 2011. 90

[242] D. J. Yeager, J. Holleman, R. Prasad, J. R. Smith, and B. P. Otis, "Neuralwisp: A wirelessly powered neural interface with 1-m range," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 3, no. 6, pp. 379–387, 2009. 90

[243] J. Wang, J. Xiong, X. Chen, H. Jiang, R. K. Balan, and D. Fang, "Tagscan: Simultaneous target imaging and material identification with commodity rfid devices," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '17. New York, NY, USA: ACM, 2017, pp. 288–300. [Online]. Available: http://doi.acm.org/10.1145/3117811.3117830 90

[244] S. Pradhan, E. Chai, K. Sundaresan, L. Qiu, M. A. Khojastepour, and S. Rangarajan, "Rio: A pervasive rfid-based touch gesture interface," in *Proceedings of the 23rd Annual International*

Conference on Mobile Computing and Networking, ser. MobiCom '17. New York, NY, USA: ACM, 2017, pp. 261–274. [Online]. Available: http://doi.acm.org/10.1145/3117811.3117818 90

[245] U. Ha, Y. Ma, Z. Zhong, T.-M. Hsu, and F. Adib, "Learning food quality and safety from wireless stickers," in *Proceedings of the 17th ACM Workshop on Hot Topics in Networks*, ser. HotNets '18. New York, NY, USA: ACM, 2018, pp. 106–112. [Online]. Available: http://doi.acm.org/10.1145/3286062.3286078 90

[246] X. Huang, K. Kumar, M. K. Jawed, A. M. Nasab, Z. Ye, W. Shan, and C. Majidi, "Chasing biomimetic locomotion speeds: Creating untethered soft robots with shape memory alloy actuators," *Science Robotics*, vol. 3, no. 25, p. eaau7557, 2018. 90

[247] E. J. Markvicka, M. D. Bartlett, X. Huang, and C. Majidi, "An autonomously electrically self-healing liquid metal–elastomer composite for robust soft-matter robotics and electronics," *Nature materials*, p. 1, 2018. 90

[248] R. Liu, Q. Shao, S. Wang, C. Ru, D. Balkcom, and X. Zhou, "Reconstructing human joint motion with computational fabrics," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 3, no. 1, pp. 19:1–19:26, Mar. 2019. [Online]. Available: http://doi.acm.org/10.1145/3314406 91

[249] Y. Liu, M. Pharr, and G. A. Salvatore, "Lab-on-skin: a review of flexible and stretchable electronics for wearable health monitoring," *ACS nano*, vol. 11, no. 10, pp. 9614–9635, 2017. 91

[250] K.-I. Jang, S. Y. Han, S. Xu, K. E. Mathewson, Y. Zhang, J.-W. Jeong, G.-T. Kim, R. C. Webb, J. W. Lee, T. J. Dawidczyk *et al.*, "Rugged and breathable forms of stretchable electronics with adherent composite substrates for transcutaneous monitoring," *Nature communications*, vol. 5, p. 4779, 2014. 91

[251] C. Occhiuzzi, C. Paggi, and G. Marrocco, "Passive rfid strain-sensor based on meander-line antennas," *IEEE Transactions on Antennas and Propagation*, vol. 59, no. 12, pp. 4836–4840, 2011. 91, 98

[252] L. Teng, K. Pan, M. P. Nemitz, R. Song, Z. Hu, and A. A. Stokes, "Soft radio-frequency identification sensors: Wireless long-range strain sensors using radio-frequency identification," *Soft robotics*, 2018. 91

[253] T. Le, R. A. Bahr, M. M. Tentzeris, B. Song, and C.-p. Wong, "A novel chipless rfid-based stretchable and wearable hand gesture sensor," in *Microwave Conference (EuMC), 2015 European*. IEEE, 2015, pp. 371–374. 91

[254] J. Kim, Z. Wang, and W. S. Kim, "Stretchable rfid for wireless strain sensing with silver nano ink," *IEEE Sensors Journal*, vol. 14, no. 12, pp. 4395–4401, 2014. 91

[255] H.-L. C. Kao, C. Holz, A. Roseway, A. Calvo, and C. Schmandt, "Duoskin: Rapidly prototyping on-skin user interfaces using skin-friendly materials," in *Proceedings of the 2016 ACM International Symposium on Wearable Computers*, ser. ISWC '16. New York, NY, USA: ACM, 2016, pp. 16–23. [Online]. Available: http://doi.acm.org/10.1145/2971763.2971777 91

[256] A. Withana, D. Groeger, and J. Steimle, "Tacttoo: A thin and feel-through tattoo for on-skin tactile output," in *The 31st Annual ACM Symposium on User Interface Software and Technology*. ACM, 2018, pp. 365–378. 91

[257] M. Weigel, A. S. Nittala, A. Olwal, and J. Steimle, "Skinmarks: Enabling interactions on body landmarks using conformal skin electronics," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2017, pp. 3095–3105. 91

[258] J. Lo, D. J. L. Lee, N. Wong, D. Bui, and E. Paulos, "Skintillates: Designing and creating epidermal interactions," in *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. ACM, 2016, pp. 853–864. 91

[259] M. Weigel, T. Lu, G. Bailly, A. Oulasvirta, C. Majidi, and J. Steimle, "Iskin: flexible, stretchable and visually customizable on-body touch sensors for mobile computing," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2015, pp. 2991–3000. 91

[260] E. Markvicka, G. Wang, Y.-C. Lee, G. Laput, C. Majidi, and L. Yao, "Electrodermis: Fully untethered, stretchable, and highly-customizable electronic bandages," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ser. CHI '19. New York, NY, USA: ACM, 2019, pp. 632:1–632:10. [Online]. Available: http://doi.acm.org/10.1145/3290605.3300862 91

[261] B. J. Betts and C. Jorgensen, "Small vocabulary recognition using surface electromyography in an acoustically harsh environment," 2005. 91

[262] H. Manabe, A. Hiraiwa, and T. Sugimura, ""unvoiced speech recognition using emg - mime speech recognition"," in *CHI '03 Extended Abstracts on Human Factors in Computing Systems*, ser. CHI EA '03. New York, NY, USA: ACM, 2003, pp. 794–795. [Online]. Available: http://doi.acm.org/10.1145/765891.765996 91

[263] A. Kapur, S. Kapur, and P. Maes, "Alterego: A personalized wearable silent speech interface," in *23rd International Conference on Intelligent User Interfaces*. ACM, 2018, pp. 43–53. 91

[264] M. Fukumoto, "Silentvoice: Unnoticeable voice input by ingressive speech," in *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '18. New York, NY, USA: ACM, 2018, pp. 237–246. [Online]. Available: http://doi.acm.org/10.1145/3242587.3242603 91

[265] G. Wang, Y. Zou, Z. Zhou, K. Wu, and L. M. Ni, "We can hear you with wi-fi!" *IEEE Transactions on Mobile Computing*, vol. 15, no. 11, pp. 2907–2920, 2016. 91

[266] V. Young and A. Mihailidis, "Difficulties in automatic speech recognition of dysarthric speakers and implications for speech-based applications used by the elderly: A literature review," *Assistive Technology*, vol. 22, no. 2, pp. 99–112, 2010. 91

[267] P. Bhartia, K. V. S. Rao, and R. S. Tomar, *Millimeter-wave microstrip and printed circuit antennas / P. Bhartia, K.V.S. Rao, R.S. Tomar*. Artech House Boston, 1991. 93

[268] D. Jackson and N. Alexopoulos, "Analysis of planar strip geometries in a substrate-superstrate configuration," *IEEE transactions on antennas and propagation*, vol. 34, no. 12, pp. 1430–1438, 1986. 93

[269] "Impinj rfid reader - itemsense reference guide," https://developer.impinj.com/itemsense/docs/reference-guide, 12 2018. 94

[270] H. Ansoft, "Ver. 13," *Ansoft Corporation*, 2011. 95

[271] U. Tata, H. Huang, R. Carter, and J. Chiao, "Exploiting a patch antenna for strain measurements," *Measurement Science and Technology*, vol. 20, no. 1, p. 015201, 2008. 95

[272] D. O. Oyeka, J. C. Batchelor, and A. M. Ziai, "Effect of skin dielectric properties on the read range of epidermal ultra-high frequency radio-frequency identification tags," *Healthcare technology letters*, vol. 4, no. 2, pp. 78–81, 2017. 98

[273] F. Costa, S. Genovesi, and A. Monorchio, "Chipless rfid trasponders by using multi-resonant high-impedance surfaces," *2013 International Symposium on Electromagnetic Theory*, pp. 401–403, 2013. 99

[274] H.-W. Son and C.-S. Pyo, "Design of rfid tag antennas using an inductively coupled feed," *Electronics Letters*, vol. 41, no. 18, pp. 994–996, 2005. 99

[275] L. Catarinucci, R. Colella, and L. Tarricone, "Smart prototyping techniques for uhf rfid tags: electromagnetic characterization and comparison with traditional approaches," *Progress in Electromagnetics Research*, vol. 132, pp. 91–111, 2012. 101

[276] H. L. Bear, "Decoding visemes: improving machine lipreading (phd thesis)," *CoRR*, vol. abs/1710.01288, 2017. [Online]. Available: http://arxiv.org/abs/1710.01288 102

[277] C. S. Group, "The cmu pronouncing dictionary," http://www.speech.cs.cmu.edu/cgi-bin/cmudict, 2019, (Accessed on 01/17/2019). 103

[278] C. Danescu-Niculescu-Mizil and L. Lee, "Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs," in *Proceedings of the 2nd Workshop on Cognitive Modeling and Computational Linguistics*. Association for Computational Linguistics, 2011, pp. 76–87. 104

[279] "Murata electronics," https://www.mouser.com/murataelectronics/, 01 2019. 104

[280] "Xerafy dot-on," https://docs.wixstatic.com/ugd/9b3bd3_6abffdaa1ce94028b4ff2f0c40020ce7.pdf, 01 2019. 104

[281] A. Koumparoulis, G. Potamianos, Y. Mroueh, and S. J. Rennie, "Exploring roi size in deep learning based lipreading." 105

[282] A. Krizhevsky *et al.*, "Learning multiple layers of features from tiny images," Citeseer, Tech. Rep., 2009. 105

[283] Y. LeCun, "The mnist database of handwritten digits," *http://yann. lecun. com/exdb/mnist/*. 105

[284] "Velmex positioning products," https://www.velmex.com/Products/index.html, 12 2018. 106

[285] M. M. Nice, "Length of sentences as a criterion of a child's progress in speech," *Journal of Educational Psychology*, vol. 16, pp. 370–379, 1925. 112

[286] J. S. Chung, A. W. Senior, O. Vinyals, and A. Zisserman, "Lip reading sentences in the wild." in *CVPR*, 2017, pp. 3444–3453. 112

[287] M. Pantic and I. Patras, "Dynamics of facial expression: recognition of facial actions and their temporal segments from face profile image sequences," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 36, no. 2, pp. 433–449, 2006. 113

[288] H. Guo, X.-H. Zhang, J. Liang, and W.-J. Yan, "The dynamic features of lip corners in genuine and posed smiles," *Frontiers in psychology*, vol. 9, p. 202, 2018. 113

[289] M. Thomas, W. Colleran, E. Fountain, and T. Humes, "Interference cancellation in rfid systems," May 11 2006, uS Patent App. 10/981,893. 114

[290] H. Barthel, "Epcglobal–rfid standards & regulations," 2005. 114

[291] H. Jin, Z. Yang, S. Kumar, and J. I. Hong, "Towards wearable everyday body-frame tracking using passive rfids," *IMWUT*, vol. 1, pp. 145:1–145:23, 2017. 116

[292] C. Pan, Y. Ohm, J. Wang, M. J. Ford, K. Kumar, S. Kumar, and C. Majidi, "Silver-coated poly (dimethylsiloxane) beads for soft, stretchable, and thermally stable conductive elastomer composites," *ACS applied materials & interfaces*, vol. 11, no. 45, pp. 42 561–42 570, 2019. 116