

©Copyright 2021

Vikram Iyer

Creating the Internet of Biological and Bioinspired Things

Vikram Iyer

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2021

Reading Committee:

Shyam Gollakota, Chair

Joshua R. Smith

Sawyer Fuller

Program Authorized to Offer Degree:
Electrical and Computer Engineering

University of Washington

Abstract

Creating the Internet of Biological and Bioinspired Things

Vikram Iyer

Chair of the Supervisory Committee:
Professor Shyam Gollakota
Computer Science and Engineering

Biological systems can perform incredible feats in comparison to their electronic counterparts. Plants like dandelions can disperse their seeds over a kilometer in the wind, and small insects like bumblebees can see, smell, communicate, and fly around the world, despite their tiny size. Enabling some of these capabilities for the Internet of Things (IoT) and cyber-physical systems would be transformative for applications ranging from large-scale sensor deployments to micro-drones, biological tracking, and robotic implants.

In this dissertation, we will explain how by taking an interdisciplinary approach spanning battery-free systems, wireless networking, sensing, and biology, we can create programmable systems for the Internet of biological and bio-inspired things. We present the first battery-free wireless sensors, inspired by dandelion seeds, that can be dispersed by the wind to automate deployment of large-scale sensor networks. We will then describe how integrating programmable wireless sensors with live animals like bumblebees can enable mobility for IoT devices, and how this technique has been used for real-world applications like tracking invasive “murder” hornets. Finally, we will present an energy-efficient insect-scale steerable vision system inspired by animal head motion that can ride on the back of a live beetle and enable tiny terrestrial robots to see.

Building these systems requires addressing the combined technical challenges of miniaturizing computing, sensing, communication, and power for these devices. The traditional approach to miniaturization has been to develop custom integrated circuits for each of these components. In

contrast, this thesis explains how we can use programmable general purpose computing devices to rapidly develop a variety of novel and innovative miniaturized wireless systems. This approach has the potential to exponentially increase innovation in a range of domains including smart agriculture, biological tracking, microrobots, and implanted devices by empowering anyone with a computer engineering degree to work on miniaturized systems without the need to develop custom silicon.

TABLE OF CONTENTS

	Page
List of Figures	iii
Chapter 1: Introduction	1
1.1 Creating miniaturized wireless systems using general purpose computers	2
1.2 Organization	14
Chapter 2: Wind dispersal of battery-free wireless devices	15
2.1 Introduction	15
2.2 Results	17
2.3 Discussion	44
2.4 Videos	48
Chapter 3: Internet of Biological Things: A Flying Wireless Platform on Live Insects	49
3.1 Introduction	49
3.2 Application Scenarios	53
3.3 Living IoT Design	54
3.4 Evaluation	68
3.5 Related Work	77
3.6 Discussion	80
Chapter 4: Wireless steerable vision for live insects and insect-scale robots	82
4.1 Introduction	82
4.2 Results	87
4.3 Discussion	109
4.4 Videos	113
Chapter 5: Lessons learned & General design principles	115

Chapter 6: Conclusion and Future Directions	121
Bibliography	124

LIST OF FIGURES

Figure Number	Page
1.1 Interdisciplinary approach. Building the Internet of biological and bio-inspired things requires an interdisciplinary approach combining battery-free systems, wireless networking, microfabrication, actuation, and biological systems.	3
1.2 Wind-dispersed battery-free wireless sensors. A) Battery-free wireless sensing device shown next to a dandelion seed. B) Battery-free wireless sensor hovering over a vertical wind tunnel.	5
1.3 Internet of biological things: A flying wireless platform on live insects. A) Our IoT platform with computation, communication, sensing, and localization capabilities attached to a live bumblebee (<i>Bombus vosnesenskii</i>) B) Bumblebee flying freely in a plastic enclosure while carrying a wireless sensor. C) Amplitude of the beamformed, periodic RF waveform transmitted by the access point. D) Multi-antenna access point setup transmitting signals to the bee for self-localization. . . .	8
1.4 Wireless steerable vision for live insects and insect-scale robots. A) Wireless steerable camera attached to a live darkling beetle (<i>Eleodes nigrina</i>) climbing a tree. B) Wireless steerable camera attached to a vibration motor driven terrestrial robot shown next to a US penny for scale.	10
1.5 Radio tracking invasive “murder” hornets. A) Wireless sensor used for tracking including a Bluetooth enabled microcontroller and temperature sensor. B) Attaching a sensor to a live asian giant hornet. C) Tagged hornet resting on an apple tree prior to flight. D) Entomologists from WSDA and USDA searching for tagged hornet with radio receivers. Image credits (B-D): Karla Salp, WSDA.	13
2.1 Wind-dispersed battery-free wireless devices. A) Battery-free wireless sensing device shown next to a dandelion seed. B) Top view of our battery-free wireless sensing device showing its solar cells and drag enhancing structure next to an elm seed. C) Circuit showing the electronic components including sensors that can be attached. D) The circuit is made on a flexible substrate that can be folded. E) Circuit folded into a cube placed next to a dandelion seed.	18

2.2	Effect of outer ring on terminal velocity. A) Terminal velocities of 26% fill structures with different diameters using a 30 mg payload. B) Terminal velocities of 55% fill structures with different diameters with and without an outer ring using a 30 mg payload.	19
2.3	Drag structure payload tests. A) Drag coefficient (C_D) versus Reynolds number (Re) for 75 designs (3 film thicknesses \times 5 diameters \times 5 payloads, $N=5$ trials per design, error bars= $\pm\sigma$). B) Terminal velocity versus payload for the same data shown in A ($N=5$ trials per design, error bars= $\pm\sigma$).	21
2.4	Drag structure tests. A) Designs of different sizes and fill patterns. B) Terminal velocity for designs from C on 12 μm films ($N\geq 5$ trials per design, error bars= $\pm\sigma$). C,D) Probability of upright landing versus fill percentage and projected area ($N=50$ trials per design).	22
2.5	Additional fill patterns. Additional fill patterns used to vary porosity. Changes include concentric circles connecting horizontal spokes extending out from the center with various spacings, as well as additional cuts in the concentric circles between spokes	23
2.6	Fabrication process. Fabrication process used to produce the curved drag enhancing structures. Thin film patterns are first cut with a laser and then the curvature is defined by placing them in a curved mold. A thin layer of double-sided tape is used to define the shape and attach the solar cells.	24
2.7	Backscatter communication. A) A transmitter broadcasts a single frequency tone, the backscatter device toggles an RF switch connected to an antenna, and a receiver decodes the backscattered data. B) Error rate of backscattered packets received on an SX1276 radio at different power levels and bitrates in benchtop tests ($N=500$ packets/point, error bars= $\pm\sigma$). C) Outdoor range test: transmitter and receiver are 20 m apart and backscatter device is moved away equidistant from both. D) In-air packet delivery rate outdoors vs. range at 4.8 kbps with a single transmissions and 1x re-transmission ($N\geq 100$ packets per point, error bars= $\pm\sigma$).	27
2.8	SX1276 receiver on-off keying packet structure. A) Packet structure for the SX1276 receiver operating in on-off keying mode. B) Minimal packet structure used for testing the lower bit rates of 1.5 and 4.8 kbps.	29
2.9	Battery-free wireless sensor. A) Wakeup sequence in continuous operation mode. B,C) Startup time and throughput in continuous mode under outdoor light conditions, $N\geq 5$, error bars= $\pm\sigma$. D) Component weights. E) Sensor weight and power measurements in different operating modes.	33

2.10	Indoor light power harvesting results. A) Cold start time required to charge the storage capacitor from zero to the time at which it can backscatter its first packet; measured using indoor LED lights ($N \geq 5$, error bars= $\pm\sigma$). B) Throughput of the sensor device after startup at different illumination levels ($N \geq 5$, error bars= $\pm\sigma$).	34
2.11	Supporting transmissions from multiple devices. A) Natural variance of startup time for outdoor measurements. B) Recordings from similar light conditions show natural variation reduces collisions. C) Network simulation of up to 1000 nodes each sending 100 packets. Node start and inter packet transmission times are normally distributed based on power harvesting data ($N=100$ packets per node \times 10 repetitions=1000). D) Simulation with normally distributed startup time assuming 1.4 kLux followed by 10 s inter-packets delay ($N=1000$). E) Addition of random delay between packets. F) Simulation with random delays up to 8 packet lengths with random start times assuming 1.4 kLux ($N=1000$).	37
2.12	Drone mechanism to drop the sensors. A) Small quadrotor with the mechanism used to drop devices from altitudes over 22 m outdoors. B) Drop mechanism consisting of a solenoid wirelessly triggered by a Bluetooth transmission. Upon triggering, the supercapacitor keeps the solenoid retracted for a few seconds allowing the sensor to fall freely. While this mechanism was designed to drop a single device for this evaluation, it could be easily be adapted to simultaneously drop multiple devices.	40
2.13	Outdoor wind dispersal. A) Landing locations in outdoor drop tests from increasing heights. B,C) Wind speed versus distance for 5 m and 22 m drop heights. ($\pm\sigma$, $N > 200$ windspeed values) D) Wind speed versus distance for two different structures ($\pm\sigma$, $N > 200$ windspeed values).	41
2.14	Outdoor sensor evaluation. A,B) Outdoor temperature and relative humidity measurements for two sensors over 6 hrs showing successful cold start and correction for sun exposure using the onboard light sensor. C) Outdoor light sensor readings compared to reference solar radiation data. D) Outdoor pressure sensor data compared to a reference sensor. E) Magnetometer readings at 1 Hz from a passing car and experimental setup. F) Accelerometer readings at 1 Hz from a device in a vertical wind tunnel compared to anemometer readings.	42
2.15	Sensor calibration setup using a weather station. A) Picture of weather station used for reference temperature, humidity, solar radiation, and pressure measurements and test sensor node placement. B) Correlation of small, onboard light sensor with reference solar radiation. C) Correlation of small, onboard temperature sensor (HDC2010) with reference. D) Correlation of onboard temperature sensor (BME280) with reference.	43

2.16	Prototype Full Duplex Cancellation of narrowband signals. A) Full duplex cancellation block diagram. B) Results of applying full duplex cancellation to reduce transmitter interference for backscatter.	45
3.1	Living IoT platform on a bumblebee. (Left) Live bumblebee (<i>Bombus sitkensis</i>) carrying our platform which includes computing, sensing, wireless communication, and self-localization; (Right) Bee flying while carrying our platform.	50
3.2	Living IoT architecture. Insect-borne sensor packages can self-locate using transmissions from APs and collect location-dependent data using on-board sensors. Data is stored in onboard memory as the bee flies and data are uploaded to the AP when the bee is back at the hive.	53
3.3	Living IoT electronics. Complete electronics package including an antenna, envelope detector, backscatter transmitter, and sensor shown on a US penny for scale (left). Custom miniature antenna (right) also shown on a US penny.	55
3.4	Envelope detector performance. Miniaturized envelope detector included onboard Living IoT (left). A plot evaluating its performance compared to a similar design fabricated on a standard thickness PCB with larger component demonstrates miniaturization does not compromise performance.	57
3.5	Envelope detector circuit diagram. Circuit diagram for our envelope detector design. The circuit consists first of tuning elements for the antenna followed by a three stage voltage multiplier that extracts the envelope of the incoming RF signal.	59
3.6	Self-localization architecture. Our self-localization system uses multiple antennas at the AP and transmits signals with periodically increasing phase delays from the antennas. The result beamforms the output over a range of angles such that the maximum amplitude received on the bee indicates the angle to the AP.	60
3.7	Multipath simulation. Simulation results showing the effect of multipath. Different multipath conditions are shown as a ratio to the strongest path $R = \frac{\sum_{i=2}^M a_i}{a_1}$. Access points with 2–5 transmit antennas are simulated and the resulting angular error for different multipath conditions is shown for each.	62
3.8	Living IoT backscatter hardware. A block diagram of the backscatter hardware used in Living IoT is shown on the left. The corresponding light-weight hardware prototype with these components is shown on the right.	66
3.9	Stationary self-localization angular accuracy. The accuracy of the localization system is measured for multiple antennas at distances of up to 80 m in an open field (left). The angular error versus range is plotted for each antenna configuration (right).	67
3.10	Mobile 2D localization angular accuracy. The receiver is attached to a drone flying at different speeds to evaluate the impact of mobility. The resulting angular errors versus range for each flight speed at each AP location are shown.	68

3.11	Localization accuracy on farm. The 2D localization accuracy is evaluated by measuring signals at 10 points across a small farm (left) with two APs. A CDF of the localization accuracy is shown to the right.	70
3.12	Localization accuracy on bee. The 2D localization accuracy is evaluated by measuring signals from an envelope detector attached to a live bee kept in a plastic enclosure at 10 points across a small farm (left) with two APs. A CDF of the localization accuracy is shown to the right.	71
3.13	Effect of flapping wings on envelope detector. A small envelope detector is shown attached to an insect-scale flapping wing robot (left). The effect on the envelope detector at different flapping frequencies is evaluated (right).	73
3.14	Backscatter performance. The uncoded BER on a static and mobile bee is shown (left). These values are low and comparable to prior backscatter designs [147] and the bee can upload data when it is back at the hive. The impact of the small antenna performance is shown to the right in a comparison of BER versus range.	74
3.15	Recharging batteries using RF harvesting. A patch antenna transmitting power is placed next to a beehive contained within a cardboard box (left). The charging current and time are shown to the right for different RF transmit powers.	75
3.16	Living IoT sensors. Close up images of the sensors included in Living IoT. A combined temperature and humidity sensor is shown (left) along with an ambient light sensor (right).	76
3.17	Solar harvesting for battery-free platform. A small solar cell is shown above a bee (left). The power produced by this solar cell in different light conditions is shown to the right.	79
4.1	Steerable wireless camera on a beetle and a small robot. A) Wirelessly steerable camera system attached to a live beetle. B) A wireless, power-autonomous terrestrial robot with steerable vision. The camera streams video to a smartphone, which can command the robot to move and pan camera. C) Exploded view showing all components of the system including the Bluetooth, camera and optics, robotic head, high voltage electronics, and battery. D) Diagram showing the components of the steering mechanism. E,F) Close-up diagram of hinge motion as the camera pans.	86

4.2	Smartphone interface. A) Image showing the smartphone GUI for recording video and controlling the camera angle on a Samsung Galaxy S9. The GUI includes a menu for selecting the device to connect to, an image viewer, and controls to operate the camera in a single image capture mode, continuous streaming mode, or accelerometer-triggered capture mode. The GUI also includes data logging and debugging features and the following menus for setting parameters: B) menu for selecting between four preset image resolutions, C) menu for setting the robotic head angle, and D) menu for selecting Bluetooth bit rate.	87
4.3	Full system block diagram. A) Block diagram of the backpack carried by the insect. B) Block diagram of the robot. Both of these systems are composed of a Bluetooth enabled microcontroller (MCU), robotic head and power electronics, camera, motion sensors and power system. The robot includes two vibration motors and two buck converters to control their voltage and speed via Bluetooth. A voltage regulator supplies power to the MCU and boost converter from the LiPo battery. The MCU both provides a wireless interface and control signals for sensors and vibration motors. The MCU also controls the boost converter and switches to produce the output voltages required to move the piezo.	89
4.4	Evaluating the wireless camera and arm performance. A) Labeled diagram showing peizo actuator operation. B) Low voltage input current versus high voltage output generated by the boost converter. C) Sample 160×120 images showing the performance of the camera using a 1.5 mm diameter 1 mm focal length lens (Edmund Optics 43394). D) Sample 160×120 images showing the performance of the camera using a 3.8 mm diameter 2.33 mm focal length lens (Panasonic EYL-GUDM128). E) Frame rate versus line-of-sight range from the insect. Error bars indicate mean $\pm 1\sigma$ (N=10 frames); the rate remains constant until the sensitivity limit of the wireless link, and then begins to degrade. F) Battery life when continuously streaming 160×120 images at different rates, with and without robotic head motion, and different batteries. G) Weight breakdown of the Bluetooth, camera and robotic head, as well as each component's percentage of total weight.	91
4.5	Boost converter circuit schematic and drive signal waveforms. A) Boost converter and switch network used to drive the piezo arm. Power is provided from a LiPo battery and control signals are generated by the MCU. B) Boost converter output V_{Boost} is connected to the top piezo and V_{PZT} , is connected to the center carbon fiber layer of the actuator. Voltage is varied in 5 steps. The center plot shows the switch control voltages. The PNP transistor Q_3 first shorts V_{PZT} to V_{Boost} , for positive angles. Both transistors Q_2 and Q_3 are then off for 0° angle. NPN Q_2 is turned ON shorting V_{PZT} to V_{gnd} , while V_{Boost} increases for negative angles. Output angles are shown in the bottom plot.	92

4.6	Comparison table for image sensors. Comparison of commercial image sensor specifications including their resolution, power consumption and other characteristics.	95
4.7	Sample images comparing lenses. Sample images comparing 1.5 mm diameter and 3.8 mm diameter plano-convex lenses imaging a standard test pattern (1951 USAF resolution test chart, MIL-STD-150A). The test pattern was printed and held at the specified distance away from the camera. Images were captured indoors at a resolution of 160×120 at illuminance levels 200 lux and 5 lux. The comparable image quality at 5 lux demonstrates the system can operate in low light conditions.	98
4.8	Frame rate versus distance and resolution using Bluetooth. A) Frame rate measured versus range between the small Bluetooth chip and smartphone at two data rates. Error bars indicate mean $\pm 1\sigma$ (N=50 frames). B) Maximum frame rate at different resolutions.	99
4.9	Field evaluation of accelerometer triggered camera. A) Close-up image of the wireless camera without a micro-robotic arm on the back of a death feigning beetle. B) Experiment site with a grove of trees, a dry stream bed, and gravel paths. The beetle walked freely in 4 different locations. C) Images from the camera showing a person walking. D) Beetle motions detected by the accelerometer per minute over the 363 min experiment. The inactivity explains the improvement in battery life over continuous streaming. E) Top 25 intervals between accelerometer triggers when the system is in sleep mode.	102
4.10	Field evaluation of micro-robotic arm on darkling beetle capturing panorama images. A) Wireless camera system with the micro-robotic arm attached to a darkling beetle. The camera was set to capture five images as it rotated in 15° steps to capture a panorama. B) Aerial view of the experiment site in a gravel parking lot showing buses and trucks in the Southwest corner. C) Panorama showing the trucks and buses composed of 5 images captured by the insect-mounted camera while rotating 60°. D) Beetle motions detected by the accelerometer per minute. E) Top 25 intervals of beetle inactivity.	104
4.11	Comparison with prior power autonomous robots. Comparison of prior small, power autonomous robots including their size scale, mass, power consumption and wireless interface.	105

4.12	Evaluation of the insect scale wireless robot. A,B) Close-up images of the robot with vibration motors, three legs, a battery, Bluetooth chip, camera and robotic head. C) Effect of carrying payload greater than the camera system and the battery on robot speed, while running the vibration motors at a constant power. Error bars indicate mean $\pm 1\sigma$ (N=5 trials). D) Robot power consumption for different motion types and speeds. E) Energy consumption for different angles when using the motors versus head to turn the camera. F) Battery life for different robot speeds without video streaming, with 1 fps video streaming, and 1 fps video streaming while panning the camera.	106
4.13	Distortion during motion. A) Image showing the bottom corner of a Rubik's cube from the robot's perspective. B,C) Images captured during robot motion. The side-to-side distortion results from the rolling shutter of the camera which samples data line by line while the vibration motors shake the robot.	109
4.14	Navigation and focusing on another moving robot. A) Overhead view of robot's path when a human operator uses the camera to navigate. B) Sample images from the robot's camera during navigation. C) Side view showing our wireless robot is stationary with the camera looking at another robot marked with an arrow moving left. Background shows numbers incrementing from right to left to show motion. D) Top view of scene shows the moving robot going across the field of view of the camera. The camera rotates in discrete steps to maintain focus on the moving robot. E) Images captured at each angle.	110

ACKNOWLEDGMENTS

I would like to thank all of my collaborators across the various departments and labs at UW and other organizations who helped make this interdisciplinary work possible as well as all of my family and friends who have supported and encouraged me.

Chapter 1

INTRODUCTION

Living organisms can perform incredible engineering feats. A bumblebee can store enough energy to fly for hours, it has vision to navigate, it can communicate and perform sensing. All of these capabilities are contained in a small (sub-centimeter), power-efficient, and lightweight (sub-gram) package. Even simpler organisms like dandelions, which cannot move, have evolved mechanisms to disperse their seeds for over a kilometer in the right conditions [19].

By comparison, there remains a large gap between the capabilities of these biological systems and state-of-the-art Internet of Things (IoT) and embedded systems. Many of IoT systems are larger, heavier, and lack mobility. My thesis asks the following question: can we create tiny battery-free wireless devices that can move around the world, and float in the air similar to dandelion seeds? This technology could enable automated deployment of large scale sensor networks with hundreds of devices in remote, hard-to-reach areas like forests and glaciers. Further, if we can create insect-scale wireless sensors, then we can attach them to tiny insects like bees, beetles and murder hornets and transform biological research for insect-scale organisms. In fact, if we could go a step further and add actuation, we can start building power-autonomous insect-scale wireless robots to explore small spaces like inside pipes or under heavy equipment.

This thesis demonstrates that it is possible to close the gap between biology and computing systems and create the Internet of biological and bio-inspired things. We will show how to create bio-inspired battery-free wireless sensors that can disperse in wind across a large area similar to dandelion seeds. We will then create the Internet of biological things that can be used for wildlife tracking of small animals and insects. Finally, by integrating wireless systems with actuators, we will demonstrate the smallest power-autonomous terrestrial robot with steerable vision.

The common enabling technology across all these systems is a set of techniques that allows

us to build tiny wireless systems using programmable, general purpose computing devices. This approach of using general purpose computing platforms enables rapid prototyping and innovation that was not previously possible at this scale.

1.1 Creating miniaturized wireless systems using general purpose computers

Building end-to-end wireless sensing devices involves addressing numerous technical challenges. They require digital processors to do computing, a wireless link like a radio, sensors to measure the quantity of interest, and a power source. The traditional approach to miniaturizing such systems has been to build custom integrated circuits (ICs). The heterogeneous process requirements of digital circuits like memory and processors, analog sensor interface circuits and radios, sensors with micro-electromechanical (MEMs) components, and power sources make it difficult to integrate all of these components onto a single monolithic IC, and necessitates multiple custom chips.

This approach however introduces a number of practical challenges for researchers which has throttled the pace of innovation. First, custom chips have high fabrication costs and times often leading to long, multi-year design cycles. Second, application specific integrated circuits (ASICs) are often designed for a specific purpose and difficult to adapt to new scenarios or technological changes. Third, the expertise required to design a chip takes years to develop making this domain inaccessible to researchers in the broader computing and engineering community. Even within circuits, building complete systems requires expertise from multiple sub-specialties across analog and digital design, and even MEMs and materials science to design sensors and batteries.

These challenges have thus far limited development of miniaturized wireless sensing devices to only a few research groups, or multiple groups that have had to join together over a few years to build a working system. While the concept of “smart dust” has existed for over 20 years [67], only a handful of end-to-end systems have been developed in this time. This is similar to the state of wireless research before the creation of software radios, where the need to design custom radios was holding back a whole community of researchers.

The key technical insight of this thesis is to instead use programmable, general purpose computing chips like microcontrollers to build millimeter scale wireless sensors. For example, the

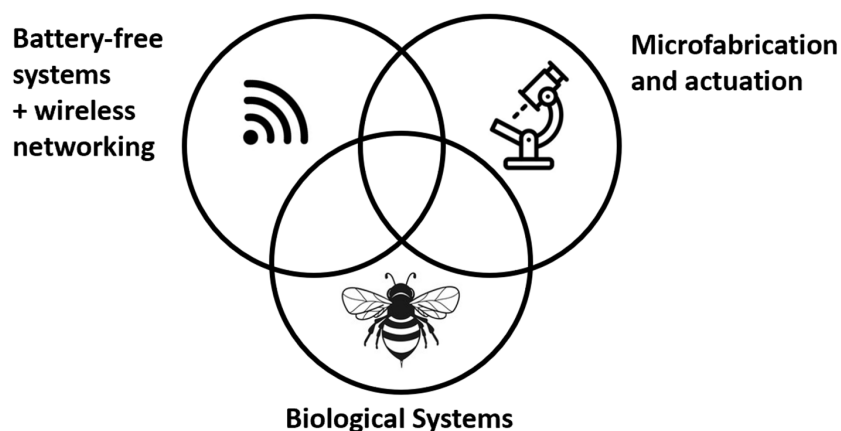


Figure 1.1: Interdisciplinary approach. Building the Internet of biological and bio-inspired things requires an interdisciplinary approach combining battery-free systems, wireless networking, microfabrication, actuation, and biological systems.

smallest commercial microcontroller used in this work is only 1.5×1.4 mm. While it has limited computation and lacks common hardware interfaces for sensors, its programmability allows for a software implementations of the common Inter-Integrated Circuit (I²C) protocol to connect to sensors. While the chip lacks a radio, later on this thesis will explain how to use a technique called backscatter to decouple the radio onto another device, and just implement the digital operations of encoding the packet on a microcontroller. The power savings from these techniques also enables the use of harvested energy from sources like a small solar cells to power the system.

The true power of this approach is that by eliminating the need to design custom silicon, it empowers a much broader community of researchers across computer science and engineering to actually build tiny wireless devices. Reducing this barrier has the potential to exponentially increase the pace of innovation across a range of domains including smart agriculture, biological tracking, microrobots, and implanted devices.

Building on these techniques, we can start creating the Internet of biological and bioinspired things that takes an interdisciplinary approach spanning battery-free systems and wireless networking, microfabrication and actuation to build devices that can move, and an understanding of biology to integrate these electronics with living organisms. Taking such an interdisciplinary approach illustrated in Fig 1.1 allows us to create entirely new inventions and build innovative systems that push the boundaries of what is possible with technology. This thesis presents three such systems:

- *Bio-inspired wind-dispersal of wireless sensors.* The lightest-weight system presented in this dissertation demonstrates how we can learn from something as simple as a dandelion seed to create the first system to disperse battery-free sensors in the wind across a large area. By minimizing weight to tens of milligrams and emulating the drag enhancing structures of dandelion seeds, we demonstrate these wireless sensors can travel up to 100 m in the wind. The structures also enable the sensors to land upright to harvest solar power, and their designs can be varied to mimic variation between individual seeds to prevent clustering. We demonstrate outdoor experiments showing these wireless sensors that can measure temperature, humidity, light, pressure, acceleration, and magnetic fields and enable various sensing applications.
- *Internet of Biological Things.* In our second system, we piggyback wireless sensors on live insects like bumblebees to enable mobility for IoT devices that far exceeds the limits of drone battery life. Insects like bees can support larger payloads than the wind dispersal system described above allowing for the use of small batteries but still face significant size, weight, and power miniaturization challenges. Our platform combines wireless communication, sensing, and a novel low power self-localization technique to perform opportunistic sensing as bees fly around environments like farms. This Living IoT platform has been integrated with organisms like bees, beetles and hornets and has helped in the discovery of the first “murder” hornet nest in the United States, in collaboration with WA department of agriculture.
- *Bio-inspired microrobot with wireless vision.* The final system presented in this work investigates the next larger class of sensing devices that have higher computational and bandwidth requirements, i.e., cameras. We take a bio-inspired approach to integrating high bandwidth sensors like cameras onto these tiny wireless systems. Specifically, we present the first wireless steerable camera that emulates head motion in animals to enable wider field of view at low power and size. This system enables integration of cameras onto small insects for novel behavioral studies and for the first time extends vision to power autonomous microrobots.

Below, we summarize the high level ideas and contributions of each of these systems.




Figure 1.2: Wind-dispersed battery-free wireless sensors. A) Battery-free wireless sensing device shown next to a dandelion seed. B) Battery-free wireless sensor hovering over a vertical wind tunnel.

1.1.1 Bio-inspired wind dispersal of battery-free wireless devices

Plants cover a significant fraction of the earth's land mass despite most species having limited to no mobility. To transport their propagules, many plants have evolved mechanisms to disperse their seeds using wind. A dandelion seed, for example, has a bristly filament structure that decreases its terminal velocity and helps orient the seed as it wafts to the ground. Inspired by these seed structures, we designed the first system to achieve wind dispersal of battery-free wireless sensing devices to enable automated deployment of sensor networks. This could enable numerous applications such as deploying large-scale sensor networks for smart agriculture, environmental sensing networks in remote locations like glaciers, and sensors to detect air pollution or forest fires.

Our millimeter-scale devices shown in Fig 1.2A,B weigh 30 mg and are designed on a flexible substrate using programmable, off-the-shelf parts to enable scalability and flexibility for a variety of sensing and computing applications. The system is powered using lightweight solar cells and an energy harvesting circuit robust to low and variable light conditions. Power harvesting eliminates the need for heavy batteries and allows for designing a system with the absolute smallest size and weight. The wireless sensor uses a backscatter communication link for data transmission. Backscatter is a communication technique that allows a device to transmit data using reflected radio frequency (RF) signals. In addition to significantly reducing the power required to communicate, backscatter can be used to enable communication without an onboard radio. This allows us

to build a complete wireless sensing device using only a microcontroller that reads a sensor, and then toggles an RF switch connected to antenna to produce a modulated reflection that communicates the data. We demonstrate end-to-end outdoor experiments showing our sensors can measure temperature, humidity, light, pressure, acceleration, and magnetic fields.

To achieve wide-area dispersal and upright landing necessary for solar power harvesting, we developed dandelion inspired thin film drag-enhancing porous structures. As a seed falls, the force of gravity pulls it down while air resistance produces an upward drag force. At the point where these two forces cancel each other out, it experiences zero net force and falls with a constant terminal velocity. Intuitively, if a seed falls slower with a lower terminal velocity it can be carried to longer distances in the wind. A detailed discussion of terminal velocity is presented in chapter 2, but at a high level seeds with low mass and large area fall slowest. We perform detailed empirical characterization of the design space by testing 75 different structures with varying diameters, material thicknesses, and payloads. We show that our designs can achieve a terminal velocity of 0.87 ± 0.02 m/s and aerodynamic stability with over 95% probability of upright landing to make sure the device can harvest solar power. Our results in outdoor environments demonstrate that these devices can travel 50-100 m in a gentle to moderate breeze. Finally, in natural systems, variance in individual seed morphology causes some seeds to fall closer and others to travel farther. We adopt a similar approach and show how we can modulate the structures' porosity and diameter to achieve dispersal variation across devices.

Relation to Prior Work: Our wireless sensors themselves differ from prior work in that they are battery and radio free, making them ultra-lightweight (< 70 mg needed for wind dispersal). While prior designs have demonstrated millimeter scale wireless computing and sensing platforms using custom ICs [64, 21, 38] they do not consider strategies for efficient wind dispersal. Our method presents a highly accessible platform for researchers with diverse backgrounds to build on and rapidly prototype new applications at millimeter-scale without the significant time, expense, and expertise required to design custom silicon. Prior attempts at automated sensor network dispersal are orders of magnitude larger and require significantly heavier structures to produce drag that increases overall size and weight. By comparison, our approach is small and light enough that 1000s

of devices could be dropped from even small commercial quadrotors capable of carrying only 50 g. Additionally, our designs can be cut from a single sheet of material and can be rapidly fabricated or even integrated into flexible circuit fabrication processes. We also take inspiration from natural variation in seed morphology to show how these designs can be varied to achieve different dispersion distances. Combined, these techniques allow for orders of magnitude increase in the number of sensors that can be dispersed and demonstrate a method to achieve greater dispersion coverage.

1.1.2 *The Internet of Biological Things: A Flying Wireless Platform on Live Insects*

Sensor networks with devices capable of moving could enable applications ranging from precision agriculture to environmental sensing; however using mechanical drones to move sensors severely limits operation time since their flight times are limited by the energy density of current battery technology. We explore an alternative, biology-based solution: integrate sensing, computing and communication functionalities onto live flying insects like bumblebees to create a mobile IoT platform. Such an approach takes advantage of these tiny, highly efficient biological insects which are ubiquitous in many outdoor ecosystems, to provide mobility for free. Fig 1.3A,B show our platform on the back of a live bumblebee (*Bombus vosnesenskii*).

Doing this however requires addressing key technical challenges of power, size, weight and self-localization in order for the insects to perform location-dependent sensing operations as they carry our IoT payload through the environment. Bumblebees can carry approximately 100 mg of payload which allows for the use of small batteries unlike the wind dispersal system described above. Using a battery that weighs 70 mg however still places stringent constraints on the remaining electronics. Because our system does not control the flight of the bees, we need a method to locate them as they fly to correlate sensor measurements with locations. This is challenging because technologies such as the global positioning system (GPS) consume significantly more size, weight, and power than our system can afford. Instead we design and implement an ultra-lightweight and low power self-localization system. The high level idea is similar to backscatter, where the power expensive and complex components are offloaded to an access point. A multi-antenna access point transmits signals which repeatedly cycle through phase offsets between its antennas as shown in

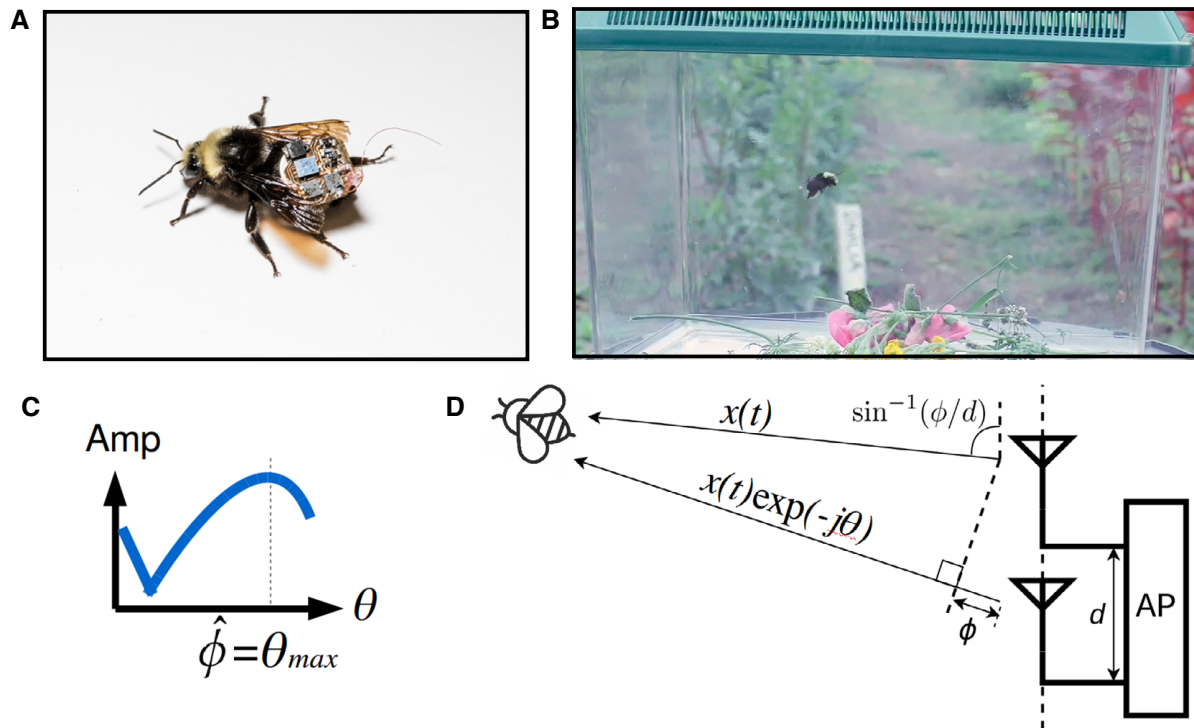


Figure 1.3: Internet of biological things: A flying wireless platform on live insects. **A)** Our IoT platform with computation, communication, sensing, and localization capabilities attached to a live bumblebee (*Bombus vosnesenskii*) **B)** Bumblebee flying freely in a plastic enclosure while carrying a wireless sensor. **C)** Amplitude of the beamformed, periodic RF waveform transmitted by the access point. **D)** Multi-antenna access point setup transmitting signals to the bee for self-localization.

Fig 1.3C,D. When these signals constructively interfere in space, they will produce a higher amplitude signal. On the bee, a passive envelope detector circuit can be used to extract the amplitude of the RF signal. The highest amplitude in this beamforming cycle indicates the angle of the bee with respect to the access point. By combining angles from two access points at known locations, we can determine the location of the bee.

We develop and deploy our platform on bumblebees which includes backscatter communication, low-power self-localization hardware, sensors, and a battery. We show that the system can periodically take sensor and location measurements as the bee flies and store them in onboard memory. Then, data can be uploaded when the bee returns to its nest close to an access point. We show that our platform is capable of sensing, backscattering data at 1 kbps when the insects are back at the hive, and localizing itself up to distances of 80 m from the access points, all within a

total weight budget of 102 mg.

Relation to Prior Work: Researchers have attempted multiple methods to track bees in the past including by attaching laser activated transponders [139], radio-frequency identification (RFID) tags [52, 109], and high power radar systems [113, 27]. The RFID and laser based approaches are limited to centimeter ranges, while the radar based approach is difficult to scale to large number of bees. Additionally, none of these works have onboard general-purpose computing or sensing capabilities. In contrast to this prior work, our design demonstrates the first low-power self-localization technique for flying insects, a general purpose platform that enables computing, communication and sensing on flying insects and, demonstrates for the first time that insects such as bees can be used to carry general purpose sensors in lieu of drones.

1.1.3 Wireless steerable vision for live insects and insect-scale robots

Vision serves as an essential sensory input for insects but consumes significant energy resources. The cost to support sensitive photoreceptors has led many insects to develop high visual acuity in only small retinal regions and evolve instead to move their visual systems independent of their bodies through head motion. By understanding the tradeoffs made by insect vision systems in nature, we can more intelligently design vision for insect-scale robotics in a way that balances energy, computation and mass.

A naïve approach to building a vision system for small robots would be to leverage advances in miniaturization made for smartphone cameras. This would seem an intuitive choice due to their small size and megapixel resolution [156]; however, the processing and energy requirements needed to support these cameras necessitate powerful processors and prohibitively large batteries. We instead look to biology and explore the trade-offs that evolution has made in the visual systems of insects to inform the design of vision for insect-scale robotics.

Visual systems of insects like flies have lower resolution than large animals like humans but can consume 2.5 to 13% of body mass [8]. Furthermore, supporting these large organs represents a substantial energetic cost—the retina of a blowfly alone consumes 8% of its resting metabolism [100]. The energy cost to support these sensitive photoreceptors has led many insects to develop high

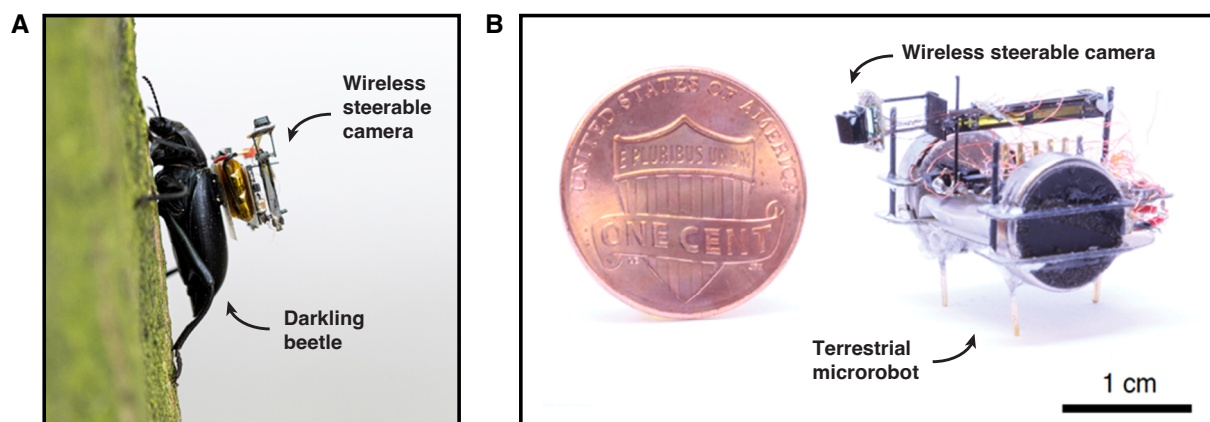


Figure 1.4: Wireless steerable vision for live insects and insect-scale robots. **A)** Wireless steerable camera attached to a live darkling beetle (*Eleodes nigrina*) climbing a tree. **B)** Wireless steerable camera attached to a vibration motor driven terrestrial robot shown next to a US penny for scale.

visual acuity in only small retinal regions [48, 152, 49]. To compensate for a smaller visual field, insects have evolved to move their visual systems independent of their bodies through head motion to expand their field of view by dynamically scanning or to maintain focus on moving objects (e.g., prey or potential mates). In contrast to moving the whole body, this adaptation allows these animals to gather more visual information in an energy-efficient manner [77]. In addition, this added degree of freedom is used by some insects to guide steering as the first part of a motor program that can precede body turns [15], infer depth or motion information [76, 152, 49], or orient their gaze in a direction independent of their movement direction [14].

Inspired by this, we design a fully wireless, power autonomous, mechanically steerable vision system that imitates head motion in a form factor small enough to mount on the back of a live beetle or a similarly sized terrestrial robot shown in Fig 1.4A,B. The wireless vision system consists of a programmable microcontroller which implements a custom software interface to read data from an ultra-low power CMOS image sensor with a lightweight carbon fiber lens assembly. The microcontroller then transmits the data to a smartphone or other Bluetooth compatible device using its onboard radio. Just like insects, a larger image sensor would consume additional size, weight and power as well as wireless bandwidth to transmit the data. We improve the camera's field of

view by mounting it at the end of an actuator that can pan it to the right or left. The actuator leverages the piezoelectric effect, a material property that causes mechanical deformation in a material in response to an applied electric field. This property is used to produce a bending motion, which is coupled to a hinge mechanism that pans the camera. A detailed description of the actuator design is presented in chapter 4. The high voltage (200 V) signals required to drive the piezo are generated onboard using a lightweight boost converter controlled by the microcontroller.

Our electronics and actuator weigh 248 mg and can steer the camera over 60° based on commands from a smartphone. The camera streams ‘first person’ 160×120 monochrome video at 1–5 frames per second (fps) to a Bluetooth radio from up to 120 m away. We mount this vision system on two species of freely walking live beetles, demonstrating that triggering image capture using an on-board accelerometer achieves operational times of up to 6 hours with a 10 mAh battery. We also built a small, terrestrial robot (1.6×2 cm) that can move at up to 3.5 cm/s, support vision and operate for 63–260 mins. Our results demonstrate that steerable vision can enable object tracking and wide angle views for 26–84 times lower energy than moving the whole robot, which is similar in behavior to biological systems.

Relation to Prior Work: Existing wireless camera solutions require significant amounts of power and weight. For small robots, a heavy payload requires more energy to maintain speed, reducing their operational time. Similarly, the addition of a heavy payload can limit an insect’s ability to move. Due to these constraints prior insect mounted sensing and control systems have not previously demonstrated wireless vision [28, 119, 145, 60]. Furthermore, no prior insect-scale terrestrial robots have included wireless cameras. Larger aerial robots like the 10 cm Delfly micro which weighs 3 g have demonstrated wireless cameras, but their camera sub-system is too power intensive and heavy to be carried by smaller vehicles or insects [30]. Similarly, the small vision systems developed for wireless capsule endoscopy robots require hundreds of milliwatts which make them incompatible with small and lightweight batteries [22, 31, 39] (details in Chapter 4). Our approach of using commercially available chips has the potential to rapidly scale and be adapted to a wide variety of uses across the microrobotics community.

1.1.4 *Impact of this thesis*

These works demonstrate a flexible method for developing millimeter scale wireless sensing systems using general purpose computing chips. As stated above, the most powerful aspect of this approach is that by eliminating the need to design custom silicon, it empowers a broad set of researchers across the broad computer science and engineering community to start developing miniaturized wireless and sensing devices. Similar to the impact of software radios, this approach has the potential to exponentially increase the pace of innovation across a range of domains including smart agriculture, biological tracking, and implanted devices.

Moreover, this approach could be transformative for miniaturized robotic systems as well. Larger robots routinely use vision for all kinds of tasks, but for robots the size of a penny the power and weight constraints make traditional cameras infeasible. This dissertation demonstrates for the first time that this new tiny robots can “see” and our open-source software and hardware serves as a primitive for building all kinds of functionality like navigation, mapping, object detection and more.

“Murder” hornet tracking. These technologies have already been adapted for real world applications an example of which is discussed in detail below. In 2020, following the discovery of the first asian giant hornets (*Vespa mandarinia* or “murder hornets”) in the US by the Washington Department of Agriculture (WSDA), the technologies in this dissertation were adapted to track live hornets. This species of hornet is a large predatory insect native to parts of Asia that preys on smaller insects such as bees and wasps [151]. Because it is invasive in North America and faces no competition or natural predators, they have the potential to quickly spread across the Pacific Northwest and eventually across North America and devastate native bee populations that have not co-evolved natural defenses [101].

Unlike other insects like gypsy moths which can be targeted by pesticides, this approach has significant potential to harm native bee and wasp populations as well. The traditional approach to locating an asian giant hornets nest has been to set up a large number of traps and progressively hone in on the location, or by catching an individual and watching the direction it flies. This is

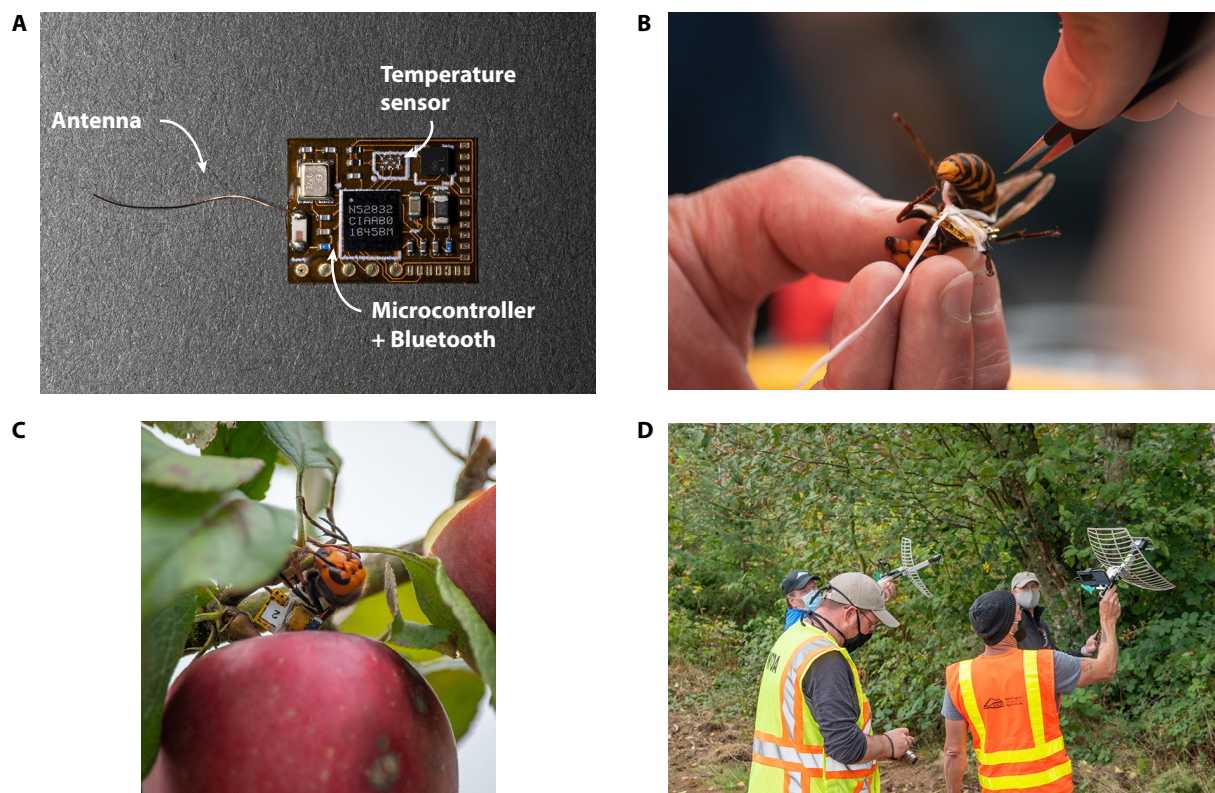


Figure 1.5: Radio tracking invasive “murder” hornets. **A)** Wireless sensor used for tracking including a Bluetooth enabled microcontroller and temperature sensor. **B)** Attaching a sensor to a live asian giant hornet. **C)** Tagged hornet resting on an apple tree prior to flight. **D)** Entomologists from WSDA and USDA searching for tagged hornet with radio receivers. Image credits (B-D): Karla Salp, WSDA.

challenging to control a small population where there may be only a few established colonies in a large geographic area.

To address this, a radio tracking system was developed and deployed on live hornets to follow them back to their nest in collaboration with WSDA. The system consists of a small microcontroller integrated with a Bluetooth transmitter adapted from my work [58]. The microcontroller shown in Fig 1.5A periodically samples a temperature sensor which can be used to detect when the hornet has entered its nest (typically warmer than ambient temperatures due to the large number of animals). The system was programmed to sample the temperature sensor twice per second and transmit the data as a Bluetooth advertisement. The data and signal strength were received on a directional

15 dBi parabolic grid antenna connected to a Bluetooth receiver and data was displayed on a smartphone. By using Bluetooth, additional community members were able to aid in the search using commodity smartphones.

The hornets were cooled on ice for sedation. On the first tracking attempt cyanoacrylate glue was used to attach the tag, however the animal was highly active even after sedation and glue flowed onto its wing during attachment rendering it unable to fly. Subsequent attempts instead attached the tags with a small loop of dental floss around the petiole of the animal (connection between the abdomen and the thorax) where the body narrows as shown in Fig 1.5B. The insect was placed on an apple tree Fig 1.5C and quickly flew $\approx 10 - 20$ m into the air, made a characteristic orientation loop and flew straight in the direction of its nest. Entomologists from WSDA and USDA tracking the hornets are shown in Fig 1.5D. The third live tracking attempt resulted in discovery of the location and destruction of the first *V. mandarinia* nest in the US. Multiple studies building on this approach are currently under development including studies of energetics in hummingbirds and a platform for studying bioacoustics for small birds.

1.2 Organization

The remainder of this dissertation is organized as follows. Chapter 2 presents a detailed discussion of dandelion inspired battery-free wireless sensors. This includes experimental characterization and design methodology for the drag structures, backscatter communication techniques implemented on a small microcontroller, and evaluation of sensors. Chapter 3 presents Living IoT, a wireless platform complete with communication, computation, sensing, and localization small enough to fly on the back of a live bumblebee. Further details of the self-localization system and characterization are presented. Chapter 4 presents a wireless steerable camera small enough to ride on the back of a live beetle as well as the smallest, power-autonomous terrestrial robot with vision. This section explains the camera interface, actuator and boost converter design, and field experiments with insects and robots. Chapter 5 discusses the lessons learned from these works in terms of general principles for building miniaturized wireless sensors with programmable computing platforms. Chapter 6 concludes and discusses a vision for future directions of this research.

Chapter 2

WIND DISPERSAL OF BATTERY-FREE WIRELESS DEVICES

2.1 Introduction

Plants have evolved various mechanisms to use wind for seed dispersal over a wide area [26, 82, 45, 46] including creating light-weight diaspores with plumose or comose structures that act as drag-enhancing parachutes [5, 19]. Asteraceae plants, like the common dandelion, produce plumed seeds containing a pappus, which is a bundle of bristly filaments [26]. The pappus increases the drag experienced by the diaspores, decreasing their terminal velocity. This helps prolong their descent, and allows greater opportunity for lateral wind currents to disperse them outward from their release point. Empirical studies have shown that dispersal distance is inversely proportional to terminal velocity [7]; in dry, windy and warm conditions, dandelion seeds can travel as far as one kilometer [26, 45]. The bristly filament structures also enable the ability to flip in mid-air to achieve a stable descent in a defined, upright orientation [127]. Successful seed dispersal also requires prevention of clustering so that not all seeds land in the same area. While fluctuation in wind patterns allows for some level of dispersal variation, the unique morphology of individual seeds provides an additional factor that varies travel distance [121].

These properties are highly attractive for deployment of sensor networks which have a similar goal: disperse large numbers of low-cost sensors over a broad geographic area to enable wide-area sensing. While numerous communication protocols have been developed for sensor networks, there is recent interest in the problem of automating widespread sensor deployment [98, 108, 58]. Since manual deployment is difficult to scale, passive aerial wind dispersal presents an attractive alternative. Prior attempts of automated sensor network dispersal however have orders of magnitude larger and heavier aerodynamic structures [98, 108] or use batteries for power [98, 108, 58]. In this work we investigate designs inspired by plumed seeds to develop battery-free wireless sensing

devices that disperse in the wind.

Many wind-dispersed seeds with plume structures seen in nature have evolved to be small and light-weight, as these are more efficient for producing drag with small payloads than winged seeds [19]. For example, a dandelion seed measures roughly 14 mm and weighs around 1 mg [26]. Because terminal velocity scales with mass, the combination of milligram-scale weights and high drag coefficient makes long-distance wind dispersal a viable strategy. Emulating this behavior with wireless sensors however poses four key challenges. First, this requires designing and fabricating a light-weight drag-enhancing structure similar to the pappus of a dandelion seed. Second, this requires significant miniaturization of all the electronic components, including a computing platform, wireless communication module, sensor, and power source. Third, batteries present a significant scaling limitation due to the limits of current battery technologies. For example, the smallest commercial batteries used in recent miniaturized wireless sensing platforms weigh 68 mg [60, 58]. While custom solutions could be smaller, they would further reduce operational lifetime due to the limited energy density of lithium batteries. Solar power harvesting designs eliminate the need for batteries but introduce the challenge of limited and intermittent energy availability. Solar power harvesting also requires the cells to face the sun. This precludes the use of traditional parachutes that could fold or get tangled to occlude the solar cells, and adds aerodynamic stability requirements to the drag structure so it can passively re-orient itself and land upright. Fourth, to prevent clustering of the sensors after landing, we need to introduce varied morphology into individual sensors to ensure their widespread dispersion across the target area.

In this chapter, we address extreme size, weight and power (SWaP) challenges outlined above and demonstrate wind dispersal of battery-free wireless sensors (Fig 2.1A-E, Video 2.1). Specifically, we make the following key contributions: 1) We design an array of 2-D drag-enhancing structures that can be rapidly fabricated and provide detailed characterization of the design space for achieving varied terminal velocities at different device payloads. 2) A 30 mg battery-free programmable computing device that can wirelessly transmit sensor data using backscatter communication, as well as the power electronics and cold start circuitry for robust operation in low or variable light conditions. 3) We integrate an array of sensors capable of measuring parameters like

temperature, humidity, pressure, light intensity, acceleration, and pressure weighing 1.4–12.4 mg and demonstrate power-saving modes capable of operation with as little as 2.8 μW . We further demonstrate sensing applications like environmental monitoring and vehicle proximity detection.

2.2 Results

2.2.1 Wind dispersal mechanism

Successful wind dispersal requires a structure that produces drag and reduces the terminal velocity. Intuitively, the slower the seed falls, the greater chance it has for horizontal displacement by the wind. Two common drag enhancing structures employed by seeds are plumes, e.g., dandelion seeds, and wing structures such as those seen in species of maple trees (e.g. *Acer saccharinum*). Winged seeds autorotate to produce lift by generating a leading edge vortex [82] whereas plumed seeds tend to operate at lower wing loading and benefit from a sharp increase in drag coefficient at low Reynolds number [19, 7]. We explore the latter option as inspiration for miniaturized wireless sensing devices that operate in this regime due to its simplicity and scaling potential.

We have three goals: 1) develop a strategy to fabricate synthetic structures inspired by plumed seeds and investigate their payload capacity, 2) explore design strategies to introduce varied sizes and shapes into individual devices so they disperse to different distances, 3) determine the factors that affect aerodynamic stability to maximize the probability the sensors descend in the correct upright orientation to harvest solar power after impact.

Plumed seeds like dandelions consist of a bundle of bristles extending from a central point. These bristles can be modeled as a porous disk that produces a stable vortex of air above it, which increases its drag compared to a solid disk of equivalent area [26]. Using this type of structure therefore allows for increased drag with less material than a traditional parachute. Creating a synthetic structure that exactly mimics these many thin bristles is however challenging. We note however that the terminal velocity scales with the projected area: $v(t) = \sqrt{\frac{2mg}{\rho AC_D}}$, where m and A are the mass and area of the device, g is the gravitational acceleration constant, and C_D and ρ are, respectively, the drag coefficient and fluid density. We therefore approximate this structure by

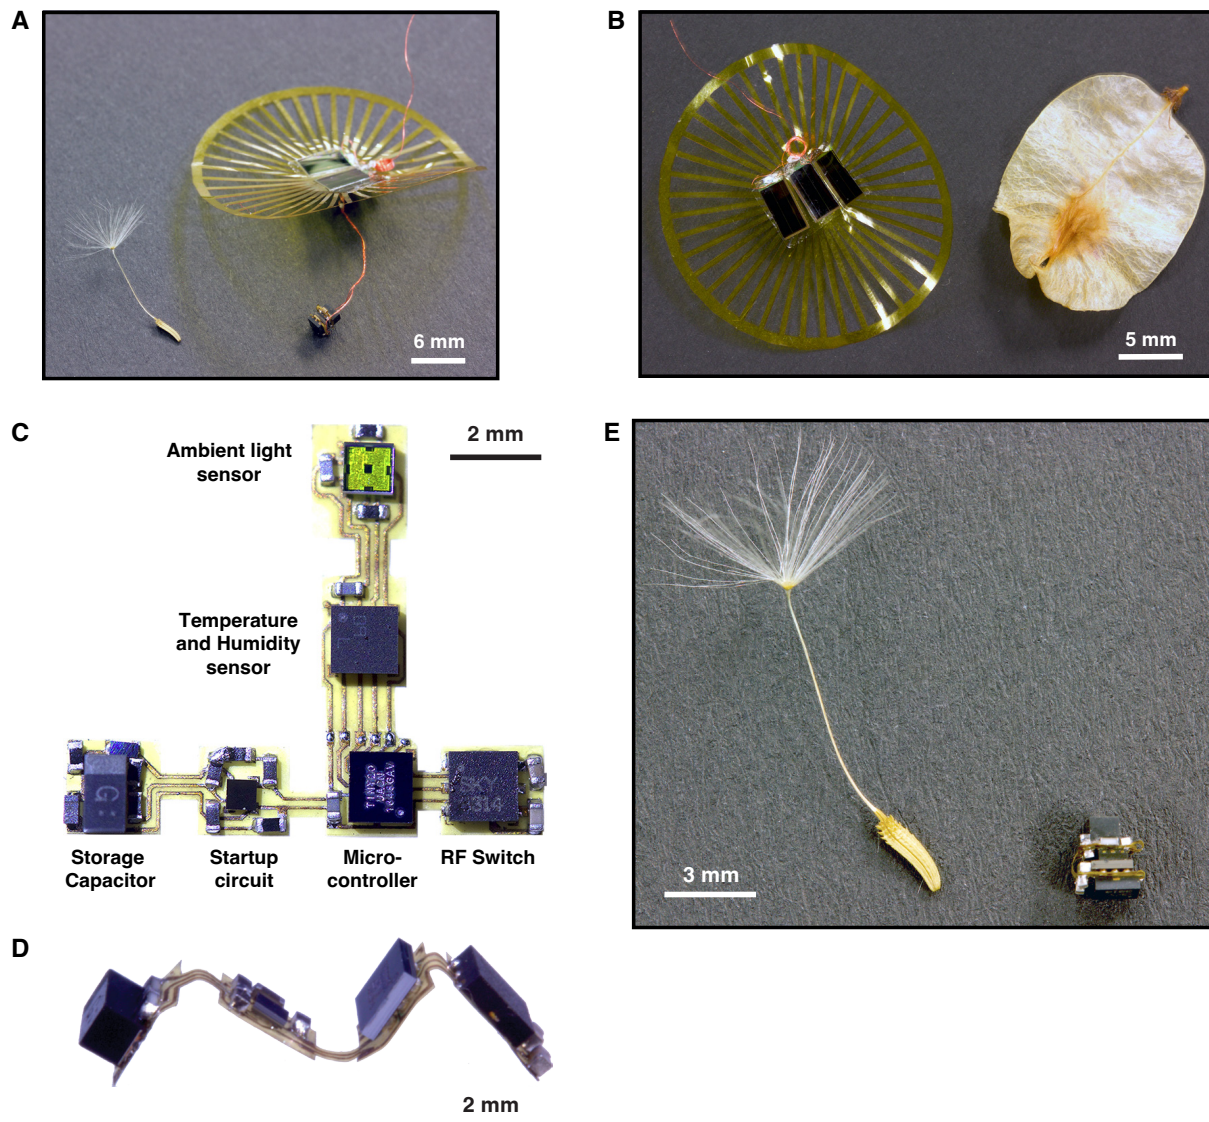


Figure 2.1: Wind-dispersed battery-free wireless devices. A) Battery-free wireless sensing device shown next to a dandelion seed. B) Top view of our battery-free wireless sensing device showing its solar cells and drag enhancing structure next to an elm seed. C) Circuit showing the electronic components including sensors that can be attached. D) The circuit is made on a flexible substrate that can be folded. E) Circuit folded into a cube placed next to a dandelion seed.

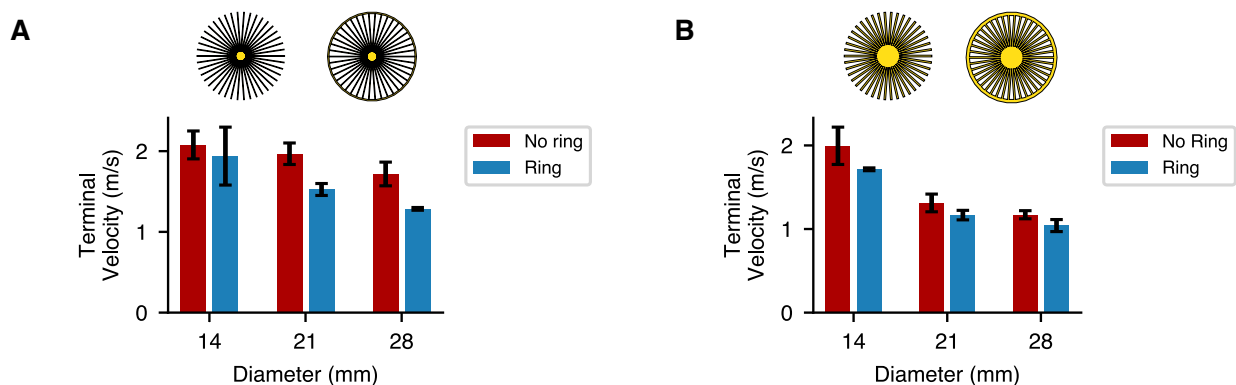


Figure 2.2: Effect of outer ring on terminal velocity. **A)** Terminal velocities of 26% fill structures with different diameters using a 30 mg payload. **B)** Terminal velocities of 55% fill structures with different diameters with and without an outer ring using a 30 mg payload.

producing a 2D projection of the desired dandelion plume consisting of a disk with thin rectangular projections extending from a central point as shown in Fig. 2.1A,B and Fig. 2.2.

To minimize overall weight and enable rapid fabrication, we produce these structures using thin polyimide films (7.5–25 μm). We use a 355 nm laser to cut patterns out of sheets of the material allowing rapid prototyping of different designs. We note that this process could be scaled for mass production as polyimide films are commonly used as the backing for flexible circuit boards, further opening potential for patterning the circuit on the structure as well.

We use this fabrication technique to investigate different designs starting with the diameter of a dandelion seed (≈ 14 mm) and increasing it to double this size (28 mm), similar to larger plumed seeds such as *Tragopogon pratensis* [19]. When attaching a 30 mg weight, we observe that the significantly greater weight than natural seeds causes the bristles to bend inward while falling. This is likely due to the limited stiffness of the thin films; we address this by adding a ring to couple the bristles together and increase stiffness. Video 2.2 shows a high speed video comparing the original design and the reduced terminal velocity of the structure with the ring; Fig 2.2 confirms that this effect is consistent across diameters.

Next, we systematically measure the terminal velocity for 75 designs (3 material thicknesses \times 5 diameters \times 5 attached payloads). We perform empirical measurements due to the complexities of accurately modeling the fluid-structure interactions for thin deformable, freely falling objects.

We test diameters of 10–50 mm for film thicknesses ranging from 7.5–25.4 μm . For each design we attach weights between 20–70 mg to model sensor payloads and perform drop tests from a height of 2 m and measure terminal velocity. Using the terminal velocity and mass we calculate $C_D = \frac{mg}{\frac{1}{2}\rho v^2 A}$, which gives a measure of how efficiently the structure produces drag. We also use these data to calculate the Reynolds number $Re = \frac{\rho v d}{\mu}$, the ratio of inertial to viscous forces in the system. Results for our synthetic structures show that C_D increases as Reynolds number decreases (see Fig. 2.3A); this follows a similar trend to natural plumed seeds [26, 19] and in general the trend observed for all objects in flow (e.g. spheres, cylinders, etc.) [146]. Note that the points with highest C_D are the smallest and lightest weight; this data supports the advantages of miniaturization for wind dispersal.

In the practical context of using these structures for wind dispersal of sensors, it is also informative to see how terminal velocity scales with sensor weight. Fig 2.3B shows this relationship for different film thicknesses. These data show first that increasing the diameter (and therefore projected area A) results in lower terminal velocity as expected; however increasing to dimensions beyond 30 mm appears to result in diminishing returns. Second, achieving terminal velocities below 1 m/s, where even light air flow to a moderate breeze (1–5 m/s) can result in wind dispersal greater than release height, constrains sensor payloads to < 70 mg. Third, both Fig 2.3A and Fig 2.3B indicate counter-intuitively that greater film thicknesses (higher mass structures) generally perform better. A high speed video of a 50 mm, 7.5 μm thick disk carrying a 70 mg payload shows it deforming significantly as it falls (Video 2.3). The edges bend inward due to the low stiffness of the thin film, and so increased diameter beyond a point does not increase projected area and explains why thicker films perform better.

We next explore other ways to vary the design to emulate variation in seed morphology and attain further control over terminal velocity. In addition to the hair-like structures extending outward from the center of the pappus, certain plumed seeds such as *T. pratensis* have additional orthogonal fibers to produce a mesh pattern which could be used to vary porosity [19]. We take inspiration from these to produce mesh infill patterns as seen in Fig 2.4A and Fig. 2.5. We wrote software to generate the desired patterns by rotating a rectangle around a central point and adding concentric

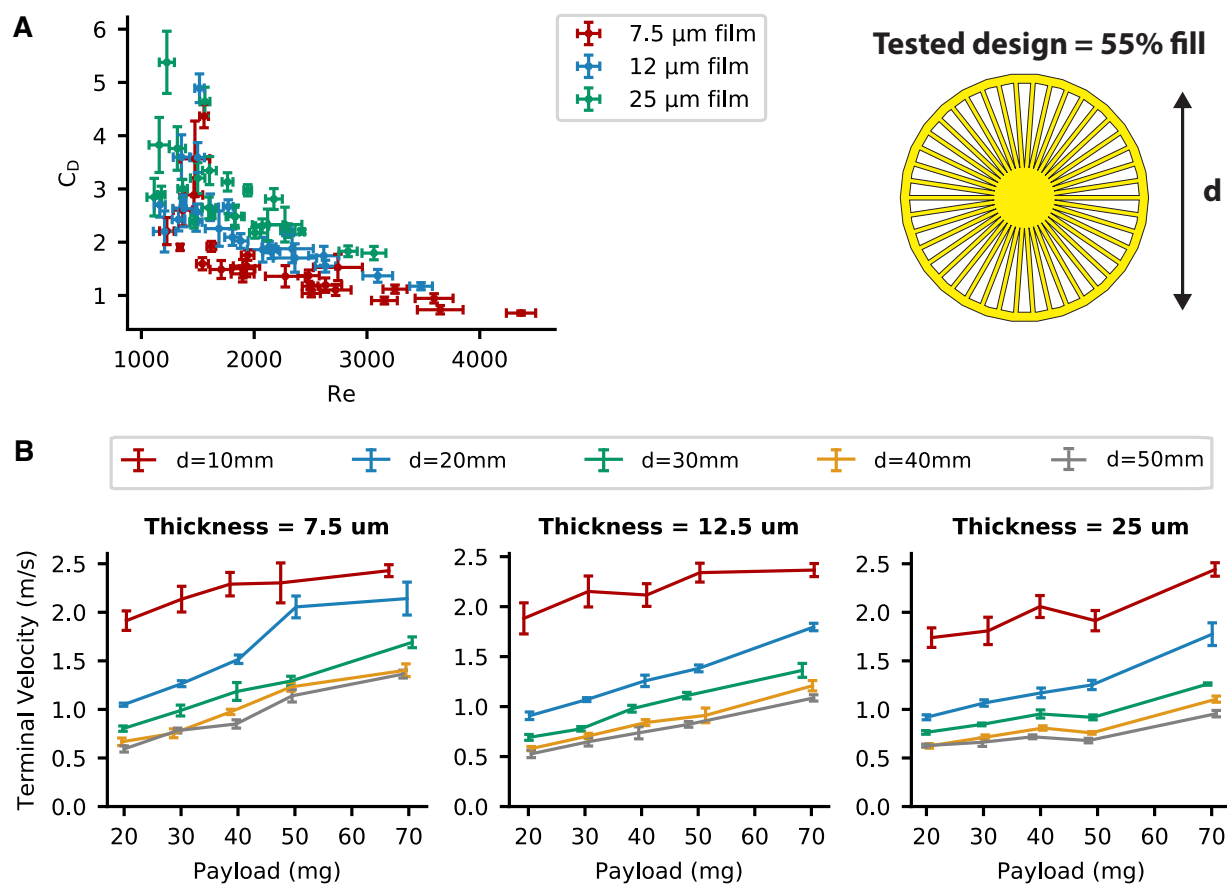


Figure 2.3: Drag structure payload tests. **A)** Drag coefficient (C_D) versus Reynolds number (Re) for 75 designs (3 film thicknesses \times 5 diameters \times 5 payloads, $N=5$ trials per design, error bars= $\pm\sigma$). **B)** Terminal velocity versus payload for the same data shown in A ($N=5$ trials per design, error bars= $\pm\sigma$).

circles to increase the filled area similar to the mesh patterns seen in other seeds and bird feathers.

We empirically measure the terminal velocity of each of the designs shown in Fig 2.4A and plot the results versus the fill percentage in Fig 2.4B. The results follow the expected inverse square root relationship with projected area, showing terminal velocity decreases as the structure's diameter increases. Increasing the filled area also decreases terminal velocity but with diminishing returns beyond 50% fill. The lowest terminal velocity in these experiments – 0.81 ± 0.04 m/s – was achieved with a 28 mm diameter disk and a 100% fill. At 68% fill, the terminal velocities were 1.07 ± 0.04 m/s and 0.87 ± 0.02 m/s with 21 mm and 28 mm disk diameters, respectively. At disk diameters (14 mm) and fill fractions (10-25%) similar to a dandelion seed, the structure has

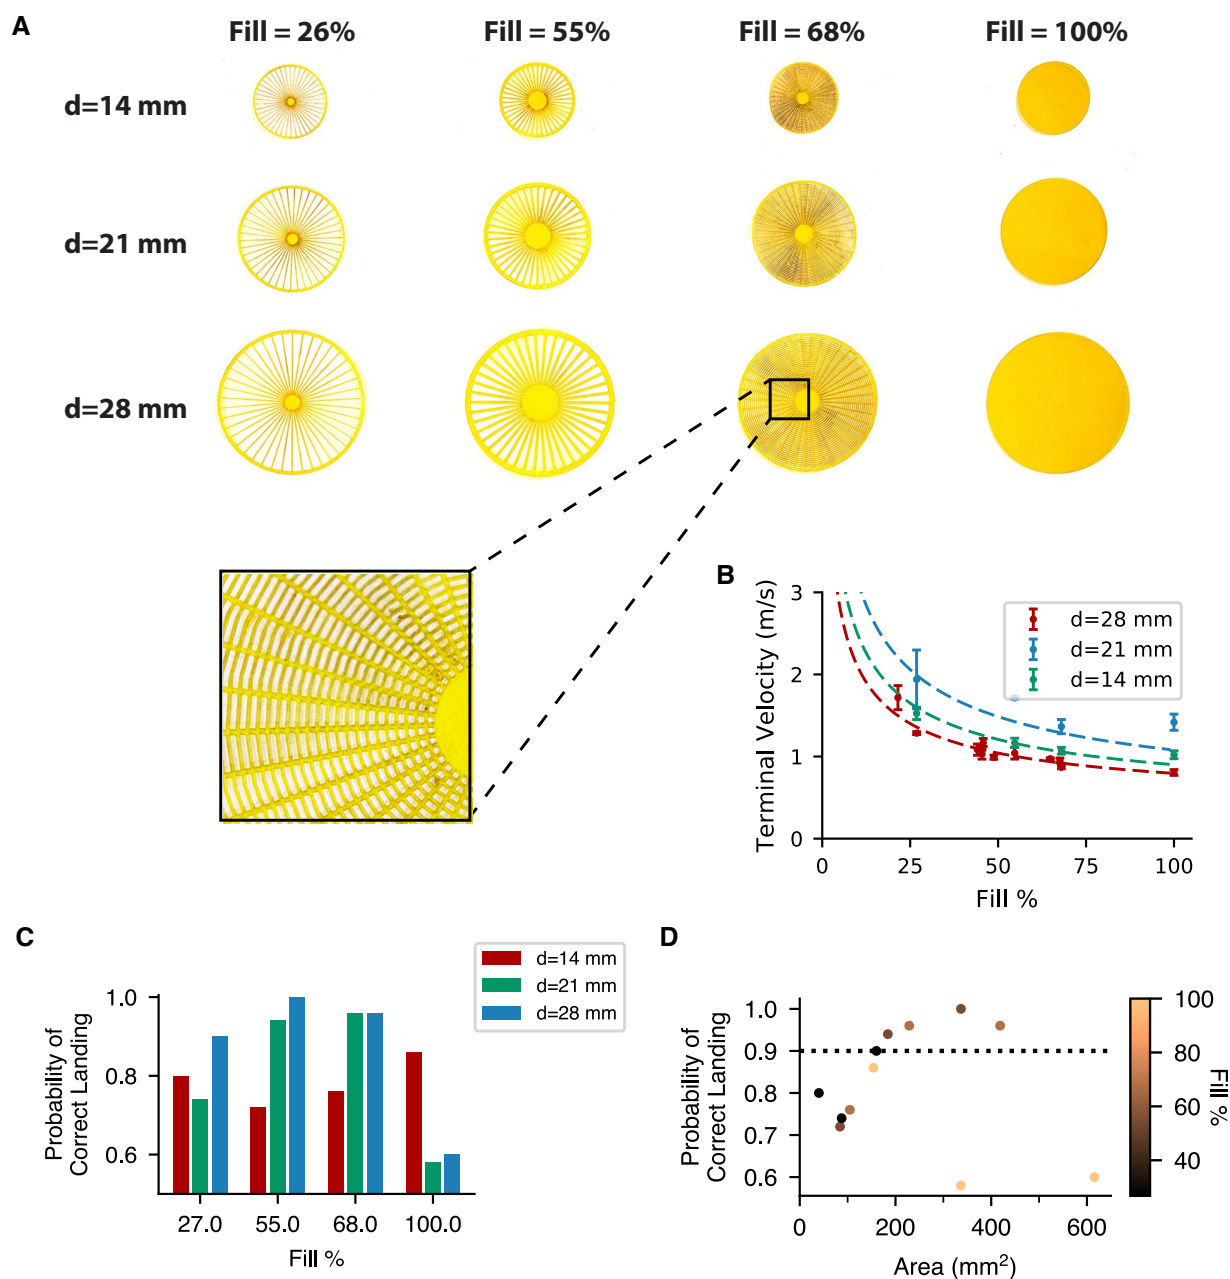


Figure 2.4: Drag structure tests. **A)** Designs of different sizes and fill patterns. **B)** Terminal velocity for designs from C on $12\ \mu\text{m}$ films ($N \geq 5$ trials per design, error bars $= \pm \sigma$). **C,D)** Probability of upright landing versus fill percentage and projected area ($N=50$ trials per design).

a terminal velocity of 1.9 ± 0.36 m/s compared to the 0.39 m/s achieved by dandelion seeds. The difference is due to the greater weight of our payload compared to the seed.




Figure 2.5: Additional fill patterns. Additional fill patterns used to vary porosity. Changes include concentric circles connecting horizontal spokes extending out from the center with various spacings, as well as additional cuts in the concentric circles between spokes

Although 100% infill results in the lowest velocity, similar to elm seeds (Fig 2.1A), a solid disk has an unstable descent and tends to flutter and rotate as it falls, (Video 2.4). In contrast, dandelion seeds maintain a stable descent and so we trade off a small decrease in terminal velocity for the porous structure with a stable descent. This behavior is similar to that of a Badminton shuttlecock. By placing the wireless device at the center and adding curvature, we can emulate this structure and define its upright orientation. To produce this curvature, we placed the patterned piece of polyimide in a curved mold; the mold consists of a 3D printed plastic cube with a 7.7 cm diameter cylinder cut out (Fig 2.6). We applied a piece of double-sided tape to make the structure hold its curvature and this doubles as a means to attach solar cells to the top of the structure. Though we produced the drag enhancing structure and electronics separately, they could be integrated by patterning the traces onto a sheet of copper-coated polyimide.

Adding curvature causes the structure to right itself when dropped upside down (Video 2.5). We computed the probability of landing in the correct orientation (i.e., curved sided up) by dropping each design with a 30 mg payload from a height of 2 m ($N=50$ trials). At larger disk diameters (≥ 21 mm), 100% infill causes instability and has an average 50% probability of landing on the




Figure 2.6: Fabrication process. Fabrication process used to produce the curved drag enhancing structures. Thin film patterns are first cut with a laser and then the curvature is defined by placing them in a curved mold. A thin layer of double-sided tape is used to define the shape and attach the solar cells.

correct side, whereas a 26-68% infill results in a 90-100% probability (Fig 2.4C). At a 14 mm diameter, the structure maintains a stable descent but does not always land in a consistent orientation. Plotting the data versus the actual area of material reveals a pattern: structures with an area below 160 mm^2 , regardless of their diameter and infill patterns, are not large enough to produce the righting behavior (Fig 2.4D). Structures above this minimum area, and without a very high infill percentage, are stable and likely to land upright.

2.2.2 Battery-free programmable wireless sensors

The results above highlight the need for miniaturized sensing electronics. The conventional solution to this is to develop custom integrated circuits (ICs), e.g., smart dust and neural dust [124, 67] and Michigan micro motes [81, 20, 21]. However, single monolithic IC solutions are often highly specific to a particular sensing task. Moreover, prior works have not demonstrated fully functional battery-free sensors capable of passive wind dispersal at small (milligram) scale and have been limited to short communication ranges. A device weight of 50 mg would allow for deployment of 1000 sensors using even a small drone that can carry 50 g.

Our approach instead uses microcontrollers, which are small computing platforms designed for general-purpose programming and easy integration with different sensors. Additionally, this method presents a highly accessible platform for researchers with diverse backgrounds to build on and rapidly prototype new applications at millimeter-scale without the significant time, expense, and expertise required to design custom silicon. We designed the first millimeter-scale, battery-

free computing devices that use microcontrollers, weigh only 28.4 mg, and are fully integrated with backscatter communication and power electronics (Fig 2.1A,B). The electronic components are assembled on a 12 μm thick flexible polyimide substrate (Dupont Pyralux AC121200). By designing each subsystem to fit within an approximately 2×2 mm square with interconnects to the next block, the circuit can be folded to fit within an 3 mm cube (Fig 2.1D,E). This dense packaging emulates the structure of many wind-dispersed seeds, where a small seed is attached to a drag-enhancing structure.

Our system is designed around the ATtiny20 microcontroller, which measures 1.56 mm \times 1.4 mm, weighs 2.2 mg, and includes an onboard temperature sensor connected to an analog-to-digital converter. However, the microcontroller lacks a radio required for wireless communication or the circuits to power it. Our recent works on battery-powered millimeter-scale wireless sensors [58, 59] use small Bluetooth radios for communication; these require external components like crystal oscillators and weigh 166–248 mg. We instead designed a backscatter communication system that allows communication from the microcontroller, without a radio.

Backscatter systems communicate using reflected signals rather than actively generating radio frequency (RF) transmissions. This significantly reduces the power required to communicate since generating high frequency RF signals can be the most power-consuming part of a communication system. A backscatter device encodes data in reflected signals by modulating its radar cross section using both the antenna and the components connected to it [114]. Doing so produces a mixing operation at the antenna: an incoming radio signal is multiplied with the switch modulation signal (Fig 2.7A). This lets us decouple the energy and size-expensive parts of a radio and move them to an external device that broadcasts a signal that the small, low-power sensor can reflect to communicate.

An impedance boundary between two media allows a fraction of the signal to pass through it and the remainder to be reflected. By choosing two different impedances, we produce two different reflective states to communicate information. We therefore toggle between an open and a short circuit state. While these two states should ideally produce a +1 and -1 , respectively, components like switches introduce losses; our implementation has a loss of -2.21 dB. We connect the output

of the backscatter circuit to a lightweight antenna. To design an antenna of minimum size, we use a piece of 43 AWG wire. To make it resonant at 915 MHz, we coil the wire for 12 turns with an 1.4 mm radius at the base. The resulting antenna has a -5 dB loss compared to typical 915 MHz monopole antenna.

Range however presents challenges for backscatter. A traditional radio transmission from a distance d incurs a path loss proportional to $\frac{1}{d^2}$. In contrast, backscatter has two links: one from the transmitter to the backscatter device (d_1), and another from the backscatter device to the receiver (d_2), incurring a path loss proportional to $\frac{1}{d_1^2 d_2^2}$. Thus, prior millimeter-scale, programmable backscatter systems like Living IoT have a 1-2 m backscatter range[60]. To improve this range, we performed coding and modulation on the microcontroller to mimic signals expected by high-sensitivity commodity radio receivers.

We built a backscatter receiver using 900 MHz *Semtech SX1276* radio chipsets, which can decode on-off keying radio signals at -117 dBm. In our system, the transmitter broadcasts a single tone. However, this chip requires specific fixed-packet structures that it is designed to decode. While our goal is to control the backscatter signal and produce the appropriate packet structure, we are constrained by our small form factor microcontroller (ATtiny20, Atmel). While the Semtech SX1276 receiver supports different modulations, like chirp spread spectrum and frequency shift keying, the microcontroller limits us to operate using its on-off keying mode to send digital data. Prior work has shown long-range backscatter communication using chirp spread spectrum coding [135], but has orders of magnitude larger weight and size and requires power-consuming, field-programmable gate arrays to perform computationally intensive operations. In contrast, our design runs on small form factor microcontrollers and meets our size, power and weight constraints. We ran the ATtiny microcontroller at 8 MHz, which constrains the bandwidth and resolution of the frequency shift keying signals it can produce. To generate on-off keying modulation a ‘1’ bit is communicated by the presence of an RF signal, and a ‘0’ bit by its absence. This can be implemented by toggling the microcontroller’s output ON and OFF.

In addition to choosing a sensitive receiver, we must also reduce the interference it will receive directly from the USRP transmitter broadcasting the RF carrier. To this end, we employ a technique

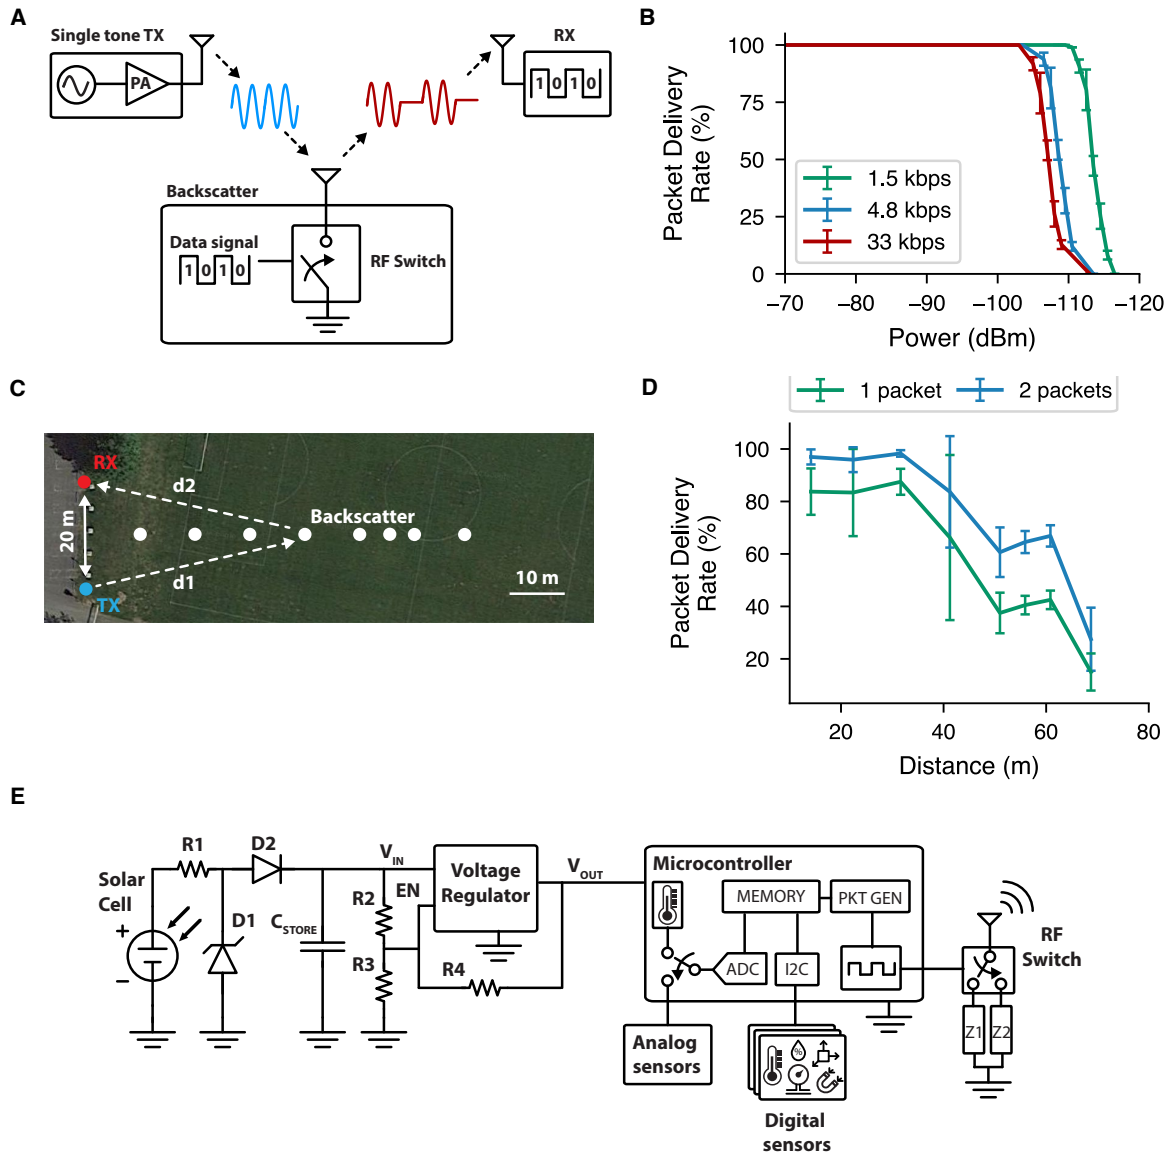


Figure 2.7: Backscatter communication. **A)** A transmitter broadcasts a single frequency tone, the backscatter device toggles an RF switch connected to an antenna, and a receiver decodes the backscattered data. **B)** Error rate of backscattered packets received on an SX1276 radio at different power levels and bitrates in benchtop tests ($N=500$ packets/point, error bars= $\pm\sigma$). **C)** Outdoor range test: transmitter and receiver are 20 m apart and backscatter device is moved away equidistant from both. **D)** In-air packet delivery rate outdoors vs. range at 4.8 kbps with a single transmissions and 1x re-transmission ($N\geq 100$ packets per point, error bars= $\pm\sigma$).

called *subcarrier modulation*, which sends data after applying an additional frequency shift. This technique adds some separation between the transmitter and the backscatter signal in the frequency domain, thus reducing interference. We leverage the fact that backscatter itself is a multiplication operation. So, to produce a frequency shift, we would ideally like to multiply the incoming signal with a sinusoid. While an RF switch does not let us produce a sine wave, we can approximate this waveform using a square wave toggling between the two states at the same rate. If the square wave is decomposed as a Fourier series, its primary component is a sine wave at the desired frequency, therefore producing the desired frequency shift. We choose a frequency offset of 4 MHz, which is determined by the microcontroller's maximum frequency.

Fig 2.8 shows the packet structure the SX1276 is designed to receive. It consists of a preamble of at least 3 bytes that is a sequence of alternating ones and zeros (0xAA or 0x55). The receiver uses this structure to identify the start of a packet. The preamble is followed by a 4-byte address field that uniquely identifies a transmitter and a optional length field for variable packet lengths. The payload is followed by an optional 3-byte cyclic redundancy check (CRC).

To generate a packet, we first sample our sensor either using the microcontroller's ADC to read analog sensors, or by sending a command using the I2C protocol to read digital data from a variety of external sensors described in further detail below. The data is sampled and stored in memory. Next, the microcontroller constructs the packet which consists of fixed preamble bytes (0x555555), followed by a 4-byte address field and the sensor data payload as shown in Fig. 2.8. We omit the CRC for 4.8 kbps and 1.5 kbps packets to reduce their length, however, this field may be used for higher bit rates (e.g., 32 kbps). To generate an on-off keying signal with subcarrier modulation, a timer on the microcontroller is configured to output two complementary 50% duty cycle pulse width modulated (PWM) signals at the desired frequency. When these two outputs are enabled and connected to the switch, they produce a backscattered signal shifted by the subcarrier frequency. An assembly function translates the binary data to the appropriate delays to achieve the desired bit rate and toggles the PWM signal to produce the on-off keying modulation.

To measure the backscatter packet delivery rate in Fig 2.7B, we use a bi-directional coupler to isolate the transmitter and receiver. A software radio (USRP E310) is configured to transmit a

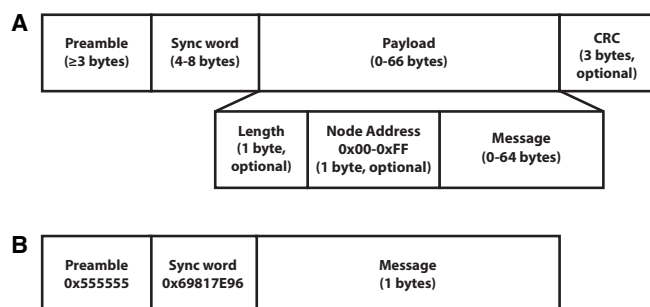


Figure 2.8: SX1276 receiver on-off keying packet structure. A) Packet structure for the SX1276 receiver operating in on-off keying mode. B) Minimal packet structure used for testing the lower bit rates of 1.5 and 4.8 kbps.

919 MHz single tone into the input port of the bi-directional coupler, and the backscatter switch is connected to the output port. The receiver (SX1276) is connected at the output-coupled port through a series of attenuators, and the input-coupled port is terminated with a 50Ω load. All devices are connected directly to the coupler's male-to-male SMA connectors and housed in metal containers to prevent weak radiated signals from propagating to the receiver. The microcontroller was programmed to send 100 packets with fixed data to measure packet error rate. The signal's output power was varied by decreasing the USRP output. The power output of the packets was measured on a spectrum analyzer (Tektronix MDO4104-6). A total of 500 packets were sent for each power level, and these experiments were repeated for different bit rates. System performance is limited by the noise floor of the receiver, which increases with the bandwidth of the filter used at the receiver. The minimum filter bandwidth required to receive 33 kbps packets was 41.7 kHz. In contrast, 1.5 kbps packets could be received with a filter bandwidth of 12.5 kHz, resulting in a lower noise floor and improved sensitivity. Further decreasing filter bandwidth resulted in higher packet loss, likely due to the frequency resolution of the microcontroller's oscillator.

For the range measurements outdoors, we set up the transmitter and receiver at one end of a field, 20 m apart. The receiver was the SX1276 module connected to an 8.5 dBi patch antenna. A USRP was configured as a transmitter and connected to a power amplifier (RF5110G, Qorvo) to boost output power to 25 dBm. The transmitter is also connected to the same 8.5 dBi patch antenna. The antennas are both placed to radiate outwards onto the field. Due to their radiation patterns,

their parallel placement helps provide isolation. Packet error rates were collected by connecting the backscatter switch and microcontroller to our form factor wire antenna. The antenna consists of 12 turns of 43 AWG wire with a 1.4 mm diameter as well as an 17 mm long straight length of wire as seen in Fig 2.1A,B. It was held at a height of 1 m and moved to increasing distances away from the transmitter and receiver along a line perpendicular to their midpoint as shown in Fig 2.7C.

The transmissions from our backscatter devices support multiple bit rates and can be successfully decoded at signal levels as low as -111 dBm (Fig 2.7B). In an open outdoor field (Fig 2.7C), 4.8 kbps backscatter transmissions were decoded with a packet delivery rate of 66-88% up to 41 m; this increased to 84-98% when each transmission was repeated twice (Fig 2.7D). Packet repetition enables reliable reception in the absence of acknowledgements.

2.2.3 *Harvesting circuitry*

Small batteries either have very limited capacities [20] or add significant weight to the payload [60]. In contrast, solar energy harvesting allows our wireless sensors to operate without a battery, which significantly reduces weight and allows for a longer lifetime; however, the variable nature of sunlight requires addressing multiple technical challenges to build a system that can operate robustly. Considering the absence of sunlight overnight, the system must be able to perform a “cold start” from a powered-off state. This is challenging since the microcontroller takes 4-5x more energy to boot up than to backscatter data. We include a capacitor that can accumulate energy from the solar cells over time and store sufficient charge to transmit a burst of packets, but when cold starting from a discharged state the capacitor will not have stored sufficient energy to buffer the startup current spike. If the solar cell output is insufficient, the microcontroller gets stuck in a loop of attempted starts. A detailed description of the circuit is given below.

We use thin film solar cells (40 μm thick, Microlink Devices) that provides up to 0.25 mW/mm². Using a solar cell still requires an energy storage element in order to operate in variable light conditions or in low light when the solar cell cannot provide enough power for continuous operation. We chose tantalum capacitors due to their high energy density, which allows for a large capacitance in a small and lightweight package. The microcontroller requires 2.7 V or higher to run at the 8 MHz

clock frequency used for subcarrier modulation described above; however, even in bright light conditions, a single 2.5×5 mm solar cell produces only 2.6 V. Instead of using a boost converter circuit which could add significant weight [63, 59], we instead connect 3 small solar cell segments in series. We note that when operating in bright environments fewer cells could be used to reduce weight. While this strategy of combining solar cells in series allows for consistent operation in low light conditions, it presents a problem for very bright light (e.g. 1 sun); in bright sunlight the cells can produce a maximum voltage of over 7 V which is higher than the other circuit components can tolerate. To allow for robust operation across this spectrum of light conditions we add an 5.6 V zener diode (D1) for protection at the input. At low voltage, current flow through the diode is 1 μ A or less, however when it approaches 5.6 V it allows excess current to flow to ground and prevents damage to the remainder of the circuit such as the capacitor and voltage regulator which are rated for a maximum of 6.3 V and 6.0 V respectively.

Upon cold-start, the chip waits for 64 ms to make sure the power supply is stable enough to power on the chip. This is followed by 6 additional oscillator cycles and 21 system clock cycles during which configuration data is loaded before the system can begin executing program instructions. In order to cold-start, the capacitor must have buffered enough energy to begin executing instructions before it can enter its low power modes for further operation. Our system therefore requires some method to sense charge level prior to being able to perform computation or logic.

To do this with minimum power and size, we utilize the enable functionality of an 0.645×0.645 mm low-dropout voltage regulator (Toshiba Semiconductor TCR3UG28ALF) to toggle the output on and off. We connect the input to the storage capacitor and the output to the microcontroller. To sense the charge level of the storage capacitor and trigger turning on the microcontroller, we use a voltage divider (Fig 2.7E). By setting the ratio of the two resistors (7.6 $M\Omega$ and 3.49 $M\Omega$), we can set the enable pin to exceed the required threshold to turn on (0.7 V) when the capacitor has reached a target voltage (3.3 V); however, this suffers the same problem as before: once the chip turns on and starts drawing power, capacitor voltage decreases and would cause the chip to turn off. To introduce the required hysteresis, we add an additional resistance (30 $M\Omega$) to pull the enable pin up to the output. At startup, the output is floating, which has minimal effect. After the

regulator turns ON, it uses this path to stay ON until encountering a negative threshold (2.1 V). This approach lets the microcontroller remain powered long enough to start up.

Fig 2.9A shows the system starting from near zero charge and operating in a burst sampling mode where it repeatedly backscatters sensor data until charge is depleted. The red line shows the capacitor voltage charging up. After it crosses the cold start threshold, the blue line, which shows voltage regulator output gets higher, powering the microcontroller. After booting, the device sends a burst of backscatter packets. Upon reaching its negative threshold, the voltage regulator disables the microcontroller so the system can avoid complete discharge. Fig 2.9B shows the time required to cold start under different outdoor lighting conditions down to 1400 lux. Below this light level, the startup circuit does not trigger periodically, e.g., at 500 lux, the first two packets were transmitted 3 min apart, but the third was transmitted after 22 min. In this burst sampling mode, the size of the storage capacitor determines the number of packets that can be sent in each burst. The 68 μF capacitor charges faster but can reliably transmit only a single packet per burst. The 100 μF capacitor is 1.5 mg heavier and charges slower but can transmit 2 packets per burst, achieving a higher throughput (Fig 2.9C). The weights of these components are shown in Fig 2.9D.

We captured the waveform in Fig 2.9A on an oscilloscope (Tektronix MDO4104-6), starting with the system completely discharged. A desk lamp (Phive, OL-1) is turned ON at $t=3.45$ s, and the the red line indicating capacitor voltage begins to increase. Upon reaching a threshold, the voltage regulator turns on (blue line). The microcontroller enters its startup phase (72 ms), after which it begins toggling the control signal to the RF switch to backscatter data (yellow line). We measured the time required for a cold start in Fig. 2.10 and Fig 2.9B under both indoor light conditions using LED lamps (Amscope LED-144 and GS Vitec LT-V9-15) and outdoor natural light. Illuminance was measured using a lux meter (Dr. Meter LX1330B) placed next to the solar cell. To determine cold start time, we recorded the waveform on the oscilloscope and used it to measure the duration from when the circuit begins charging to when the microcontroller turns ON. We then measured the time between packets to determine throughput, shown in Fig 2.9C. We performed these measurements for two different storage capacitors, 68 μF (AVX, F980G686MMA) and 100 μF (AVX, F980J107MMAAXE).

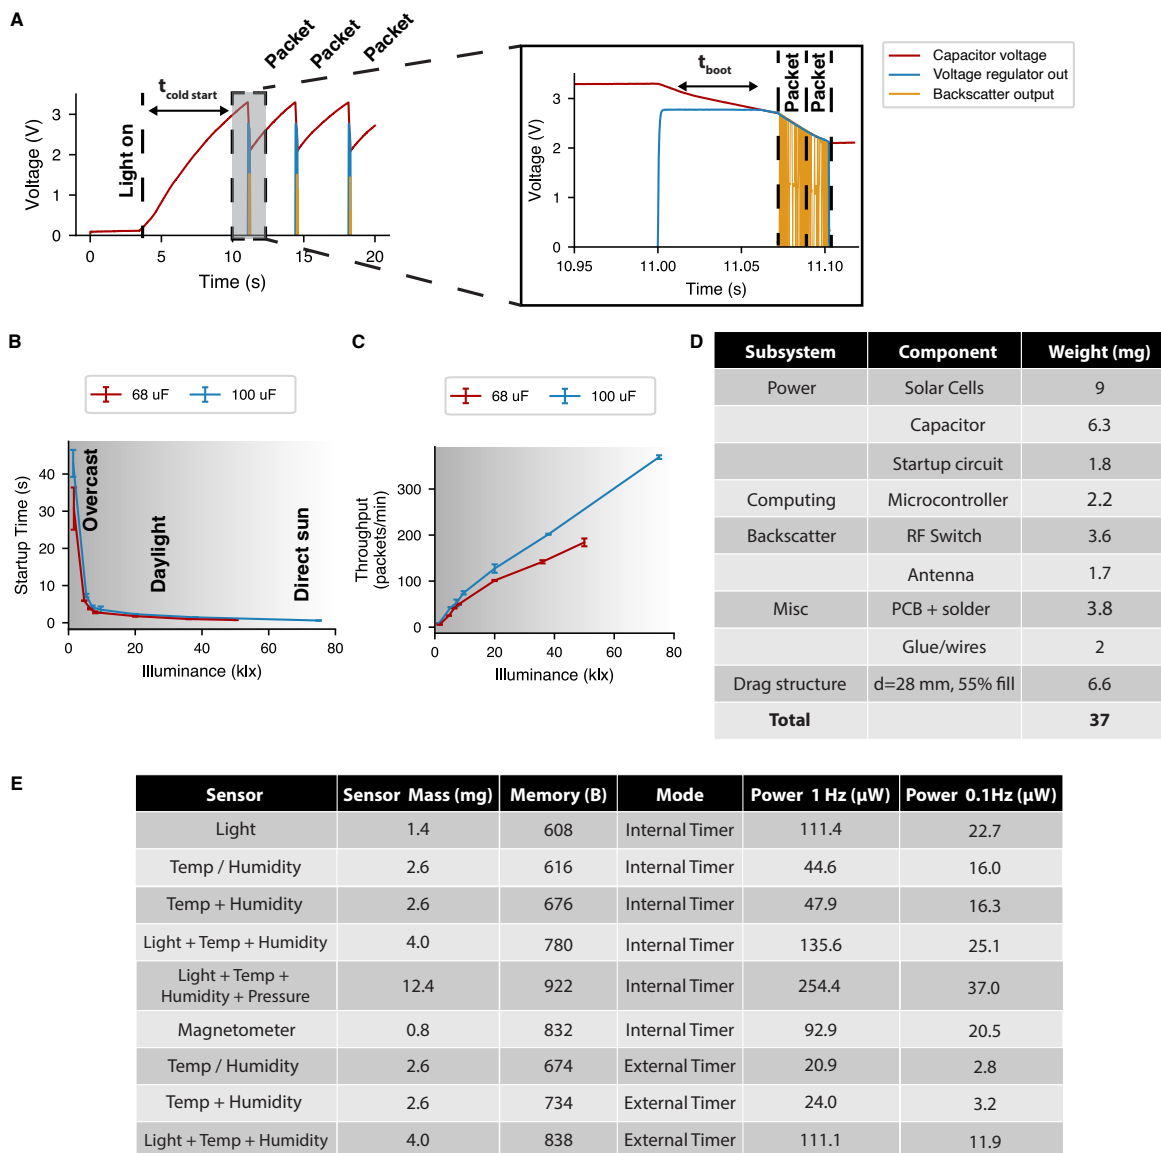


Figure 2.9: Battery-free wireless sensor. **A)** Wakeup sequence in continuous operation mode. **B,C)** Startup time and throughput in continuous mode under outdoor light conditions, $N \geq 5$, error bars= $\pm\sigma$. **D)** Component weights. **E)** Sensor weight and power measurements in different operating modes.

After startup, in sufficient light levels, we can also leverage the microcontroller’s sleep modes to sample data at a consistent rate while consuming minimal power. We demonstrate two sleep modes shown in Fig 2.9E. In these modes the microcontroller performs duty cycled operations and spends most of its time in a low power sleep mode with all of its peripherals OFF. Using either a low frequency internal timer on the microcontroller, or even lower power timers available on

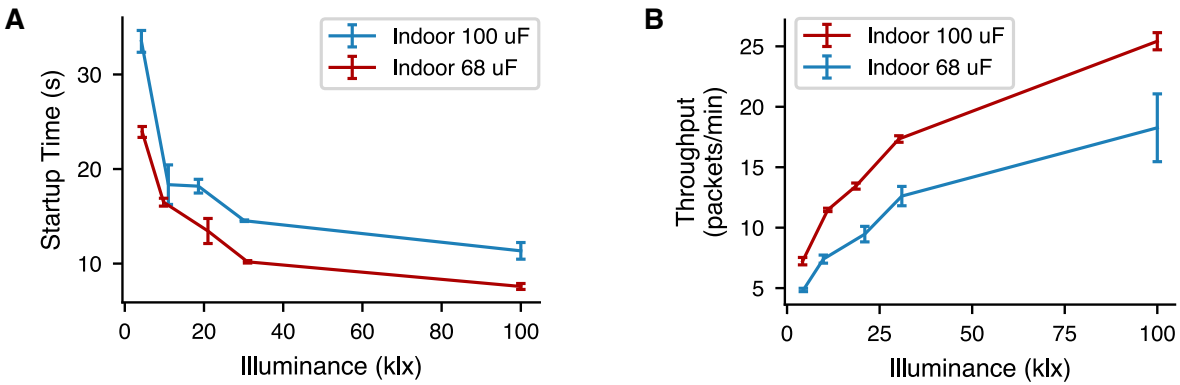


Figure 2.10: Indoor light power harvesting results. **A)** Cold start time required to charge the storage capacitor from zero to the time at which it can backscatter its first packet; measured using indoor LED lights ($N \geq 5$, error bars= $\pm\sigma$). **B)** Throughput of the sensor device after startup at different illumination levels ($N \geq 5$, error bars= $\pm\sigma$).

sensors, the system stays ON continuously, periodically waking up to sample and transmit data before returning to sleep. Implementation details of the sleep modes are given below.

2.2.4 Power consumption and Operating modes

After starting up, the system can operate in three modes.

a) Burst sampling. The CPU and all peripherals remain fully ON as long as power is available. This is most useful in scenarios with too little power to maintain the device in sleep. This mode allows the device to start up, transmit a few packets with its available energy, and turn off. The trade-off is long startup times and lack of a programmable sampling rate. This mode also allows for maximum data rate outputs and could be used in scenarios with bright sunlight that allows for continuous operation. To operate in this mode, the chip is programmed to execute instructions, e.g., sampling sensors and transmitting data, in the main CPU loop that executes repeatedly.

b) Sleep with Internal Timer. The second operation mode utilizes the lowest power sleep mode of the microcontroller which powers down the main clock and all peripherals. We use the low power Watchdog Timer to restart from this state without external hardware. This onboard timer runs at a low clock frequency and results in sleep power consumption of around 5 μA . The watchdog timer can be programmed to wake up periodically (e.g. 0.125-60 Hz). In this mode upon startup, all

peripherals are turned off, and the chip is immediately set to go to sleep and wake up after the maximum 8s delay to allow the capacitor voltage to recover from the high startup power draw. Upon the first wake-up from sleep, the device begins sampling and transmitting sensor data and the sampling rate can be increased. Each time the chip wakes up from sleep it samples sensor data, and transmits it before returning to sleep. This mode requires harvesting enough energy to maintain the timer and make it through the first sleep cycle, but does not incur the cost of repeated startup and allows for a consistent sampling rate without the addition of other components.

c) Sleep with External Timer. The chip can also be woken up from its lowest power sleep mode by a signal sent from an external chip. The advantage of this is that it does not require running the onboard timer, and allows the use of even lower power timing circuits integrated into many sensors. For example, the HDC2010 temperature and humidity sensor includes a programmable timer with a power consumption of $< 0.1 \mu\text{A}$. This order of magnitude reduction allows for regular sampling rates at the lowest power consumption as shown in Fig. 2.9E. This also allows for even lower power asynchronous, event-based sampling as the same temperature sensor, as well as other magnetometers, accelerometers, and pressure sensors described also have onboard logic to trigger a wake-up signal only when the sensor data exceeds a threshold. To enable this mode, upon startup, the chip immediately sends a command to the sensor to enable periodic wake-up, and then proceeds to turn off all features and go into its lowest power sleep mode. Upon receiving a wake-up signal, the device samples the sensor data and transmits it before returning to sleep.

We demonstrate operation in sleep modes with six different kinds of sensors on our platform capable of measuring temperature, humidity, light, pressure, magnetic fields, and acceleration. These sensors are chosen due to their small size and weight, and low average power. Each communicates data using the Inter-Integrated Circuit (I2C) protocol. We demonstrate operation at various sampling rates in the sleep modes described above for different sensors and combinations of sensors and show the results in Fig 2.9D and Table 2.1

2.2.5 *Sensor integration*

We demonstrate operation of six different sensors on our platform capable of measuring temperature, humidity, light, pressure, magnetic fields, and acceleration. We demonstrate operation at various sampling rates in the sleep modes described above for different sensors and combinations of sensors and show the results in Fig 2.9E and Table 2.1.

These sensors are chosen due to their small size and weight, and low average power. Additionally certain sensors include low power timing circuitry that can be used to wake up the microcontroller as described above. We demonstrate that each one is well within the limits of our microcontroller's computational capabilities. Each of these sensors outputs digital data over an I2C interface. Because this microcontroller does not have dedicated hardware to send I2C commands as a Primary device we implement this capability in software. Each time the microcontroller starts up or wakes up from sleep, it writes commands to the sensor to sample new data, waits briefly for the sampling to complete if necessary, and then reads the data into memory. The data is then transmitted via backscatter. We note that in the current implementation the data is read and transmitted as raw bytes and conversion to floating point outputs and higher level data processing is done at the access point to minimize the computation and power. Applications that require checking a threshold and other low complexity operations can however be implemented onboard for specific applications.

2.2.6 *Network Simulation*

If a large number of our devices were to transmit simultaneously, their packets may collide and prevent the receiver from decoding them. We perform simulations based on measured data to explore three different strategies for scaling our design to larger networks. The first two leverage natural lighting variation shown in Fig 2.11A,B and the third introduces a pseudorandom delay.

a) Natural lighting variation with no timer. In our outdoor tests of startup time, variation in light conditions causes the devices to have different startup times and charging rates that reduce the probability of collisions. Figs. 2.9A and Fig 2.11A shows the variance in startup time for different

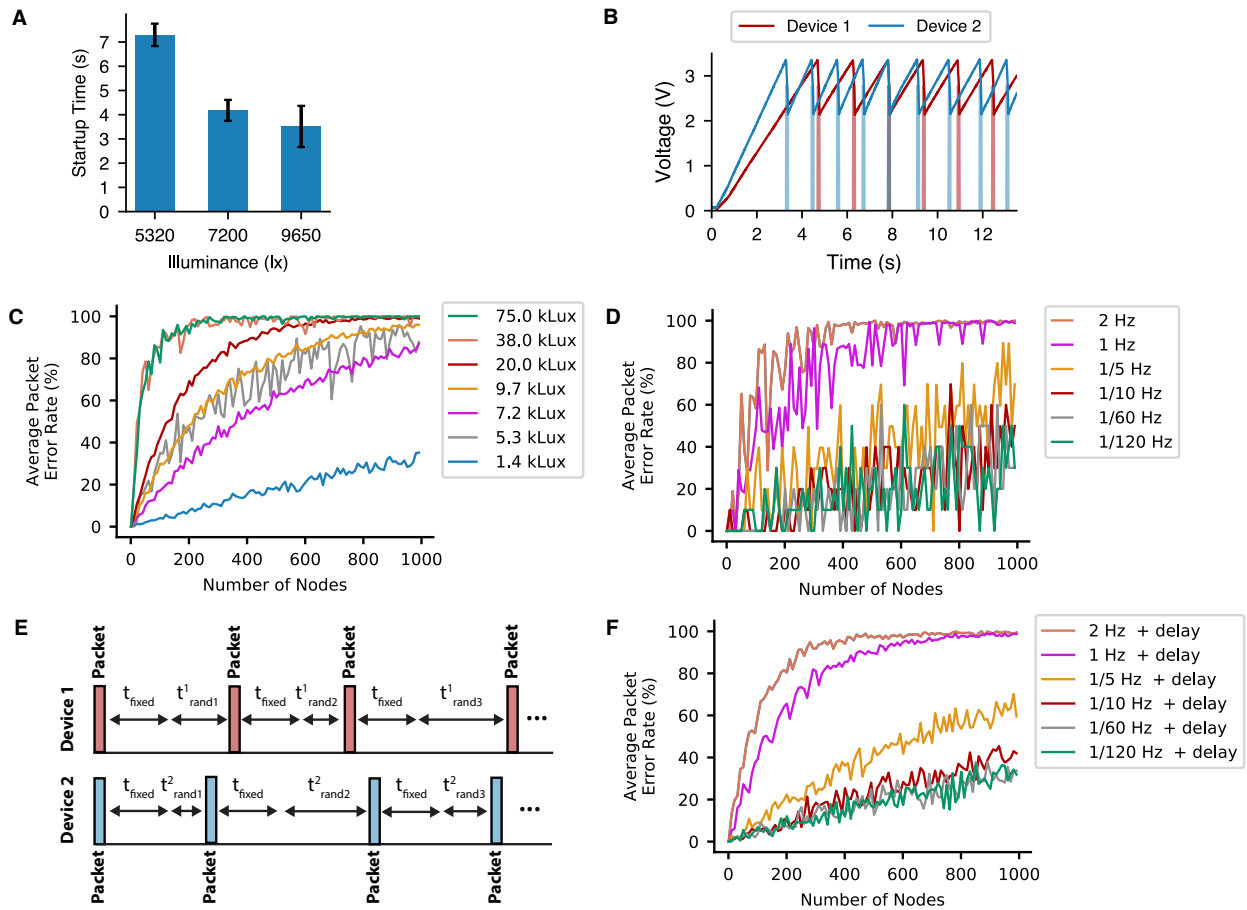


Figure 2.11: Supporting transmissions from multiple devices. **A)** Natural variance of startup time for outdoor measurements. **B)** Recordings from similar light conditions show natural variation reduces collisions. **C)** Network simulation of up to 1000 nodes each sending 100 packets. Node start and inter packet transmission times are normally distributed based on power harvesting data ($N=100$ packets per node \times 10 repetitions=1000). **D)** Simulation with normally distributed startup time assuming 1.4 kLux followed by 10 s inter-packets delay ($N=1000$). **E)** Addition of random delay between packets. **F)** Simulation with random delays up to 8 packet lengths with random start times assuming 1.4 kLux ($N=1000$).

light levels. The plots show that even under similar light conditions (trials for each light level were performed consecutively in the same location) small changes such as a passing cloud or shadow can affect the rate at which a device harvests power. In some cases we observe a standard deviation of up to 0.8 s, which is orders of magnitude larger than the 13 ms required to send a 4.8 kbps packet.

Fig 2.11B shows the startup waveforms for two trials at around 9650 lux, aligned to start at the same time. Different charging rates result in only one overlapping transmission period. Fig 2.11C shows a simulation with increasing network sizes. For each network size, 100 packets start times were generated for every node. Packet start times were sampled from a normal distribution using the mean and variance measured for each light level: $t_{start} \sim \mathcal{N}(\mu_{lux}, \sigma_{lux}^2)$. Each of the 100 subsequent packet transmission times was also generated using the measured mean and variance for the inter-packet delay. Next the packets for a randomly selected node were analyzed to see whether other transmissions were within 1 packet length indicating a collision. The count of collisions was used to compute its packet error rate. Packets are assumed to have 11 B (7 B preamble + 4 B payload sufficient for 2 sensors) and transmitted at 32 kbps. This experiment was repeated 10 times to produce an average packet error rate experienced by a node in the network for a given network size. The results in Fig 2.11C indicate that for small networks with tens of nodes this may be a feasible strategy; however network performance will vary during the day and it does not scale to large networks in bright light.

b) Startup variation with fixed rate. The timer based sleep modes on our device allow for transmitting packets at fixed intervals. A similar experiment described above was performed. Startup times were generated based on the lowest lux level (1.4 kLux) assuming all nodes in the network would be exposed to low light levels upon sunrise and continue from then on to operate in their fixed rate mode. The inter-packet delay was set to fixed values available on the external timer. We assume here that compared to the packet lengths (milliseconds), variation in the sleep timers (with delays determined by >100 kHz clocks) would be negligible. The results in Fig 2.11D show as expected that high rates result in poor performance, but at rates below 1 packet/second, performance improves noticeably but with high variance.

c) Pseudo-random Delay. To improve upon the fixed rate method described above, we also explore a pseudo-random time offset, similar to Wi-Fi, as shown in Fig 2.11E. Instead of transmitting packets at a fixed interval of t_{fixed} , each device adds an additional pseudo-random time offset t_{random} ; this can be easily generated on a microcontroller using a linear feedback shift register (LFSR) circuit. For each packet we set t_{random} to a randomly selected integer multiple of packet

lengths up to a maximum of $8t_{packet}$. The results shown in Fig 2.11F show this approach can support large networks where the data is transmitted at rates less than once every 10 seconds.

2.2.7 Outdoor and sensor evaluation

We performed experiments outdoors in varying wind conditions (Video 2.6). The 30 mg payload disk structures (28 mm diameter, 55% fill) were released from different heights using a small drone and mechanism shown in Fig. 2.12A. The structure was attached to a thin wire thread hanging from the pin of a solenoid as shown in Fig. 2.12. When a Bluetooth command was sent, the solenoid would retract and allow the structure to fall. Higher drop heights resulted in greater distances since there is both a greater chance for the sensor to be carried by a wind gust during a longer descent and more time for it to travel outward from the drop location when the wind is blowing (Fig 2.13A-C). The highest drop height – 22 m – resulted in the longest distance traveled and also had greater variance: the longer it takes for the sensor to descend, the greater the chance for wind conditions to change across different locations. For example, the point recorded at 100 m was dispersed by a strong wind gust indicated by the large variance (Fig 2.13C). During the experiments, the wind blew in the same direction along which the sensors traveled. We also evaluated the effect of different fills and diameters in the presence of relatively little wind (Fig 2.13D). The structure with a 28 mm diameter and 68% infill traveled greater distances than the one with 21 mm diameter and 55% infill.

Figs. 2.14A-D demonstrate operation of two different temperature and humidity sensors, as well as pressure, and light sensors outdoors integrated with our drag structure. The results show the sensors operating over the course of roughly 4.5 hrs until sunset, at which time the solar cells can no longer harvest sufficient power. The plot also shows that the following morning after sunrise, the device is able to successfully perform a cold start and resumes operation. The raw temperature and humidity data shown in Fig 2.14A,B differs noticeably from the reference air temperature and humidity monitor due to heating from sun exposure. While this measurement could itself be useful for measuring heat on the ground, measurements from the onboard light sensor (Fig. 2.14C) is highly correlated with solar radiation, we use it to compensate for heating. Light compensating

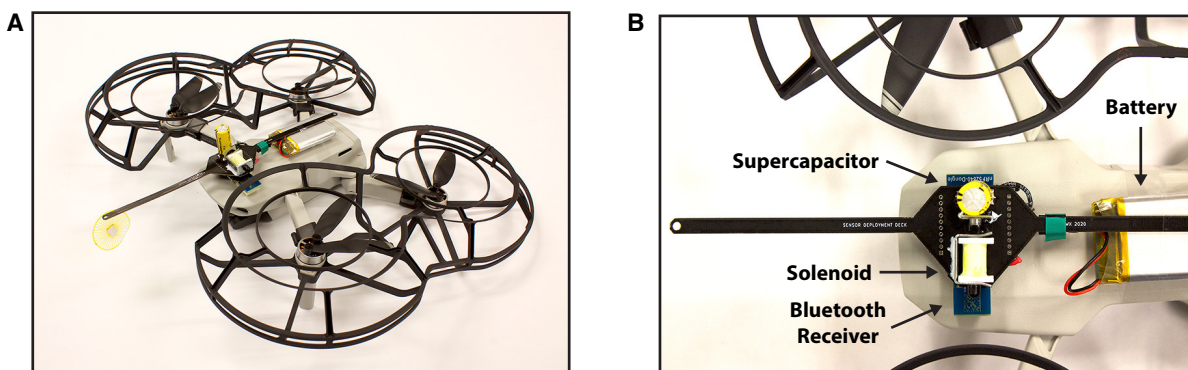


Figure 2.12: Drone mechanism to drop the sensors. A) Small quadrotor with the mechanism used to drop devices from altitudes over 22 m outdoors. B) Drop mechanism consisting of a solenoid wirelessly triggered by a Bluetooth transmission. Upon triggering, the supercapacitor keeps the solenoid retracted for a few seconds allowing the sensor to fall freely. While this mechanism was designed to drop a single device for this evaluation, it could be easily be adapted to simultaneously drop multiple devices.

temperature and humidity outputs for both sensors reduce the error. Fig 2.14D also shows air pressure data compared to a reference sensor, which is also highly correlated with a fixed offset due to the altitude of the reference sensor.

2.2.8 Light-based Temperature and Humidity correction

The weather station shown in Fig 2.15A is designed to accurately measure air temperature and includes radiation shield and 24 hr fan to cycle air through the enclosure. In contrast, our devices have minimal packaging aside from the thin polyimide film and solar cells and can therefore experience significant heating when exposed to direct sunlight for extended periods of time. By integrating a light sensor however, we can measure the effects of solar radiation and compensate for this effect. Fig 2.14C shows the output of the onboard ambient light sensor averaged over 15 min intervals to match the sampling rate of the weather station's solar radiation sensor. This plot as well as Fig 2.15B show that the two outputs are very highly correlated, showing that this sensor gives us accurate information about the amount of solar energy incident on the sensor. We next explore the relationship between the temperature error (difference between onboard sensor reading and the reference) and our light sensor output. Fig 2.15C,D show that this relationship is also highly lin-

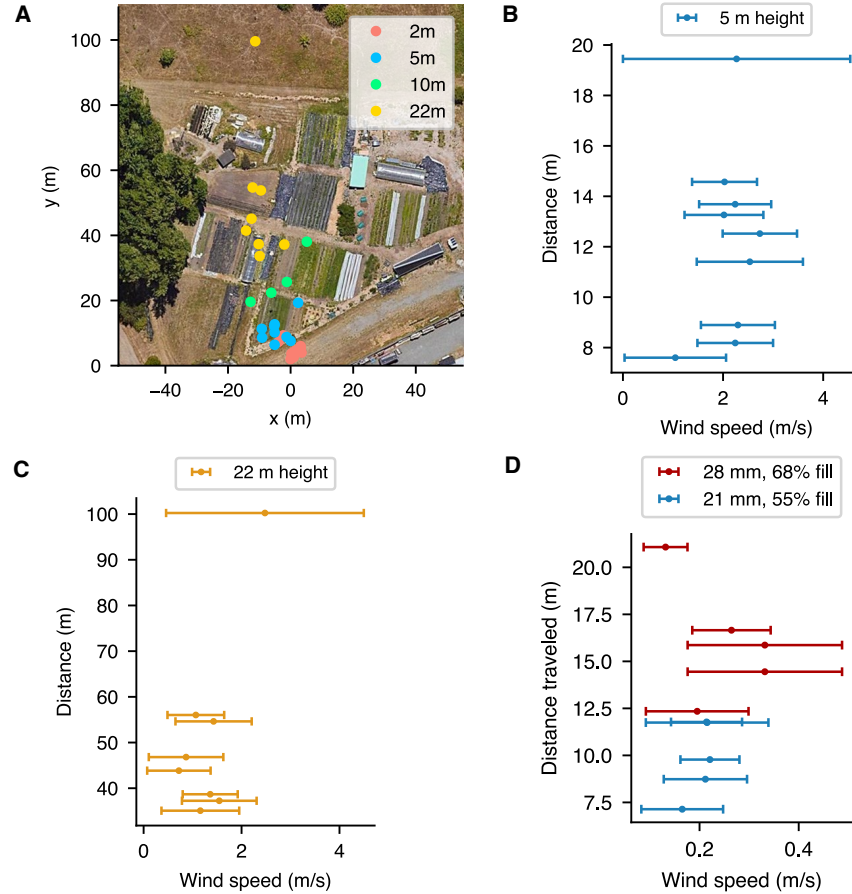


Figure 2.13: Outdoor wind dispersal. **A)** Landing locations in outdoor drop tests from increasing heights. **B,C)** Wind speed versus distance for 5 m and 22 m drop heights. ($\pm\sigma$, $N > 200$ windspeed values) **D)** Wind speed versus distance for two different structures ($\pm\sigma$, $N > 200$ windspeed values).

ear. To compensate for this effect we apply a linear regression over the temperature error for each sensor and illuminance data. Next we attempt to generate the error signal using the illuminance data and subtract it from the raw temperature readings to obtain the compensated output shown in Fig 2.14A,B. The same method is used to correct the humidity data. We note that this simple linear correction method is a proof of concept demonstrating the advantage of including multiple sensors. This method could be further refined with additional data and more complex models considering our light sensor accurately captures the effects of solar radiation. Additionally, while this correction was done with comparison to a high accuracy reference sensor measuring air temperature, future work could explore comparison to a reference sensor measuring ground temperature to

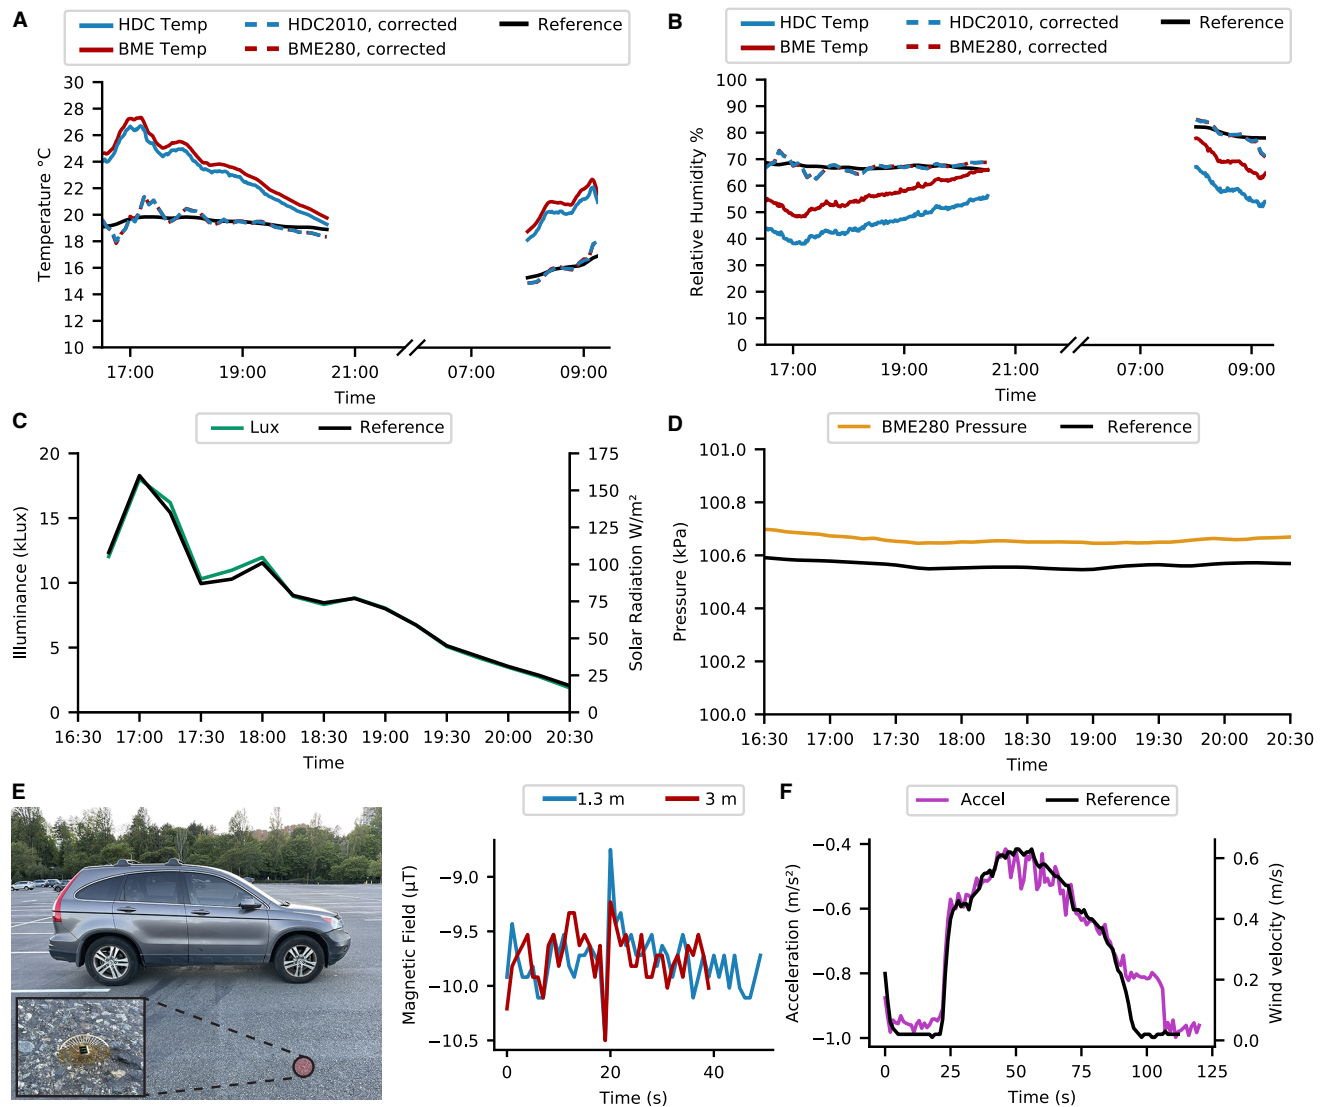


Figure 2.14: Outdoor sensor evaluation. A,B) Outdoor temperature and relative humidity measurements for two sensors over 6 hrs showing successful cold start and correction for sun exposure using the onboard light sensor. C) Outdoor light sensor readings compared to reference solar radiation data. D) Outdoor pressure sensor data compared to a reference sensor. E) Magnetometer readings at 1 Hz from a passing car and experimental setup. F) Accelerometer readings at 1 Hz from a device in a vertical wind tunnel compared to anemometer readings.

build a measurement system for distributed ground temperature measurements.

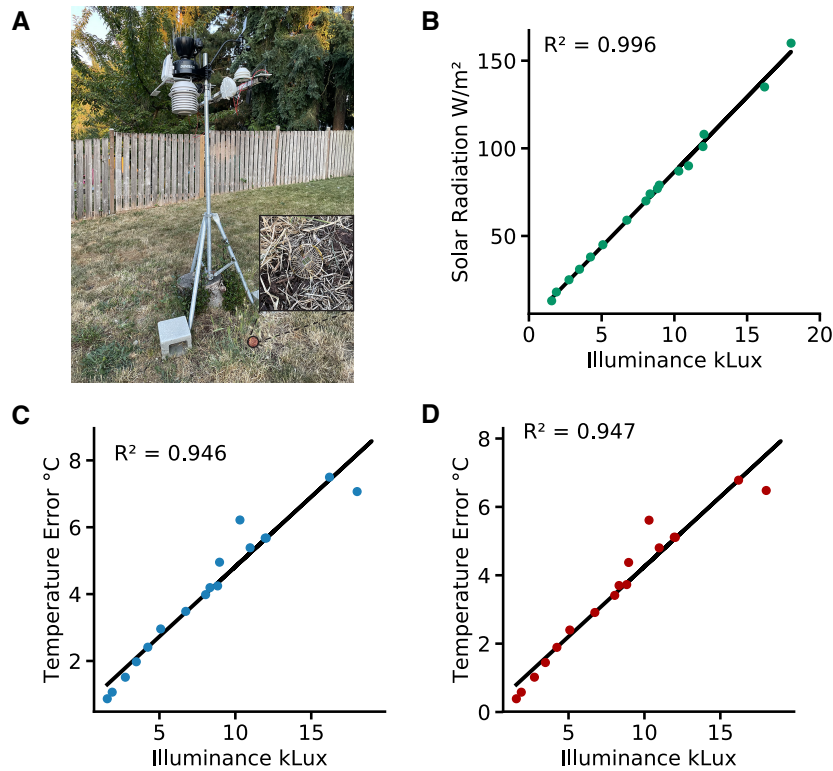


Figure 2.15: Sensor calibration setup using a weather station. **A)** Picture of weather station used for reference temperature, humidity, solar radiation, and pressure measurements and test sensor node placement. **B)** Correlation of small, onboard light sensor with reference solar radiation. **C)** Correlation of small, onboard temperature sensor (HDC2010) with reference. **D)** Correlation of onboard temperature sensor (BME280) with reference.

2.2.9 Magnetometer based car detection

We also integrate a magnetometer sensor into our devices. We place the sensor on the ground sampling at 1 Hz and show the effect of a compact SUV (Honda CR-V 2010) driving by at 16 km/hr. Fig 2.14E shows the raw data when the car passes at distances of 1.3 m and 3 m from the sensor. The plot shows a clear change in signal polarity at the time the car passes by within a radius of 3 m. Larger vehicles like trucks would produce a larger signal and could be detected at longer range, with potential uses for conservation scenarios that require detecting vehicles used for illegal logging or poaching.

2.2.10 Accelerometer based wind measurement

Finally, Fig 2.14F shows 1 Hz readings from an accelerometer integrated with our 28 mm wide 55% filled polyimide structure in Fig 2.4A. The accelerometer is suspended in a vertical wind tunnel with an anemometer (Testo 405i) placed approximately 2 cm above it in the direction of air flow. The air flow velocity of the wind tunnel is steadily increased and decreased while measured on both devices. Fig 2.14F shows the X axis measurement along with the anemometer measurements. We note that because the device was suspended the Z measurements remained roughly constant (acceleration of gravity). The X and Y measurements both varied inversely (due to the accelerometer's orientation) with the flow velocity of the wind tunnel. The results show high correlation between the accelerometer readings and the wind speed, demonstrating this sensor could give information about which direction the device is moving in or potentially serve as a means to collect distributed information about wind direction within a column of air as the devices descend.

To summarize, we present the first system to achieve wind dispersal of battery-free wireless computing devices. Our results show that it is possible to design millimeter-scale, battery-free wireless sensing devices that are programmable and can harvest power from solar energy. The structures introduced here reduce terminal velocity and correctly orient the devices after impact. Though a proof-of-concept with scope for further improvements, our system marks an important milestone towards enabling large-scale dispersal of millimeter-scale wireless devices.

2.3 Discussion

We outline various research directions that highlight the limitations of our proof-of-concept system and provide potential paths forward.

Improving communication and sensing capabilities. One of the limiting factors for range is the self-interference between the transmitter and the receiver. Using full-duplex radios that cancel the interference, can help further increase the range as well as combine the transmitter and receiver into a single device. We prototype a system that can perform this cancellation in Fig 2.16A using a bidirectional coupler (Minicircuits ZGBDC6-362HP+), phase shifter (RF Lambda RVPT0501GAC),

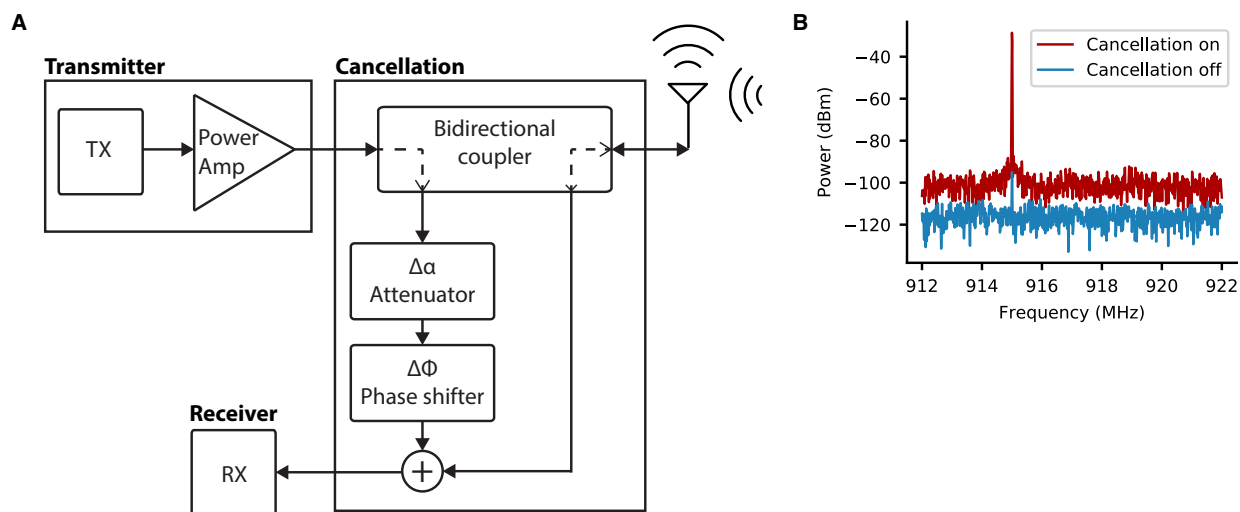


Figure 2.16: Prototype Full Duplex Cancellation of narrowband signals. A) Full duplex cancellation block diagram. B) Results of applying full duplex cancellation to reduce transmitter interference for backscatter.

and variable attenuator (Minicircuits RVA-33). Fig 2.16B shows the approximately 60 dB of cancellation achieved which is in line with prior narrowband cancellation techniques [69]. Our small antenna has a 5 dB loss compared to typical 915 MHz monopole antenna, therefore further optimizations could reduce the antenna loss and also increase the range. Communication range could also be improved by augmenting the sensor nodes to communicate with one another. Small, lightweight envelope detector circuits [60] can allow for a simple downlink to the microcontroller without adding a radio to the device to enable the devices to communicate. Devices could then relay messages, a particularly useful feature in a sensor network tasked with sensing a single event, e.g., an abnormally high temperature indicating a fire, where a triggered sensor sends an alarm signal that propagates through a network. Applications of this platform could further be augmented through integration with other sensing modules like electrochemical gas sensors [128] and cameras. We note that recent miniaturized wireless camera systems [59] currently exceed our max payload of 70 mg, however future work to improve the drag structure or reduce camera mass could enable this in the future.

Adding steerable descent and radio harvesting. One could consider better controlling the placement of the nodes with steerable descent. This would require the addition of a light-weight, low-

power actuator capable of modulating device shape. A shape memory alloy or thin film piezo material like polyvinylidene fluoride (PVDF) are potential technologies that could be used to pull one side of the device inward to change its shape and trajectory. Shape changing could also be useful after landing in order to focus solar cells towards the sun or aim a sensor in a specific direction. However, adding actuation is restricted by the power available on the device, which highlights another limitation of the current work: it can operate only during the day. Allowing these devices to work at night requires either improved energy storage or an alternative harvesting method. The latter could be enabled using radio power harvesting. For example, because the startup circuit requires only a few microwatts of power to operate, a drone flying at a height of 15 m transmitting at 36 dBm at 900 MHz would be sufficient to provide 10 μ W of power at the device. Designing this system requires careful consideration of the frequency choice needed for powering versus backscatter communication to prevent interference and maximize communication range.

Bio-degradable structures and localization. The 3D structure accounts for the largest surface area and could be made of other materials, like paper, that would biodegrade naturally. However, this still leaves unaddressed the problem of the small electronics themselves, which are not biodegradable. These electronics could be recollected by attaching a small magnet to the device. for automated collection. We perform feasibility tests by attaching small N52 magnets to our device with a 30 mg payload (1 mm \times 0.5 mm cylinder weighing 3 mg, 1 mm \times 1 mm magnet weighing 23 mg). We observe they will attract and attach to a bar magnet within 3.5 cm and 4.5 cm respectively. This opens the possibility of driving a vehicle with a large electromagnet across the deployment area to collect them. Collection times could be improved using backscatter localization capabilities to identify the approximate location of each sensor [97]. Another potential for lightweight localization is the addition of a non-linear diode to the device (3 mg), which has been used as part of harmonic radar for insect tracking [113, 27]

Table 2.1: Sensor power measurements and part numbers.

Sensor Type	Part numbers	Sensor mass (mg)	Memory (B)	Mode	On time (ms)	Peak current (mA)	Off current (uA)	Avg Current 1 Hz (uA)	Avg Current 0.1 Hz (uA)
Light	LV0104CS	1.4	540	Cont.	-	2.043	-	-	-
Temp or RH	HDC2010	2.6	688	Cont.	-	2.345	-	-	-
Temp + RH	HDC2010	2.6	728	Cont.	-	2.326	-	-	-
Light + Temp + RH	LV0104CS, HDC2010	4	832	Cont.	-	2.06	-	-	-
Light + Temp + RH + Press	LV0104CS, BME280	12.4	974	Cont.	-	2.117	-	-	-
Magnetometer	MMC56-33NJL	0.8	884	Cont.	-	5.396	-	-	-
Light	LV0104CS	1.4	608	Internal Timer	17.266	2.043	4.6	39.80	8.12
Temp or RH	HDC2010	2.6	616	Internal Timer	4.806	2.345	4.665	15.91	5.79
Temp + RH	HDC2010	2.6	676	Internal Timer	5.354	2.326	4.665	17.09	5.91
Light + Temp + RH	LV0104CS, HDC2010	4	780	Internal Timer	21.27	2.06	4.72	48.44	9.09
Light + Temp + RH + Press	LV0104CS, BME280	12.4	922	Internal Timer	40.78	2.117	4.74	90.88	13.35
Magnetometer	MMC56-33NJL	0.8	832	Internal Timer	5.076	5.396	5.82	33.18	8.56
Temp or RH	HDC2010	2.6	674	External Timer	3.137	2.29	0.295	7.48	1.01
Temp + RH	HDC2010	2.6	734	External Timer	3.627	2.29	0.295	8.60	1.13
Light + Temp + RH	LV0104CS, HDC2010	4	838	External Timer	18.96	2.077	0.295	39.67	4.23

2.4 Videos

Video links:

https://homes.cs.washington.edu/~vsiyer/thesis/wind_videos.html

Video 2.1: Vertical wind tunnel test. Video showing our millimeter-scale, battery-free devices hovering in a vertical wind tunnel with a wind velocity of 1–1.1 m/s.

Video 2.2: Effect of outer ring. High speed video (recorded at 4300 fps) comparing the descent of two 21 mm polyimide disk structures (12 μm thickness, 21 mm diameter, 26% fill) with and without an outer ring.

Video 2.3: Effect of stiffness. A high speed video (recorded at 1253 fps) showing the descent of a 50 mm, 7.5 μm thick disk carrying a 70 mg payload shows it deforming significantly as it falls due to the finite stiffness of the material.

Video 2.4: Unstable descent of solid disks. High speed video (recorded at 1253 fps) showing the unstable descent of an elm seed as well as a 28 mm polyimide disk structure (12 μm thickness, 28 mm diameter, 100% fill, 30 mg payload). Both flip multiple times during their descent making it difficult to guarantee they will land in a defined orientation.

Video 2.5: Self-righting behavior. High speed video showing the self righting behavior of our device (recorded at 1253 fps) as well as a dandelion seed (recorded at 4300 fps) when dropped upside down. Both quickly flip into their upright orientation and begin a stable descent.

Video 2.6: Outdoor wind dispersal. Video showing a slow motion release from 15 m and 22 m. The red circle indicates the position of the sensor as it wafts to the ground.

Chapter 3

INTERNET OF BIOLOGICAL THINGS: A FLYING WIRELESS PLATFORM ON LIVE INSECTS

3.1 Introduction

Mobility in sensor networks has the potential to transform agriculture by enabling smart farming applications including precision irrigation [3] and environmental sensing [80]. Sensor mobility significantly reduces the overhead of manual sensor deployment and upkeep, which remains a major barrier to adoption of smart farms [88]. Drones, which have thus far been the platform of choice for enabling mobility, are severely power constrained and last for only 5–30 minutes on a single charge due to the energy density limits of existing battery technologies [140]. This is mainly because the motors drones use for mechanical propulsion and control are power consuming and unlike digital electronics, cannot scale with Moore’s law. As a result, the majority of a drone’s power consumption is in its propulsion and control systems [36].

This work explores the idea of creating mobile wireless sensors by placing them on live insects. Using live insects such as bees is attractive for two key reasons.

- **Flight Time.** Unlike drones, flying insects use chemical energy stored in fats and carbohydrates, which have a much higher energy mass-density than batteries. This allows for much longer flight times: flies have been shown to fly for hours without food [32], while worker bees spend most daylight hours foraging for nectar and pollen [35, 150] and can fly while carrying payloads of over 100 mg [47]. Further, these animals have evolved to have aerodynamic and musculoskeletal systems that minimize power usage [34, 33].
- **Ubiquity.** Insects are nearly ubiquitous across the planet and adapted to live in diverse ecosystems, making it easy to find a species well suited for a particular environment or application. Moreover, while some are regarded as pests, others are essential to human activities. For example




Figure 3.1: Living IoT platform on a bumblebee. (Left) Live bumblebee (*Bombus sitkensis*) carrying our platform which includes computing, sensing, wireless communication, and self-localization; (Right) Bee flying while carrying our platform.

bees are needed to pollinate many commercial crops, and are in many case intentionally introduced for that purpose [99].

We piggyback on these insects to enable mobility for sensor networking applications including smart farms. Using live insects like bees however introduces two key challenges: 1) they are physically small and can only carry small payloads which severely limits options for power, computation and communication, and 2) we cannot easily control the flight of small insects like bees (see §3.5).

We present *Living IoT*, a novel general-purpose wireless sensing platform that is low-power, low-weight and can support computing and communication operations on flying insects like bees. In order to meet their stringent size and weight requirements, we present a backscatter based communication system that can be achieved with commercially available microcontrollers and has a small and light-weight form factor compatible with insects. Our design includes the capability to compute as well as sample onboard sensors. Additionally backscatter offloads many power expensive components in radios to an access-point (AP) set up near the hive, enabling low power operations.

To address the lack of flight control, our insects localize themselves in the 2D space using RF signals transmitted by access points in the environment. This self-localization architecture, similar in spirit to GPS, is attractive because it does not require the AP to estimate the locations and send them back to each insect. Thus, it can scale well with the number of insects, works at high speeds

and does not require the insects to transmit signals. More importantly, this enables location-based mobile sensors where the insect can associate location information with its sensing data as well as perform sensing operation only when they pass over the desired set of locations. The sensor data can then be uploaded to the AP when it returns within range of the hive.

Achieving self-localization on our living IoT platform is however challenging for three key reasons:

- It should be low-power in nature as we cannot run power hungry GPS radios on the tiny batteries that small insects such as bees can carry.
- Accurate wireless tracking relies on phase information [41, 144, 75] which requires a radio receiver at the insect. This is challenging because radios are power consuming. Further, we are unaware of small low-power off-the-shelf radios that provide I/Q samples of the raw RF signals or CSI data.
- Existing localization algorithms are designed for Wi-Fi chips [144, 41] or software radios that are not constrained by their computational capabilities [126, 148, 97]. In contrast, self-localization at the insect requires algorithms that can operate with a tiny antenna and a low-power microcontroller.

Our design instead eliminates power-hungry radio receiver components (e.g., high frequency oscillators) by using passive operations to perform envelope detection and extract just the signal amplitude. While this design enables the insect to receive at a low-power, it discards the phase information, which is essential for RF localization.

We present a novel technique that *extracts the angle to the AP from the amplitude information output by the envelope detector*. Our intuition is as follows: say the AP broadcasts signals to the insects from two of its antennas. By changing the relative phase between the two transmit antennas at the AP, we can create amplitude changes at the insect's envelope detector over time. These amplitude changes effectively give us multiple equations that allow us to solve for the angle to each AP. Combining the angle information from two APs allows our platform to localize itself on a farm using a passive envelope detector. In §3.3.2, we build on the above intuition and present a low-complexity algorithm that works in the presence of multipath as well as at speeds of up to

9.1 m/s.

Summary of results. Fig. 3.1 shows our hardware including the antenna, backscatter communication, wireless receiver for self-localization and the circuit board that connects the microcontroller and sensors in a lightweight form factor using micro-fabrication techniques (see §3.3.1). We attach our 102 mg platform, including a 70 mg battery, to three common species of bumblebees (*Bombus impatiens*, *Bombus vosnesenski*, and *Bombus sitkensis*).

Our results show the following.

- *Communication.* Our microcontroller-based backscatter design, placed on the bee, can transmit bits using ON-OFF keying at 1 kbps, when the bee returns to the hive.
- *Self-localization.* Across deployments in a soccer field and farm, our envelope detector design on the bees achieves a median angular resolutions of 4.6° at ranges up to 80 m. Ground-truth benchmarking with drones show that similar angular resolution can be achieved even at speeds of 9.1 m/s.
- *Power.* With the on-board 70 mg rechargeable battery, our system could last up to 7 hours while sampling its location once every 4 seconds. We also show the feasibility of fully recharging the battery back at the hive within 6 hours, using RF power.
- *Sensing.* We build prototypes including humidity, temperature, and light intensity sensors. These sensors fit within our 102 mg prototype, enabling mobile sensing using bees.

Contributions. While we are not the first to place electronics on insects — prior work in biology uses electronics on bees to understand their foraging behavior [29, 52, 109] — we show for the first time that insects such as bees can be used to carry general purpose sensors, localize themselves and perform location-based sensing operations.

To summarize, we make the following conceptual, technical and systems contributions. First, we explore the idea that insects can be used in lieu of drones to enable mobility for sensor networks. Second, we present a novel general-purpose platform that is low-weight and can support computing, communication and sensing operations on flying insects. Third, we introduce the first self-localization technique for small insects like bees using a novel algorithm that computes 2D




Figure 3.2: Living IoT architecture. Insect-borne sensor packages can self-locate using transmissions from APs and collect location-dependent data using on-board sensors. Data is stored in onboard memory as the bee flies and data are uploaded to the AP when the bee is back at the hive.

location using only the output of a passive envelope detector.

3.2 Application Scenarios

Smart farming techniques utilize a variety of sensors to measure plant health. For example, moisture and humidity sensors measure the availability of water, light sensors measure the availability and intensity of sunlight, and temperature sensors can help determine whether conditions are optimal for particular crops. Live insects like bees present an attractive option for use as mobile sensing platforms for agriculture. Bees are among nature’s best pollinators and are regularly purchased for use on farms [99], and many fruit crops *depend* on bees for pollination. These insects fly for hours foraging for food, and also fly up to individual plants, which is difficult to do with drones. Further, tracking bees alone could give important insights about pollination which is not available from commercial sensors. This includes pollination patterns that can help maintain genetic diversity [17, 10].

We outline two deployment scenarios: 1) the sensors periodically wake up from a low power sleep mode to sample sensor data and location, store the data in memory, and then return to sleep mode. As the insect goes about its business foraging for food, we get samples of data along its

trajectory. This is ideal for low power sensors such as temperature, humidity, etc. that simply require storing a single value, and 2) we can imagine other applications where we wish to acquire more detailed information at a particular location that is programmed when the bees are the hive. In this case, the system could periodically check its location, and *opportunistically* sample the sensor if it is close to the destination.

3.3 *Living IoT Design*

The idea of using flying insects as a mobile sensing platform raises a number of wireless networking and sensing challenges. At a high level, this requires an ultra-miniature sensing and computation package, memory to store data, a power source, and wireless connectivity for data transfer. Additionally, it requires a wireless localization method capable of connecting the sensor measurements to specific locations.

Our solution consists of two components: a lightweight battery powered electronics package including a sensor, RF switch used for backscatter communication and microcontroller that can be mounted on a flying bee, and multi-antenna access points with dedicated power sources that broadcast the RF signals needed for backscatter and localization. Since large bumblebee hives cover foraging areas of roughly $8,000\text{ m}^2$ [25], we target an operating range of 50–100 m for localization. The range requirements for communication however are much smaller as sensor data can be stored in onboard memory and uploaded when the bee returns to the hive.

To achieve this, our bee-mounted electronics package requires the following two components:

- *Low-power self-localization.* The first requirement is a localization algorithm that runs on the insect mounted platform and computes its location based on wireless transmissions from APs at known locations. By adopting this broadcast architecture similar to GPS, any number of insects can concurrently compute their location. We present a novel algorithm that uses a passive envelope detector to receive the RF signal and compute the location on a microcontroller.
- *Communication.* The second requirement is a form factor compatible wireless communication link to download sensor data. Our backscatter approach has the advantage of requiring only an




Figure 3.3: Living IoT electronics. Complete electronics package including an antenna, envelope detector, backscatter transmitter, and sensor shown on a US penny for scale (left). Custom miniature antenna (right) also shown on a US penny.

antenna connected to RF switches and a microcontroller, all of which can be achieved using off-the-shelf components.

In the rest of this section, we first describe our insect form factor compatible platform. We then present algorithms to estimate the location from the output of the envelope detector. Finally, we describe our backscatter design for uploading data back to the AP when the bee returns to the hive.

3.3.1 Form-Factor Living IoT Platform

Understanding form-factor requirements.

We choose bumblebees as it has been well documented that they can fly while carrying payloads of their own body weight or more. We purchased a commercially available *Bombus impatiens* colony [99]. We progressively added weight to the insects and observe that healthy workers measuring roughly 13 mm in length are capable of controlled flight while carrying loads of approximately 105 mg. When adding weight beyond this limit, they are unable to perform controlled hovering and have difficulty taking off. We perform similar experiments with wild bees specifically *Bombus vosnesenski* (Yellow faced Bumble bee) and *Bombus sitkensis* (Sitka Bumble bee). We note that these insects are slightly larger, 14 and 16 mm respectively and can carry slightly more weight. With the weights noted above, the insects were active and exhibited normal behavior. Thus, we target 105 mg or less for our platform. Of this, 70 mg is consumed by our 3V 1 mAh

rechargeable lithium ion battery [54]. This allows only 35 mg for communication, computing, sensing and self-localization.

Fabrication Method

Our platform consists of four different elements: a microcontroller, RF switches, an envelope detector, and sensor. For each of these we leverage commercially available components available in ultra-miniature packages. The core of our design is the Kinetis KL03z ARM Cortex M0+ microcontroller [123] which is available in an 2×1.61 mm package and weighs only 4.1 mg. We use this microcontroller to sample the output of the envelope detector and the sensor, and to toggle the RF switches for backscatter communication. We use two Skyworks 13314-374LF single pole dual throw switches weighing 3.3 mg each, the first to select between the envelope detector and backscatter, and the latter to toggle between the two backscatter impedances. We construct the envelope detector out of small diodes and capacitors consuming a total area of 7.26 mm^2 . We test different sensors including a TI HDC2010 humidity and temperature sensor as well as an ALS PT19 photodiode to measure light intensity. In total our whole platform weighs 102 mg and measures $6.1 \times 6.4 \text{ mm}^2$.

We must also consider the weight of the substrate such as the printed circuit board (PCB) used to create the circuit that connects these components. The weight of a $8 \times 6 \times 3.175$ mm PCB made of copper clad FR4 with standard thickness and density of 2.6 g/cm^3 is greater than 390 mg. We instead fabricate our own light-weight PCBs by laser micromachining 0.5 oz copper coated sheet of 127 μm thick FR4 and the result is illustrated in Fig. 3.3. We begin by cleaning the copper with isopropanol and placing it on a low tack adhesive [106]. We then raster out the desired pattern of copper traces using our laser micromachining system [24] followed by a low power cleaning raster over the resulting copper traces to remove dust and cut away the excess material. The resulting circuit is approximately 100 μm thick. We then populate the board with components and solder them.




Figure 3.4: Envelope detector performance. Miniaturized envelope detector included onboard Living IoT (left). A plot evaluating its performance compared to a similar design fabricated on a standard thickness PCB with larger component demonstrates miniaturization does not compromise performance.

Envelope detector

To design our envelope detector circuit, we use the 3 stage rectifier design shown in Fig 3.5 that is presented in [104]. In order to minimize size and weight, we instead use small SMS7630-061 zero bias silicon Schottky detector diodes which are available in a 0.3x0.6 mm package and 0201 capacitors each weighing less than 0.5 mg [131] and tune the circuit for operation in the 915 MHz ISM band. To understand whether our miniaturized circuit compromises performance, we compare its performance to the one in [104] made with larger components on standard thickness PCBs. We use a USRP to transmit a tone at 915 MHz and measure the rectifier output voltage with a digital multimeter. Fig. 3.4 shows the rectifier output voltage versus input power for the two designs, demonstrating we achieve comparable performance. Our receiver sensitivity of -40 dBm is similar to the values achieved in [135].

Antenna design

For simple antenna geometries, the frequency of operation directly affects the size of the required antenna. Typical antennas are a half or quarter of the wavelength in size. At 915 MHz, the wavelength is approximately 33 cm, making these antennas significantly longer than a bumblebee (8.5-16 mm) [84]. We explore the design space to achieve a high performance small and light-weight antenna.

First we consider chip antennas that attempt to miniaturize antennas by using materials such as high dielectric constant ceramics. While a 900 MHz chip antenna may only be $3.2 \times 10 \text{ mm}$, they often require a large ground plane to perform efficiently. For example, the recommended ground plane for the above chip antenna is $50 \times 120 \text{ mm}$, and the antenna requires clearance to other components [6]. Additionally, the dense ceramic materials increase the weight making the antenna 288 mg. An alternative is a wire antenna which is a significant fraction of a wavelength. Although these may be longer than the insect [47], the wires could hang off the sides or back. A 15 cm long straight wire hanging off the sides however could introduce other practical issues in the context of insect flight — thin wires can be easily tangled and caught in the insect's wings, legs, or even plants since they don't retain a rigid shape.

To address these concerns we instead design a base loaded whip antenna constructed using the 43 AWG wire. This antenna design is a common solution used in cars and handheld transmitters with electrically short antennas. Rather than having just a straight wire whip which with shortened length presents capacitive reactance, adding inductance at the base cancels this and creates a resonant antenna. We fabricate this antenna using 7 turns of 43 AWG wire with a diameter of 2 mm followed by a straight 10 mm wire. We then apply cyanoacrylate glue (Loctite 416) to the antenna to stiffen it and prevent it from losing its structure. The resulting structure is shown in Fig. 3.3, has a total weight of 4.6 mg, and length of 10 mm which is smaller than the bee and avoids the above issues of entanglement. We compare this to the performance of a 900 MHz monopole antenna. We place each antenna 3 m from a transmitter and find that our lightweight design receives roughly 6-8 dB lower power.

Attachment to Live Insects

We begin by capturing individual bees in plastic jars. For our *B. impatiens* colony we open the door to the hive and place a small cotton swab dipped in sugar syrup near the hive exit to lure bees out one at a time and trap each insect in a jar. For the other species, we capture live bees from a small farm, again by trapping them in a plastic jar while they forage on flowers. After isolating a single insect, we place the plastic jar in a freezer at approximately 0°C for 4-5 min to cold anesthetize




Figure 3.5: Envelope detector circuit diagram. Circuit diagram for our envelope detector design. The circuit consists first of tuning elements for the antenna followed by a three stage voltage multiplier that extracts the envelope of the incoming RF signal.

the insect [70]. At this point the bees typically stop moving allowing for 2–5 min of working time to attach the electronics. We then grip the insect by the sides of the thorax with a pair of tweezers and adhere our electronics package by applying a drop of cyanoacrylate glue and quickly applying a small amount of an accelerator compound (Loctite SF712) to immediately dry the glue and hold it in place. We experiment with attaching the electronics in two locations: the top of the thorax behind the head and on the upper abdomen. We find that the bees are able to fly when the weight is attached in either position, however we prefer the abdomen as it naturally droops down away from the wings when the bee is held off the ground and minimizes the risk of excess glue flowing onto the wing joints on the thorax.

3.3.2 Self-localization of insects

Living IoT’s localization system employs APs which transmit RF signals in the 900 MHz ISM band. At the insect-mounted receiver, due to the power and size constraints, living IoT uses a passive envelope detector connected to a small antenna to receive only the amplitude of the RF signals broadcast by the APs. Unlike an active radio that gives both amplitude and phase, the envelope detector only provides the amplitude of the signal but can do so with small, zero-power passive hardware components [71, 117]. The phase of the signal is however essential to achieve wireless localization.

To address this, we first extract the phase difference of the signals from each of the transmit




Figure 3.6: Self-localization architecture. Our self-localization system uses multiple antennas at the AP and transmits signals with periodically increasing phase delays from the antennas. The result beamforms the output over a range of angles such that the maximum amplitude received on the bee indicates the angle to the AP.

antennas of an AP using the amplitude output by the envelope detector. We then use this phase to compute the angle of the insect with respect to that AP. By using the angles from two APs we can identify the 2D location. Typically bees only fly a few meters above the ground and hence 3D localization is not essential. However, the technique presented in this paper can be generalized to achieve 3D localization by adding an additional AP. Next, we describe each of these techniques in further detail.

Phase from Envelope Detector

Consider a setup in which an AP transmits RF signals using two of its antennas as shown in Fig. 3.6. The two transmissions will travel different distances to the receiver and therefore combine at the insect with a phase difference corresponding to this distance difference. The small antenna at the insect receives the combined RF signals and the envelope detector outputs its amplitude. The key insight here is that the amplitude depends upon the phase difference at which the two transmitted signals combine at the receiver. For example, if the transmitted signals are perfectly in-phase they will add constructively giving the maximum amplitude; in contrast a phase difference of π will cause the two signals to cancel each other completely.

Our key insight is that by intentionally introducing an additional phase difference between the

two AP antennas, we can create amplitude changes at the receiver. We can analyze these changes to estimate the phase corresponding to the angle of the insect from the AP.

To explain this in more detail, consider $x_1(t)$ and $x_2(t)$ to be the signals transmitted from the two antennas. Let us set both these signals to $x(t) = Ae^{j\omega t}$. Assuming no multipath which we will discuss later, the signal at the receiver envelope detector can now be written as:

$$y(t) = |ax_1(t) + ae^{j\phi}x_2(t)| = aA(\sqrt{2 + 2\cos(\phi)})$$

Here a is the signal attenuation, and ϕ is the phase difference between the two paths. Note that the amplitude attenuation difference between the two signals is negligible since the separation between the two antennas is small compared to the distance between the AP and bee.

Now if the AP intentionally introduces a phase difference of θ on the second transmit antenna, i.e., $x_2(t) = x(t)e^{-j\theta}$, the signal at the receiver can be written as,

$$y(t) = |ax(t) + ae^{j(\phi-\theta)}x(t)| = aA(\sqrt{2 + 2\cos(\phi - \theta)})$$

The maximum value for the above equation happens when $\phi - \theta = 0 \text{ mod } 2\pi$. Hence, to get an estimation $\hat{\phi}$ of ϕ , we can let the AP sweep θ from $-\pi$ to π at a constant rate. At the receiver side, the envelope detector simply samples the amplitudes corresponding to each of the θ s. It then gets the sample with the maximum amplitude and infers θ_{max} from the time of that sample. Now, the $\hat{\phi}$ we are interested in is simply θ_{max} . The angle $\hat{\Theta}$ of arrival of the receiver corresponding to the two transmit antennas at the AP can be derived by $\hat{\phi} = d \sin(\hat{\Theta})$ where d is the distance between the two transmit antennas in radians. In our design, this distance between transmit antennas is set to half a wavelength.

A key consideration in the above design is that as the distance increases, noise affects the signal quality. We mitigate the effect of this noise by sampling for a longer time, i.e., the AP dwells on each phase difference θ for a longer duration. The upper bound on this duration however is determined by the motion of the bee. Our empirical results found that a duration of 50 ms per sweep across all the angles, is a good trade-off between the noise level and motion tolerance.

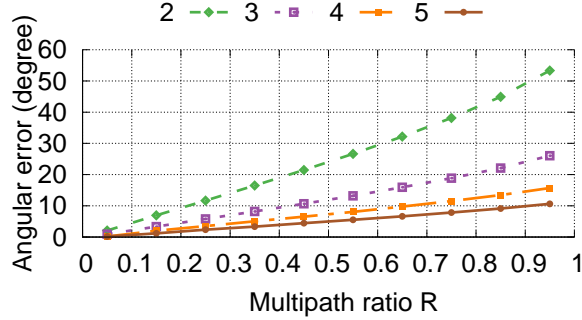


Figure 3.7: Multipath simulation. Simulation results showing the effect of multipath. Different multipath conditions are shown as a ratio to the strongest path $R = \frac{\sum_{i=2}^M a_i}{a_1}$. Access points with 2–5 transmit antennas are simulated and the resulting angular error for different multipath conditions is shown for each.

Addressing Multi-path.

The above discussion assumes that the phase difference $\hat{\phi}$ estimated at the insect can be translated into angle of arrival using amplitude variations. We however need to address the potential amplitude variations due to multipath. Unlike systems like Wi-Fi which operate indoors with mostly no LOS path, our system is designed for outdoor farm use in the natural habitat of insects. For example, in deployment scenarios such as open fields and farms there is a direct strong line of sight. However, we still need to account for the amplitude variations that result from the other paths constructively and destructively interfering with the dominant direct line of sight path. To this end, we utilize more than two antennas per AP to reduce the error in the angle caused by multipath.

Formally, suppose we use N antennas separated by half a wavelength each. Similar to the above two-antenna scenario, we introduce a phase offset of θ_i for the i th antenna. Suppose M independent multipaths exist and the first one is the direct path, the received signal can be represented as, $y_{env}(t) = |A \sum_{k=1}^M a_k \sum_{i=0}^{N-1} x(t) e^{j(i*\phi_k - \theta_i)}|$. Here a_k is the attenuation for the k th path. Let us set $\theta_i = \theta * i$, and sweep θ from $-\pi$ to π . We now get, $y_{env}(t) = A | \sum_{k=1}^M a_k \sum_{i=0}^{N-1} e^{ji(\phi_k - \theta)} |$.

We solve for θ_{max} using the same procedure as in the 2-antenna case and estimate $\hat{\phi} = \theta_{max}$, where the amplitude output by the envelope detector has the maximum value.

The key insight is that, when $\sum_{j=2}^M a_j < a_1$, which is true in most line-of-sight scenarios like the farm, the error in our estimate of the angle with respect to the AP, which is $|\hat{\phi} - \phi_1|$, is bounded.

Moreover, this error decreases linearly as the number of antennas increases. To verify this intuition, we perform a simulation where we compute this error by changing the number of antennas. We repeat this for increasing multipath ratios R , ratio of the sum of the amplitudes of all NLOS paths with respect to the amplitude of the LOS path. Assuming that the angles of indirect paths are uniformly distributed, Fig. 3.7 shows the mean error as a function of these two parameters. The plot shows that the error is less than 10° when using four antennas even if the total amplitude of all NLOS path is 60% of the amplitude of LOS path. With five antennas, we can get a similar error even when this ratio is close to 0.95. This shows that by increasing the number of antennas at the AP, we can reduce the error due to NLOS paths and achieve an accurate angle estimation.

Leveraging insect motion.

We also leverage insect motion to reduce the effects of multipath and reduce the probability of small scale fading. This is specifically useful when the bee is in motion. Since the typical speeds when the bee is in motion are less than 10 m/s, the bee does not move by more than a meter between consecutive 50 ms durations where our AP cycles across the phase values. Despite introducing slight errors because of Doppler effects, our algorithm utilizes motion to improve the accuracy by leveraging spatial diversity: the multipath combination can be significantly different for even a small displacement. Hence, we use exponential smoothing to temporally average consecutive measurements which yields a more accurate result. Formally, the final angle of arrival Θ_t of each bee is calculated as: $\Theta_t = \eta\Theta_{t-1} + (1 - \eta) \sin^{-1}(\hat{\phi}_t/\pi)$, where η is the smoothing constant which we set to 0.8. Note that the computationally expensive \sin^{-1} operation can be offloaded to the access point, as shown in Algorithm 1.

2D localization of insects

To estimate the 2D location of the insect, we employ two APs with four antennas each. Separating the two APs and placing them perpendicular to each other gives the best 2D accuracy. Given the known locations of the two APs, the two angles computed with respect to each AP, uniquely

```

for  $AP_i$  in  $\{AP_1, AP_2\}$  do
   $AP_i$  transmit  $preamble_i$ 
  for  $\Theta = -\pi/2$  to  $\pi/2$  by  $\delta$  do
    for  $j = 2$  to 4 do
      set the phase shift of the  $j$ th phase shifter to  $(j - 1) * \pi * \sin(\Theta)$ 
    end
    sleep for  $(T - T_{preamble})/\pi * \delta$ 
  end
end

```

Algorithm 1: Pseudo-code at the AP.

identifies the 2D location. The bees can either store these two angles or can also calculate their 2D location of using the intersection of the two separate angles of arrival. This intersection procedure can be implemented using a look-up table to minimize the required computation.

Specifically, the two APs intermittently transmit their sweep signal one after another. They are coarsely synchronized using TDMA so that no two sweeps will interfere with each other. We transmit two predefined orthogonal preambles using ON-OFF keying, $[1,0,1,0,1,0,1,0]$ and $[1,1,0,0,1,1,0,0]$ to identify each AP. The preambles are transmitted before every sweep of each AP respectively so that the bee can find the start of each sweep efficiently. The receiver first detects the preamble using a simple state machine, then runs the algorithm twice to get the two angles. Our pseudo-code is presented in Algorithm 2. We note that the computation for our receiver algorithm scales linearly with the sampling rate (which is in the order of kHz), and only simple arithmetic operations are involved. This makes it efficient enough to run on our microcontroller platform.

A key consideration while increasing the number of APs is that the transmissions from each of them have to be time-multiplexed. This increases the delay required to compute the location value which can be challenging especially when the bee is highly mobile. Given that each AP sweeps across various phase values in 50 ms segments, the delay to compute the 2D location is around 100 ms. Assuming a speed of 3 m/s, this translates to a motion of around 30 cm which is within

```

while True do
    receive and store amplitude samples during the last  $3T$  time into memory as  $S$ 
    if preamble1 detected in  $S_{1..2T}$  located at  $i$  then
         $p_1 = \operatorname{argmax}_{j \in (T_{\text{preamble}}, T)} S_{i+j}$ 
         $\text{angle}_1 = (p_1 - T_{\text{preamble}}) / (T - T_{\text{preamble}}) * \pi - \pi/2$ 
        locate preamble2 in  $S_{i+T..3T}$  at  $i'$ 
         $p_2 = \operatorname{argmax}_{j \in (T_{\text{preamble}}, T)} S_{i'+j}$ 
         $\text{angle}_2 = (p_2 - T_{\text{preamble}}) / (T - T_{\text{preamble}}) * \pi - \pi/2$ 
        (If Euclidean position is needed) get  $x$  and  $y$  from  $\text{angle}_1$  and  $\text{angle}_2$  using
            look-up table
        output ( $\text{angle}_1, \text{angle}_2$ ) or  $(x, y)$ 
    end
end

```

Algorithm 2: Pseudo-code of Living IoT platform.

the error of our location estimates and hence does not significantly affect our accuracies.

Localization power consumption.

In order to minimize power consumption, we duty cycle our localization and sensor by periodically performing a measurement and then returning to a low power mode until the next measurement. Sampling the envelope detector and performing computation for localization requires a peak current of 1.6 mA followed by return to a low power sleep mode with a current of 100 μA . If we check our 2D location every 4s, the average power consumption is 138 μA . For our battery capacity of 1 mAh, this results in potential battery life of up to 7 hrs on our rechargeable battery.

3.3.3 Backscatter Communication Design

Backscatter requires three components: an antenna, a switch to modulate the reflected signal, and a control signal for the switch with the desired bits [61, 86], two RF switches, and the microcontroller

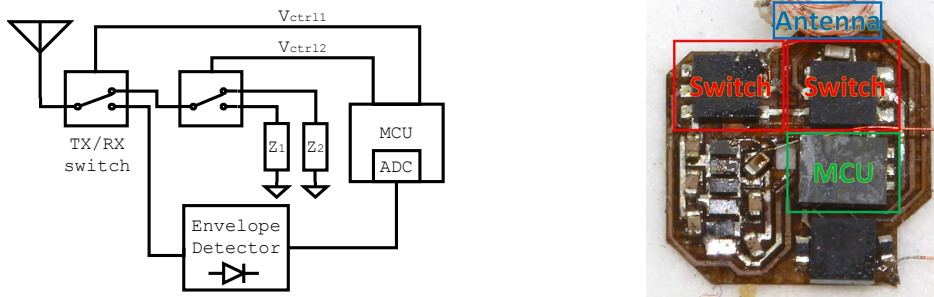


Figure 3.8: Living IoT backscatter hardware. A block diagram of the backscatter hardware used in Living IoT is shown on the left. The corresponding light-weight hardware prototype with these components is shown on the right.

as shown in Fig. 3.8. We connect the antenna to the input port of the first switch which selects between the envelope detector and the second switch used for backscatter. We use the second switch to select between two impedance states to create the backscatter signal. In order to minimize extra components we simply use the open and short impedance states which we implement by leaving one port disconnected and connecting a 51 pF capacitor to ground on the other port.

At a high level, the AP transmits a continuous wave signal from one of its antennas, which the living IoT platform backscatters with the sensor data when the bee is back near the hive. We use ON-OFF keying modulation to encode bits using backscatter. Backscatter however presents the challenge of needing to receive the weak reflected signal in the presence of the high power RF signal from the AP. To address this using a technique called subcarrier modulation [61, 72]: we use the microcontroller to generate a square wave at a rate of 2 MHz which shifts the frequency of our backscatter signal 2 MHz away from the high power signal from the AP. By doing this, the AP can then filter out the transmitted signal thereby significantly reducing the receiver noise floor and improving our communication range. We then modulate this signal by toggling it ON and OFF to produce an ON-OFF keying pattern to encode bits.

To implement this with minimal power consumption, we disable all unnecessary peripherals on the device and reduce the clock frequency to 8 MHz. We then generate a 2 MHz square wave using a timer module in PWM mode at a fixed 50% duty cycle for sub-carrier modulation. We find that enabling and disabling the timer module in software incurs delays, so we instead take advantage

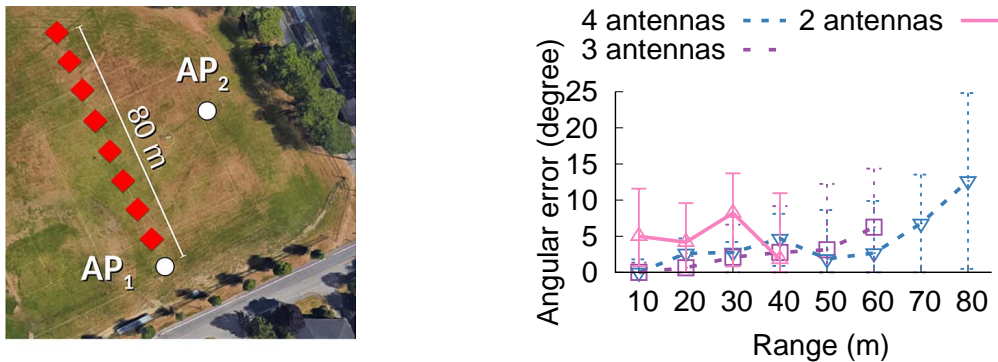


Figure 3.9: Stationary self-localization angular accuracy. The accuracy of the localization system is measured for multiple antennas at distances of up to 80 m in an open field (left). The angular error versus range is plotted for each antenna configuration (right).

of our first switch: we use a lower rate control signal from a GPIO pin on the microcontroller to toggle the first switch to the envelope detector and disconnect the backscatter signal. Using this method we can send data at a rate of 1 kbps which is sufficient for our application where the sensor readings and angle values we need to transmit are only 1-2 bytes.

We note the following about our communication design.

- **Backscatter range.** In our design the insects store the sensing data and upload it to the AP via backscatter when the insect is back at the hive. This only requires a range of a few meters. One can, in future designs, consider increasing the backscatter range using coding to hundreds of meters [135].
- **Downlink and MAC protocol.** We reuse the envelope detector for downlink communication from the AP to the bees. Our design uses a simple MAC protocol where, once the bees are back at the hive in the night, the AP queries each of the insects one after the other, using the downlink. The insect then uses backscatter to respond with the sensor data and corresponding location information.
- **Effect of interference.** A concern with using an envelope detector is that it is not frequency selective and can therefore see signals across a broad range of frequencies at 900 MHz. This however does not significantly affect our performance for two key reasons. First, our application is designed for smart farms where 900 MHz transmitters including cordless phones, LoRa and RFID

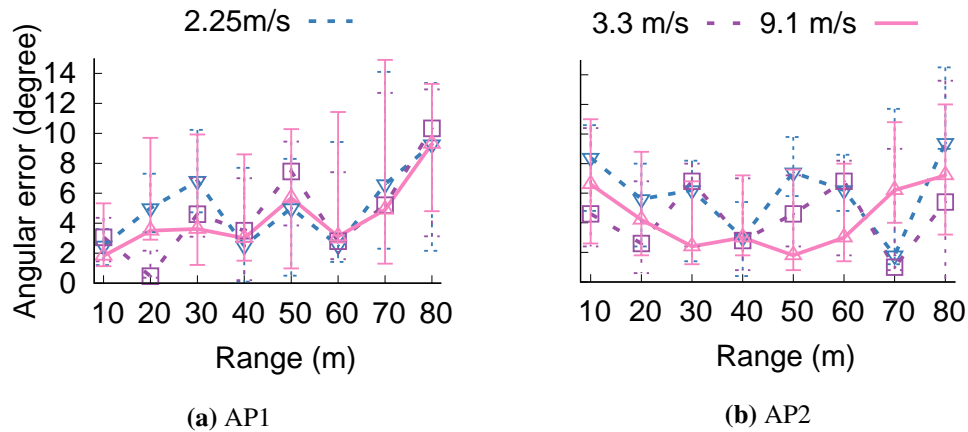


Figure 3.10: Mobile 2D localization angular accuracy. The receiver is attached to a drone flying at different speeds to evaluate the impact of mobility. The resulting angular errors versus range for each flight speed at each AP location are shown.

readers are currently uncommon. Second, most 900 MHz wireless deployments including LoRa and SIGFOX are designed for sensitivities of less than -118 dBm. These lower power signals do not register at our envelope detector which has a -40 dBm sensitivity.

- **Power requirements.** Our micro-controller based backscatter requires a peak current of 1.8 mA when transmitting at 1 kbps. Unlike localization which is performed periodically, data is only uploaded once upon returning to the hive and does not otherwise consume power. Offloading 10 sensor measurements and the corresponding angle data would only require running our backscatter transmitter for 32 ms and is therefore not a concern from a power perspective.

3.4 Evaluation

We evaluate various aspects of our living IoT platform.

3.4.1 Self-Localization Accuracy

We first evaluate our low-power localization algorithm using a static deployment of living IoT in different environments. Since it is difficult to get the ground truth location with a free flying bee,

we use a drone to evaluate the 3D accuracy of our technique at different speeds. Then, we place the living IoT platform on a live, wild caught bee that can fly freely within a plastic enclosure and run our localization algorithm at various points across a farm. Finally, to benchmark the effect of different wing speeds, we run experiments with a robotic wing of similar dimensions.

Stationary experiments.

For our experiments, each AP consists of one USRP-N210 connected to a four way power splitter followed by three phase shifters, each of which introduces a phase shift controlled by an NI myDAQ digital-analog converter. Along with the original signal, the four outputs are amplified to 28 dBm by a Qorvo RF5110G power amplifier and then connected to four 2 dBi monopole antennas separated by 12 cm each. The cable lengths are carefully calibrated so that no extra phase offset is introduced.

Soccer field deployment. We first conduct our range evaluation outdoors on an open soccer field measuring approximately 100×100 m. We place an AP at one end of the field and move our receiver along a straight line away from the AP up to a maximum distance of 80 m. Fig. 3.9 plots the angular error as a function of distance up to the point at which the receiver had insufficient SNR to decode. The plots show that our low-power envelope detector platform can compute the angle of the AP up to 80 m from the AP with four antennas. This range could, in future designs, be improved by increasing the sweep time at AP which in our implementation is 50 ms. Also, the angular error improves with the number of antennas at the AP. Further, with four antennas the range of the design increases because of antenna diversity gains.

Farm deployment. We placed one 4-antenna AP at the center of two perpendicular edges of an 90×120 m farm. We then place the living IoT platform at multiple locations around the farm. We repeat this experiment multiple times at each location and compute the x and y axis errors as well as the 2D location by combining the angles from the two APs. Fig. 3.11 plots the ground truth locations and the CDF of the error at all locations. The plots show that the median 2D localization error is 1.9 m. For context, while prior AoA work localizes devices at sub-meter resolutions [126], the range at which the experiments were conducted is limited to less than 10 m.

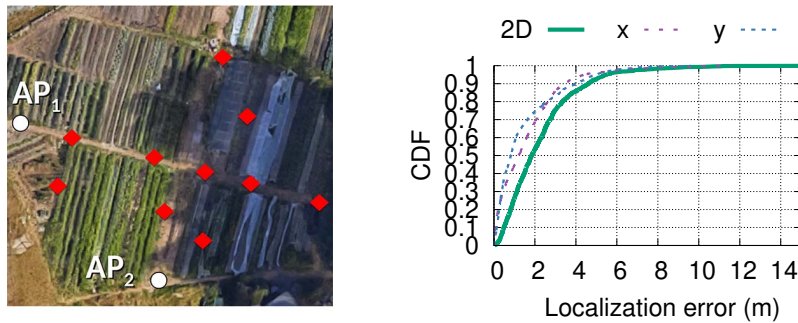


Figure 3.11: Localization accuracy on farm. The 2D localization accuracy is evaluated by measuring signals at 10 points across a small farm (left) with two APs. A CDF of the localization accuracy is shown to the right.

As the distance increases, as expected, the same angular error results in a larger localization errors. This meter-scale resolution however is sufficient for our smart farm application.

High-speed experiments.

Since we cannot control the bee motion it is difficult to run systematic accuracy experiments at different speeds on an actual bee since we do not know the ground truth at long ranges in outdoor environments. Instead, we place our platform on a DJI Phantom 3 drone. We place the first 4-antenna AP at one end of the soccer field and the second 4-antenna AP as shown in Fig. 3.9. We run experiments at three different speeds using the Phantom 3's wireless controller. We then use the drone's flight record which contains its GPS, altitude, accelerometer and compass data as the ground truth. Fig. 3.10 plots the angular accuracy at three different speeds. We plot the results for both the APs. The average angular accuracy at different distances are similar to the stationary angular accuracies observed in the previous experiment. The fluctuations in error across all distances are due to the changing multipath over time, however the exponential smoothing technique offsets this making some results better than the stationary scenario.

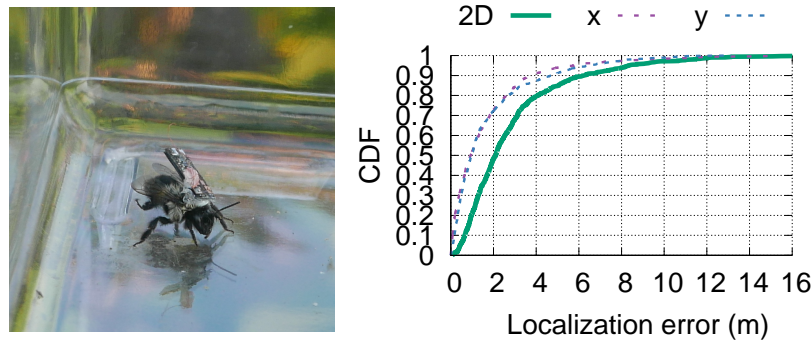


Figure 3.12: Localization accuracy on bee. The 2D localization accuracy is evaluated by measuring signals from an envelope detector attached to a live bee kept in a plastic enclosure at 10 points across a small farm (left) with two APs. A CDF of the localization accuracy is shown to the right.

Wild bees observations.

Over the course of our outdoor experiments we observe many wild bees in both the soccer field and the farm. In total we observe 3 species including honeybees (*Apis mellifera*) as well as two species of bumblebees (*Bombus vosnesenski* and *Bombus sitkensis*). In areas of the farm with flowers in bloom, we observe as many as 20 or more in a single 1 m^2 area. Even on the open field we observe many wild bees foraging in weeds with small flowers (e.g., dandelions). We note that while they typically fly at average speeds of 1-5 m/s when going longer distances [102, 83], they settle in one region slowly hovering or landing at individual flowers.

Experiments with bees.

Next, we evaluate our self localization algorithm with our living IoT platform placed on a live bumblebee as shown in Fig. 3.12. We conduct this experiment with the bee in a plastic container with a volume of $30 \times 20 \times 20$ cm, where the bee freely moved around and flew. We run experiments again in the farm deployment where we placed the two 4-antenna APs as shown in Fig. 3.11. We then place the plastic container with the mobile bee at different locations in the farm. The location of the plastic container was taken as the ground truth for computing the localization error. Fig. 3.12 plots the CDF of the localization accuracy with our living IoT platform on the bumblebee. This as

expected shows that the accuracies are similar to prior experiments and demonstrates the feasibility of self-localization on a live and mobile bee.

Effect of flapping wings.

While a drone serves as an excellent platform to perform systematic experiments with known ground truth data, its flight mechanics are different from that of an insect. So, we next evaluate the effect of flapping wings on localization accuracy. Unlike propeller driven drones, insects use flapping wings to generate lift. Since the envelope detector is mounted close to the wings of the insect which move periodically and cause vibration in the body of the insect at the flapping frequency we next evaluate how flapping wings affects the signal at the envelope detector.

In order to isolate the effect of flapping from flight motion of a real insect in a systematic manner, we instead attach our envelope detector and antenna to an insect scale robot design with flapping wings that is inspired by [153] and shown in Fig. 3.13. The robot has a wingspan of 35 mm which is similar to a bumble bee (32 mm) and are designed with hinges to mimic the wing kinematics of real flying insects.

The robot's wings are driven by 2 piezoelectric bimorph actuators which we control to flap the wings at different frequencies to determine whether the wing motion itself negatively impacts our localization performance. We then incremented the phase difference at the transmitter antennas by discrete steps of thirty degrees and monitored the change in the amplitude of the signal received at the envelope detector. Fig 3.13 plots the raw received signals at the output of the envelope detector. The figure does not show a noticeable degradation in SNR due to wing motion or vibration of the body as a whole. Further, the different amplitudes created by the phase changes at the AP, appear intact at the output of the envelope detector and are largely independent of the flapping frequency of the insect's wings.

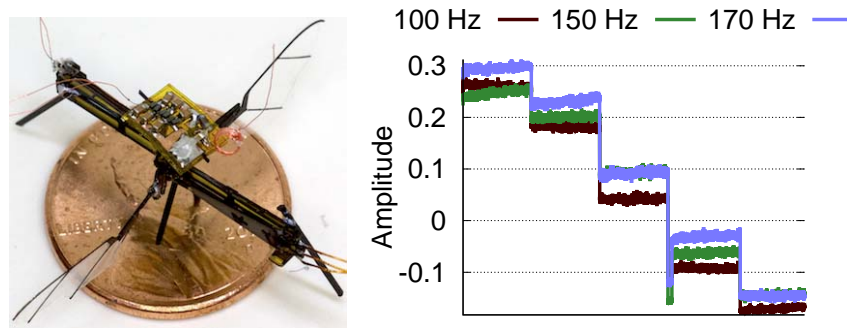


Figure 3.13: Effect of flapping wings on envelope detector. A small envelope detector is shown attached to an insect-scale flapping wing robot (left). The effect on the envelope detector at different flapping frequencies is evaluated (right).

3.4.2 Backscatter Evaluation

We implement our access point using two USRP software radios. For the transmitter, we configure a USRP to transmit a continuous tone at 915 MHz and connect its output to a Qorvo RF5110G power amplifier which outputs to a 6 dBi patch antenna. We use the same model of patch antenna for our receiver and place it parallel to the transmitter antenna separated by a distance of 0.5 m. We set the receiving USRP to a frequency 2 MHz away from the transmitter corresponding to the subcarrier generated by backscatter. We then apply a band pass filter in software to remove unwanted noise.

We evaluate our setup in our application scenario of the bee uploading its data when it is near the hive. We evaluate this first with a dead bee placed at fixed locations, and then on a live bee in a plastic enclosure. We ensure that our evaluation board is light enough to enable the bee to fly and walk around the enclosure; we glue the board to the top of the insect’s abdomen. We place our AP approximately 1 m from a bee hive separated from the AP by an enclosure consisting of 2 layers of insect netting. We configure the AP to continuously transmit 100 bits and measure the error rate at 20 locations within 2 m of the hive. We place the hive on a 1 m high table and include points on every side of the hive to simulate bees taking a variety of trajectories to approach the hive. We plot the uncoded BER in Fig 3.14. The plot shows that the live bee has a similar performance to




Figure 3.14: Backscatter performance. The uncoded BER on a static and mobile bee is shown (left). These values are low and comparable to prior backscatter designs [147] and the bee can upload data when it is back at the hive. The impact of the small antenna performance is shown to the right in a comparison of BER versus range.

that of the dead bee demonstrating that our backscatter system can operate in the presence of bee motion. We note that the uncoded BERs are similar to prior backscatter works [147] and can be further reduced by applying error correcting codes [86, 135].

Finally, to understand the effect of the antenna without complicating factors of a live insect and its motion, we evaluate in a stationary case, with both our lightweight wire antenna as well as a standard monopole antenna. We place our AP at a fixed location in an outdoor environment and configure our backscatter device to transmit a known sequence of 200 bits at a rate of 1 kbps. We then move it to increasing distances away from the AP and record the received bits at each location. Fig. 3.14 plots the uncoded BER versus distances and show that our light-weight antenna can achieve low BERs upto 5 meters. Beyond 5 m, the monopole antenna performs better than the light-weight antenna. This 5 m range is however sufficient for our application where the bees upload the data when they are back at the hive.

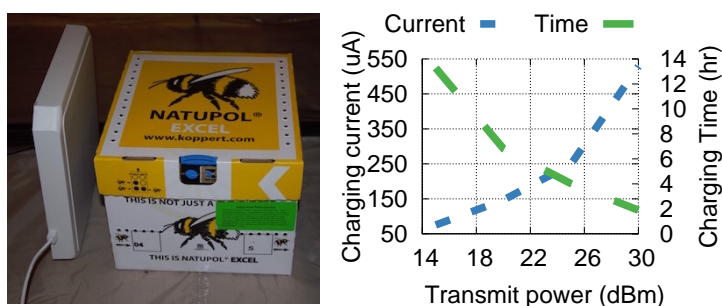


Figure 3.15: Recharging batteries using RF harvesting. A patch antenna transmitting power is placed next to a beehive contained within a cardboard box (left). The charging current and time are shown to the right for different RF transmit powers.

3.4.3 Power harvesting Feasibility

Recharging batteries using RF

Since certain insects like bees return to a single colony or hive we can use this as a central charging point. To evaluate the feasibility of this approach, we first measure the attenuation of 900 MHz signals through a real bumble bee hive (Natupol, Koppert) [99]. We place a 900 MHz 6 dBi patch antenna AP in different locations near the hive including above, below, and on either side. We then place our platform on the opposite side. We find that placement of the transmit AP below the hive causes significant attenuation as these commercial hives contain a package of sugar water below the hive to provide a food source. We also find that placing the AP above the hive results in the lowest path loss of approximately 15 dB, while the side and back are 3 and 6 dB worse respectively. Fig. 3.15 plots the measured harvested power versus transmitted RF power from the top of the hive. The results show that a 20 dBm AP could charge a 1 mAh battery in about 6 hours.

Solar harvesting for battery-free platform.

Insects such as bumblebees are most active during daylight hours. This makes solar power harvesting a particularly attractive option as these creatures naturally fly in outdoor environments with abundant sunlight. We evaluate an 8 mg 3x3 mm photovoltaic cell [93] shown in Fig. 3.17(a) by measuring the output of the cell under a microscope light [4] with controllable intensity. We mea-




Figure 3.16: Living IoT sensors. Close up images of the sensors included in Living IoT. A combined temperature and humidity sensor is shown (left) along with an ambient light sensor (right).

sure the output current and voltage of the cell and plot the results for illuminance values ranging from 1000-20,000 Lux in Fig. 3.17(b). At 1000 Lux which is representative of an overcast day the PV cell can harvest 1 μW of power and up to 50 μW on clear sunny days. This shows the potential for a battery-free design that could replace the 70 mg battery with a solar cell and small storage capacitor.

3.4.4 Sensor Peripherals

We integrate three sensors: temperature, humidity, and light intensity. These sensors are commercially available in small packages compatible with our weight budget and operate at low power. For performing temperature and humidity measurements we use the TI HDC2010 IC which weighs 3 mg and includes the physical sensor as well as an integrated ADC and digital interface as shown in fig. 3.16. This chip is capable of providing high accuracy measurements while consuming as little as 0.55 μA of current for measuring both sensors once per second. Additionally, its low power sleep mode allows for greater power savings by leaving it inactive until our algorithm detects we are in a location we wish to sense. We also measure light intensity using the ALS PT19 ambient light sensor. This chip weighs only 1.5 mg, but does not include integrated readout electronics. Instead, we use the built in ADC on our microcontroller to periodically sample the analog output.

During operation, living IoT must also be able to log its sensor data until returning to the hive to upload it. The humidity and temperature values require 11 bits each, while the light intensity

requires 12 bits. Our microcontroller has up to 32 KB of onboard flash memory. Using 2 bytes for each measurement, 1 byte for each angle, is a total of 4 bytes for a sensor measurement and 2 angle values. This would allow us to store measurements once every 5s for over 10 hrs. This is sufficient to cover daylight hours during which bees are active and foraging. Alternatively, we can leverage our localization technique to selectively log data at higher rates near specific target locations.

3.5 Related Work

Tracking bees. Prior biological research has explored the problem of tracking bees to understand their behavior and help explain the decline in their worldwide population. [139] attaches bees with laser light-activated microtransponders which can be detected by laser readers at distances of up to 10 mm from the reader. To know when the bees enter and leave the hive, a 10 by 10 mm plastic tube walkway was attached to the hive entrance, with two laser readers at the top of the tube. Intel and the Australian Commonwealth Scientific and Industrial Research Organization (CSIRO) are exploring the possibility of attaching RFID tags on bees to track them at a range of less than 10 cm allowing them to notice when the bees enter and leave the hive [52, 109]. Backscatter has also been used with large dragonflies [141]; it uses a custom silicon design and also does not have RF localization capabilities.

Radar based systems [113, 27] use a pulse radar with a 25 KW transmit power at 9 GHz as the reader and use a passive analog diode on an insect that creates harmonics at 18 GHz which the reader uses to track the insect at much longer distances [113]. The key challenges with this design are three-fold: 1) the analog design with the diode always creates a 18 GHz harmonic and hence cannot support more than one bee, 2) these passive analog diodes do not support general-purpose sensors on the bee, and 3) the bee cannot localize itself and has no computing capabilities.

Finally, [47] uses a 200 mg active pulse radar transmitter at 100-200 MHz powered by a small battery on a large species of European bumblebee. These transmitters send small pulses which the reader uses to track the bee. This does not satisfy our requirements for five reasons: 1) across our experiments the bees common in our environment, were unable to lift more than 105 mg, 2) the bees do not have the localization information and so it does not support self-localization,

3) it is unclear if the design will scale to large numbers of bees since they all must transmit the corresponding pulses, requiring a scalable MAC protocol, 4) transmitting high power pulses at the bee would limit other operations that can be performed and 5) it does not carry general-purpose sensors.

In contrast to this prior work, we make the following contributions: i) design the first low-power self-localization technique for flying insects, ii) present a general purpose platform that enables computing, communication and sensing on aerial insects and, iii) demonstrate for the first time that insects such as bees can be used to carry general purpose sensors and thus enable mobility, in lieu of drones.

Controlling insect flight. Researchers have shown how to control the motion of larger insects such as beetles, dragonflies and Locusts [1, 120, 91, 103]. An interesting future research direction would be to develop low power flight control for small insects like bees. This however is challenging since while the nervous system of larger insects like beetles is well understood, bees are at least an order of magnitude smaller in size and it has not been shown that their flight can be controlled. Achieving this for small insects like bees is an open problem and, if solved, can augment our work.

Power harvesting from large insects. Prior work has tried to harvest power from large insects such as moths and beetles [125]. [2] utilizes a piezoelectric transducer which converts the vibratory motion created by the insect's flight into electrical power to harvest 7–60 μW from beetles and hawk moths. [18] harvests 0.8 mW from moth vibrations using a magnetic induction generator. Finally, [62] harvests around 1 μW from the chemical energy stored within a moth's hemolymph. These systems however use moths and beetles that are more than 10 times larger than bees. Instead we show RF harvesting to recharge the battery and the feasibility of using solar cells to enable battery-free designs.

Bio-inspired aerial robots. The robotic community has spent the last two decades in the design of insect-scale aerial robots [107, 87, 37] that are similar in size to houseflies and mimic the wing propulsion of insects. Despite significant research [16, 153, 89, 40, 44], these robots are largely tethered to a wire to power and control them, since they consume hundreds of milliwatts of power

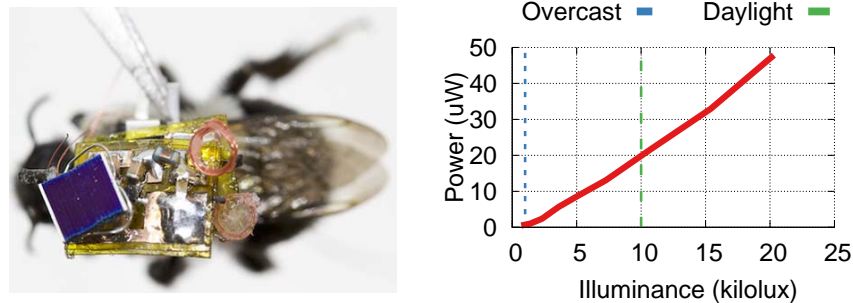


Figure 3.17: Solar harvesting for battery-free platform. A small solar cell is shown above a bee (left). The power produced by this solar cell in different light conditions is shown to the right.

for the mechanical propulsion and cannot carry batteries given the weight requirements. While recent work shows the feasibility of wireless power using lasers at the range of a meter [89], providing hundreds of milliwatts of power at 80–100 m is challenging and difficult to scale beyond a single aerial robot. Our idea instead is to leverage biological insects which can be thought of as efficient biological machines that provide flight and piggyback communication, computing and sensing on top of them.

Size-constrained sensor systems. [64, 21, 38] designed die-stacked sensor platforms that have the key building blocks of miniaturized sensor nodes, such as data transceivers, energy harvesters, power management units, and digital logic circuits. This approach requires custom ICs which limits availability to researchers elsewhere as well as programmability. We take an alternate approach by designing a system using commercial off the shelf components based on a microcontroller that is programmable to create a more modular platform. In addition, we also demonstrate a self-localization technique using only an envelope detector.

[85, 105, 155] design localization systems for sensor networks with a focus on decreasing cost, power, or improving accuracy; however these techniques all require active radios and cannot easily be scaled down to the size and power requirements for use on insects like bees.

Our prior work [97] designs a sub-centimeter sized backscatter device based on a microcontroller and localizes it using LoRa backscatter [135, 50] at a software radio using non-linear optimization on I/Q samples. In addition to being focused on a novel mobile insect application, our

design differs from [97] in four key ways: First, our localization algorithm is not based on backscatter and does not occur at the software radio. Instead we design a localization algorithm that runs on the low-power device using the output of an envelope detector, without access to the I/Q samples. Second, our backscatter design does not use LoRa transmissions, which significantly simplifies our design. Third, a key limitation of the design in [97] is that it cannot currently scale to more than one to two devices. In contrast, our self-localization algorithm is similar in spirit to GPS and hence can scale to a large number of insects at the same time. Fourth, unlike non-linear optimizations that cannot effectively run on our low-power microcontroller, we design a low-complexity algorithm to estimate the 2D location at the Living IoT platform.

3.6 Discussion

This paper explores the use of insects in lieu of drones to enable mobility for sensor networks. Making this vision pervasive, however, requires addressing additional challenges.

- **Weather dependency.** Bees hibernate during winter in cooler climates. This however correlates with plant growth and increased activity on farms in warmer seasons.
- **E-waste.** Piggybacking on insects could lead to electronic waste being scattered around the farms when the insects eventually live out their lifespan. There are three approaches to addressing this problem: i) ensure that the electronics are removed before the expected lifetime of the bees, ii) the electronics can be localized for a while even after the bees die which can be used for cleanup purposes and iii) use biodegradable electronics [53, 137] in the design of living IoT.
- **Fabrication.** Our current prototype requires manual fabrication and attachment to the insects. However we use commercially available parts which allows for easy scaling of the electronics fabrication and our process for gluing the electronics to the insects is similar to the process for attaching tracking markers to bees in commercial hives. Additionally, researchers have also shown that insects can survive common microfabrication processes such as deposition of conductive material in a vacuum chamber [129, 130] and performing surgeries at different stages of an insect's life cycle. This suggests potential approaches for mass attachment of electronics or fabricating

devices on insects themselves. Further work on implanting devices at different stages of an insect's life cycle also has potential to improve fabrication [111].

- **Camera sensing.** A future research direction is to integrate cameras [95] with the Living IoT platform. This can be useful for smart farm applications like canopy monitoring. Centeye image sensors and cameras such as the Himax HM01B0 offer a potential path for achieving such a camera-based sensing system within our weight/power budget.

- **IACUC requirements.** Finally, we note that while working with insects and other invertebrates is not governed by the Institutional Animal Care and Use Committee (IACUC) policies of our institution and insect consciousness is poorly understood [92], we do our best to follow the three Rs of animal research [116]. We minimize use of insects by benchmarking each aspect of our system using drones to simulate flight and robots to simulate flapping wings. For our experiments we perform no surgical modification, we simply attach weight to the exoskeleton and we observe no significant changes in behavior after the procedure. Additionally we use only a small number of insects for our experiments and remove our electronics package after completing experiments. We also release wild caught insects back in the area they were captured after completing experiments.

Chapter 4

WIRELESS STEERABLE VISION FOR LIVE INSECTS AND INSECT-SCALE ROBOTS

4.1 Introduction

In animals, vision provides a means to perceive the world at a distance, providing crucial information used for navigation, communication, finding food, mating, and detecting threats. Animals of all sizes from insects to humans have evolved visual systems specific to their needs. In contrast, for man-made robots, vision has been limited to larger systems and integrating cameras onto small, insect-scale robots remains challenging [132]. In this work, we seek to design a robotic vision system that rivals that of similarly-sized natural systems, such as those found in insects.

A naïve approach to building a vision system for small robots would be to leverage advances in miniaturization made for smartphone cameras. This would seem an intuitive choice due to their small size and megapixel resolution [156]; however the processing and energy requirements needed to support these cameras necessitates powerful processors and prohibitively large batteries. We instead look to biology and explore the trade-offs evolution has made in the visual systems of insects to inform the design of vision for insect-scale robotics.

The fraction of resources devoted to vision varies significantly in different animals. These range from simple eyespots composed of two cells found in zooplankton larvae [65] to complex eyes capable of seeing color and high resolution found in larger animals like humans. In an adult human, the mass of the eyes and visual cortex of the brain account for around 0.6% of body mass [23, 43]. In contrast, the visual systems of insects like flies have lower resolution than large animals like humans but can consume 2.5-13% of body mass [8]. Further, supporting these large organs represents a significant energetic cost — the retina of a blowfly alone consumes 8% of its resting metabolism [100]. The energy cost to support the sensitive photoreceptors has led many

insects to develop high visual acuity in only small retinal regions [48, 152, 49].

To balance this trade off of smaller visual field, insects have evolved to move their visual system independent of the body through head motion. Blowflies[49], moths [152], mantids[76], locusts[142], and many others move their visual system to expand their field of view by dynamically scanning, or to maintain focus on moving objects (e.g., prey or potential mates). In contrast to moving the whole body, this adaptation allows these animals to gather more visual information in an energy efficient manner [77]. Additionally, this added degree of freedom is used by some insects to guide steering as the first part of a motor program that can precede body turns[15], infer depth or motion information[76, 152, 49], or to orient their gaze in a direction independent of their movement direction[14].

Electronic image sensors follow a similar trend to these biological eyes as simpler, lower resolution image sensors consume less mass, less energy, and require less computation. In many scenarios of interest for small mobile cameras however, image resolution is a limiting factor. Rather than compromising resolution or field of view, we explore the approach used in nature and design a mechanically steerable vision system that imitates head motion. We incorporate an ultra-miniature piezoelectric cantilever actuator and microfabricated lever arm to steer the camera. By incorporating steering, our system provides much higher image resolution than is possible with a wide-angle lens covering the same visual field. We further show that our actuator is substantially more efficient than moving the body of an insect or robot, minimizes impact on battery life, and allows for maintaining fixation on moving objects.

Using this approach, we develop a fully wireless mechanically steerable vision system in a form factor small enough to mount on the back of a live beetle or a similarly sized terrestrial robot. On live insects, a wireless first-person view represents a new capability that could be used in studies of insect behavior such as how they perceive and interact with each other as well as with their environment, outside of controlled lab settings. For insect-scale robots, wireless vision provides a rich source of information about the shape and texture of the environment that is commonly used in larger robots, but is challenging on smaller, resource constrained platforms.

Building steerable wireless vision at this scale is however challenging in practice due to the

extreme size, weight and power (SWaP) limitations. In addition to being small, insects and small robots have a severely constrained payload capacity. For small robots, a heavy payload requires more energy to maintain speed, reducing their operational time. Similarly, the addition of a heavy payload can limit an insect’s ability to move. For these reasons, insect-mounted sensing and control systems have not previously demonstrated wireless vision [28, 119, 145, 60]. Furthermore, no prior insect-scale terrestrial robots have included wireless cameras. Larger bio-inspired aerial robots like the 10 cm Delfly micro which weighs 3 g have demonstrated wireless cameras, but their camera sub-system weighs more than a gram and consumes 200 mW for camera operation alone[30]. Similarly, the small vision systems developed for wireless capsule endoscopy robots require hundreds of milliwatts [22, 31, 39]. Designing a light-weight, low-power vision system could therefore have broad applications to improving battery life across multiple domains.

The requirements of these prior vision systems are also prohibitive for payload limited robots since their camera power consumption necessitates large, high drain batteries and exceeds the power required for locomotion in small terrestrial robots[154, 66]. While recent advances in low-power wireless systems have shown that it is possible to significantly reduce the power consumption for video streaming using backscatter[95, 118], these solutions currently weigh tens of grams and have a limited communication ranges of around 5 meters before significant degradation in image quality. Other recent works on flapping wing insect robots have demonstrated small form-factor light-weight cameras but required a wire tether for computing and communication[9, 90].

In this work, we demonstrate insect-scale steerable wireless vision (see Video 4.1). Specifically, we make the following key contributions:

1. We designed a 84 mg wireless camera system that can stream 160×120 monochrome video at 1–5 frames per second (fps) to a Bluetooth radio (e.g. smartphone) from distances of up to 120 m away using 4.4–18 mW of power as well as have the ability to send back commands to control the steerable head via Bluetooth in real time (see Fig. 4.2).
2. We mounted this camera on a 35 mg mechanically steerable “head” capable of panning the camera over a range of 60° . This mechanism includes a piezoelectric actuator driven by an

onboard 96 mg boost converter circuit to provide the required high voltage signals. Further, to enable low-power operation, the actuator is designed to hold its angle for over 1 min after being powered off.

3. We demonstrated real-time video streaming from the back of live insects (see Fig. 4.1A). We performed field experiments with two species of freely walking live beetles, demonstrating that triggering image capture using an on-board accelerometer achieves operational times of up to 6 hours with a 10 mAh battery.
4. We also use this vision system to demonstrate the smallest, power-autonomous terrestrial robot with a camera (see Fig. 4.1B). This robot consists of a 1.6×2 cm light-weight frame and two vibration motors, and reduces the power required for locomotion (9.25 – 33.3 mW) to the same order of magnitude as its Bluetooth communication (16 mW). Our results demonstrate that to capture images across a larger field of view from the robot, moving the mechanical head is 26–84 times more energy efficient than moving the whole robot.

This work presents a wireless steerable vision system that can be carried by insects like darkling beetles and small robots. Making this compatible with smaller insects like bumblebees and flies with payload limits of 100-200 mg and 10-50 mg respectively [60], requires another order of magnitude reduction in both power and weight. Achieving this is the next frontier for insect-scale vision systems.




Figure 4.1: Steerable wireless camera on a beetle and a small robot. **A)** Wirelessly steerable camera system attached to a live beetle. **B)** A wireless, power-autonomous terrestrial robot with steerable vision. The camera streams video to a smartphone, which can command the robot to move and pan camera. **C)** Exploded view showing all components of the system including the Bluetooth, camera and optics, robotic head, high voltage electronics, and battery. **D)** Diagram showing the components of the steering mechanism. **E,F)** Close-up diagram of hinge motion as the camera pans.

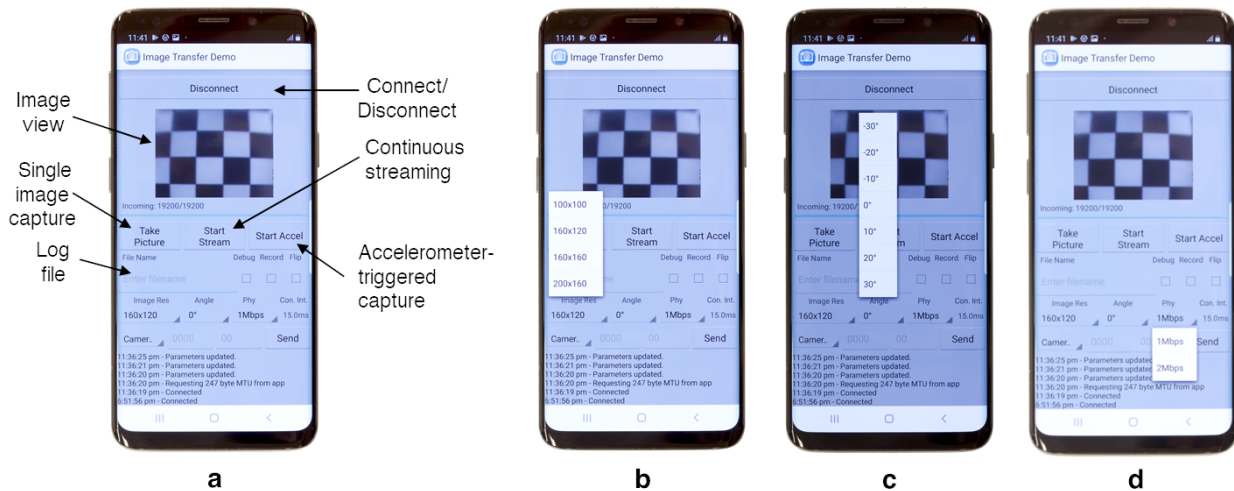


Figure 4.2: Smartphone interface. A) Image showing the smartphone GUI for recording video and controlling the camera angle on a Samsung Galaxy S9. The GUI includes a menu for selecting the device to connect to, an image viewer, and controls to operate the camera in a single image capture mode, continuous streaming mode, or accelerometer-triggered capture mode. The GUI also includes data logging and debugging features and the following menus for setting parameters: B) menu for selecting between four preset image resolutions, C) menu for setting the robotic head angle, and D) menu for selecting Bluetooth bit rate.

4.2 Results

4.2.1 Low-power mechanically steerable head

Insects move their heads using complex systems of muscles allowing for multiple degrees of freedom adapted for specific tasks [133]. Replicating this with synthetic parts raises significant challenges in terms of the added weight, power, control, and manufacturing complexity required to create a system with many actuators. We observe however that well-studied insects such as *Manduca sexta* (hawkmoths) and *Calliphora erythrocephala* (blowflies) primarily move their heads along one axis with greater speed and range of motion [152, 49]. Therefore, we adopt a simplified model of one degree of freedom motion.

Due to their small size, insects can carry limited payloads — hawkmoths have difficulty carrying more than 1 g [28]. We therefore seek to minimize the total system weight and limit our design space to sub-gram components. Further, inspired by our model insects, hawkmoths and

blowflies, which move their heads $30 - 90^\circ$ [152, 49], we set this as our target angular range for camera steering. Finally, we could optimize for power without compromising angular range by selecting an actuator that can hold its angle without continuous power input. This would allow the camera to capture visual data at an angle without draining the battery

Prior works have shown that piezo-electric actuators are an excellent candidate in terms of weight, size, and power-efficiency[89, 63]. Potential alternatives at this scale include shape memory alloy actuators[74] or small electromagnetic coils [12] but both are inefficient due to thermal dissipation. We model these piezo actuators as capacitors with low-leakage providing a storage element for charge and allowing them to hold their state. Theoretically, the energy required to steer a piezo is simply the energy needed to charge the capacitor. Since, this capacitor is on the order of several nanofarads and the required piezo voltages are on the order of 100s of volts, the energy consumption is less than 1 mJ. Additionally, because of their ultra-low leakages, they can stay at the steered angle for several seconds to minutes. The mechanism is explained in detail below.

4.2.2 *Microrobotic arm mechanical design*

The microrobotic arm is designed around a bimorph piezo actuator, which consists of a trapezoidal layer of carbon fiber (Toray M46J) measuring 1.75×12.6 mm with the same shape $127 \mu\text{m}$ thick layers of PZT (Piezo.com, T105-H4NO-2929) on either side. The design is inspired by the mechanisms used in flapping-wing micro-robots[89, 63]. Without a load, applying a high voltage (250 V) causes approximately $\pm 300 \mu\text{m}$ of deflection at the tip of the actuator. This sub-millimeter deflection is converted to the camera angle change using a transmission mechanism, shown in Fig. 4.4A, that amplifies this to a larger angle change orthogonal to the tip. The actuator is placed in a frame constructed from layers of carbon fiber and polyimide film (DuPont Kapton), which anchors its base at one end and contains the transmission mechanism that is affixed to its tip at the opposite end. Because the camera's width would cause it to interfere with the actuator's motion, the camera assembly is mounted 3.7 mm away from the hinge at the end of two carbon fiber rods. These rods are oriented orthogonal to the hinge in order to produce a right-left panning motion while keeping the structure oriented lengthwise along the insect.

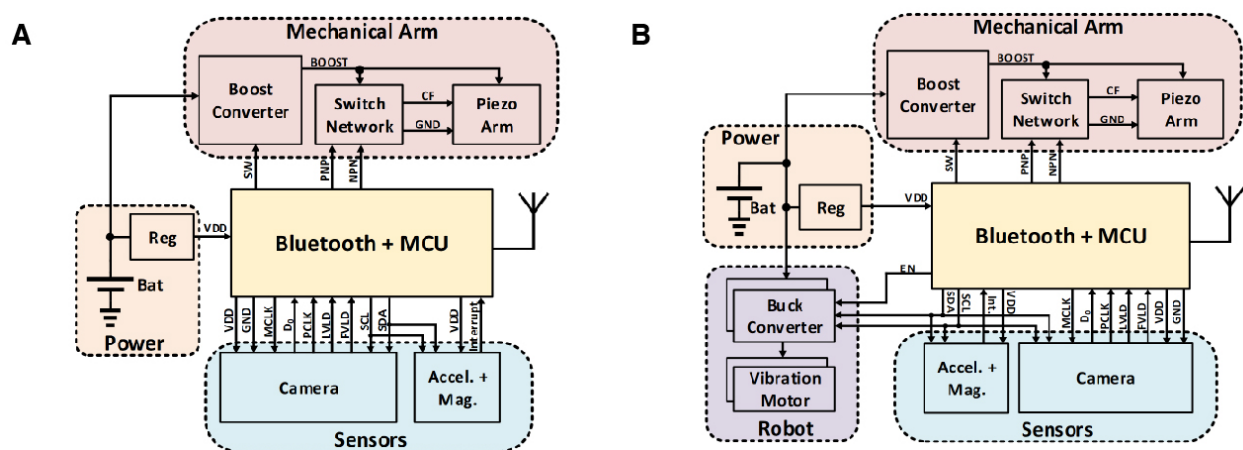


Figure 4.3: Full system block diagram. A) Block diagram of the backpack carried by the insect. B) Block diagram of the robot. Both of these systems are composed of a Bluetooth enabled microcontroller (MCU), robotic head and power electronics, camera, motion sensors and power system. The robot includes two vibration motors and two buck converters to control their voltage and speed via Bluetooth. A voltage regulator supplies power to the MCU and boost converter from the LiPo battery. The MCU both provides a wireless interface and control signals for sensors and vibration motors. The MCU also controls the boost converter and switches to produce the output voltages required to move the piezo.

To prevent the maximum angle from being limited by the output force of the piezo, only the camera and lens assembly are attached to the hinge to minimize the mass it needs to move. However, this requires a flexible interconnect capable of carrying the nine signals for power, clocks, and data from the camera to the Bluetooth chip (Fig. 4.3). We experimented with different wire thicknesses and found 43 AWG wires attached to either side of the camera were too stiff and reduced the range of motion by 10° or more in both directions. Thinner 51 AWG wires attached to one side of the camera did not limit the motion and let the camera move over a 60° range (see Video 4.6). Additionally, a flexible interconnect made of $24\ \mu\text{m}$ thick copper-coated polyimide (DuPont Pyralux AC121200E) with 9 parallel $30\ \mu\text{m}$ wide traces cut so each could flex separately provided an alternative solution with reduced potential for entanglement of individual wires.

All components were custom fabricated using a laser micromachining process. The piezo and un-cured carbon fiber were cut by the laser and laminated together using a heat press. Similarly,

the frame and transmission were produced using the same method; during lamination in the heat press, the layers were adhered together using thin sheets of a modified acrylic adhesive (FR1500 Pyralux) placed between them. The final assembly was produced by cutting away extraneous support structures, manually folding the frame into shape, placing the actuator in the frame, attaching the camera, and securing all joints together using CA glue.

To generate the high voltage signals (>200 V) required to drive the piezo actuator, we designed and fabricated a custom lightweight boost converter circuit described in detail below.

4.2.3 Boost converter circuit

The boost converter circuit used to generate the high voltage drive signals is shown in Fig. 4.5A. The control signals to switch the transistor were provided by the microcontroller (MCU) included on the Bluetooth chip. Through circuit simulation, we found that the power consumption of the boost converter was dominated by conductive losses resulting from the parasitic resistance of transistor J1 in the ON state and the coupled inductor, as well as the rise and fall times of the pulses applied to the transistor. By using a GaN-FET (Efficient Power Conversion EPC2012C), we reduced the ON resistance from 450 m Ω to less than 110 m Ω . Decreasing the ON resistance below this value and increasing the drive strength of the MCU switching signal V_{SW} did not improve performance, which suggests further performance gains are limited by the inductor's wire resistance (Coilcraft LPR4012). Moreover, output voltage is determined by the duty-cycle and load current at the boost converter output – decreasing boost output load current increases voltage. The output load current is composed of a diode leakage current, PNP and NPN leakage currents, and the current passing through the R3 resistor when the PNP switch is ON. Therefore, to maximize the output voltage, we used a low leakage diode and BJT transistors and an $R3 = 68$ M Ω resistor. Increasing R3 above this value does not increase the voltage noticeably.

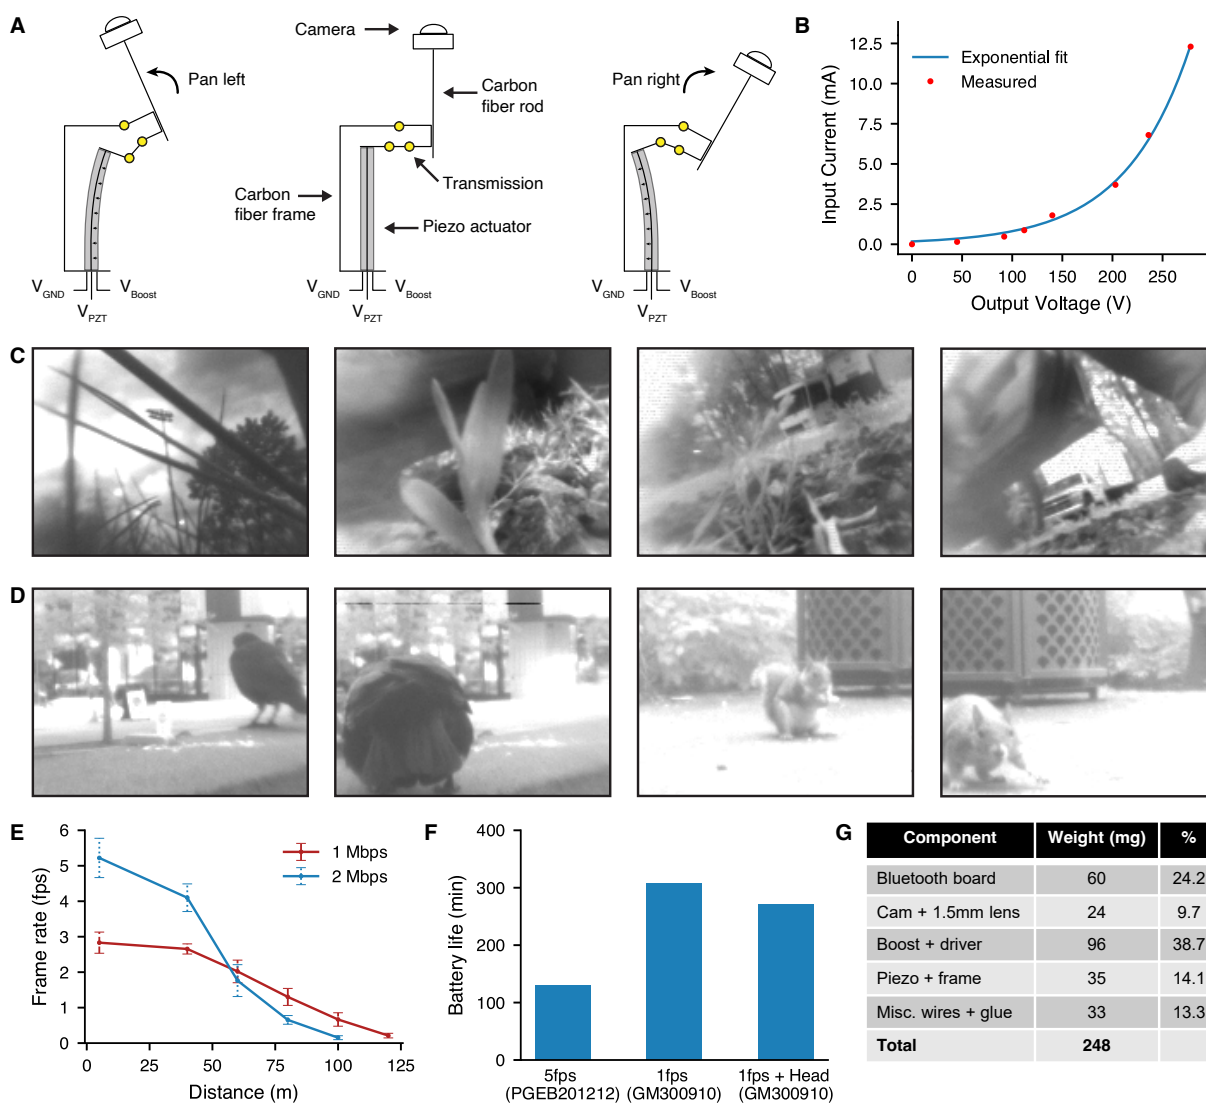


Figure 4.4: Evaluating the wireless camera and arm performance. **A)** Labeled diagram showing piezo actuator operation. **B)** Low voltage input current versus high voltage output generated by the boost converter. **C)** Sample 160×120 images showing the performance of the camera using a 1.5 mm diameter 1 mm focal length lens (Edmund Optics 43394). **D)** Sample 160×120 images showing the performance of the camera using a 3.8 mm diameter 2.33 mm focal length lens (Panasonic EYLGUDM128). **E)** Frame rate versus line-of-sight range from the insect. Error bars indicate mean $\pm 1\sigma$ ($N=10$ frames); the rate remains constant until the sensitivity limit of the wireless link, and then begins to degrade. **F)** Battery life when continuously streaming 160×120 images at different rates, with and without robotic head motion, and different batteries. **G)** Weight breakdown of the Bluetooth, camera and robotic head, as well as each component's percentage of total weight.

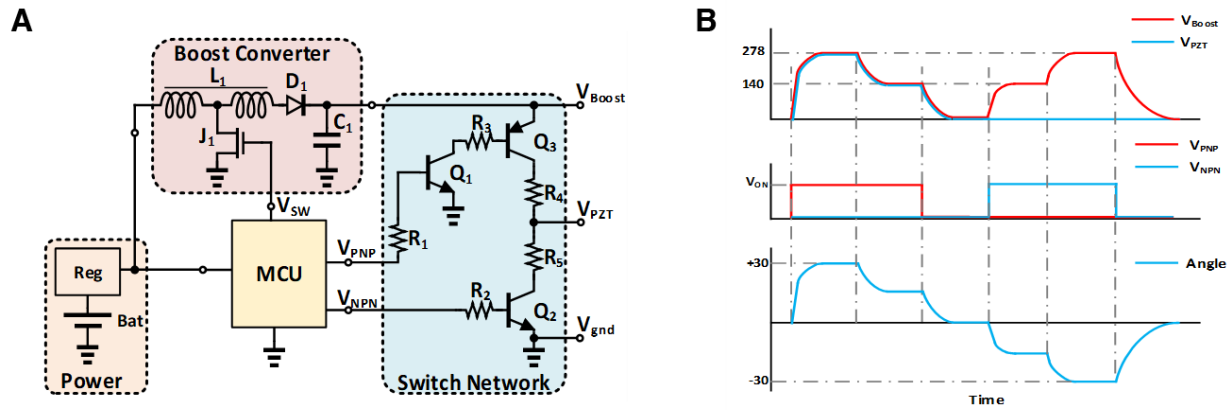


Figure 4.5: Boost converter circuit schematic and drive signal waveforms. A) Boost converter and switch network used to drive the piezo arm. Power is provided from a LiPo battery and control signals are generated by the MCU. B) Boost converter output V_{Boost} is connected to the top piezo and V_{PZT} , is connected to the center carbon fiber layer of the actuator. Voltage is varied in 5 steps. The center plot shows the switch control voltages. The PNP transistor Q_3 first shorts V_{PZT} to V_{Boost} , for positive angles. Both transistors Q_2 and Q_3 are then off for 0° angle. NPN Q_2 is turned ON shorting V_{PZT} to V_{gnd} , while V_{Boost} increases for negative angles. Output angles are shown in the bottom plot.

The current required to run the boost converter was measured at eight output voltages. Fig. 4.4B shows that producing linearly increasing output voltages from the boost converter requires exponentially higher current. To minimize energy consumption, the boost converter should therefore be operated at the lowest output voltage necessary to produce the desired angle motion. The bimorph actuator was constructed with both piezo elements poled in the same direction, meaning a positive voltage across the top piezo would cause the structure to bend in one direction, and a positive voltage across the bottom piezo would cause bending in the opposite direction. Maximum actuator deflection, and therefore angle change of the camera in either direction, is achieved by applying the highest possible voltage to one piezo and 0 V to the other. The circuit shown in Fig. 4.5A accomplishes this by using a switch network to connect the center layer of the bimorph to the top or bottom piezo layer. Fig. 4.5B shows the waveform and illustrates how using this technique, the boost voltage is only set to the minimum value required to achieve a particular angle, and the direction is set by toggling transistors Q_2 and Q_3 . The resulting total boost converter current was

measured for each angle using a digital multimeter (DMM, Fluke 287) for each angle output.

This circuit was fabricated using the same laser fabrication method noted previously. An 127 μm thick sheet of FR4 coated in 0.5 oz copper was covered in a layer of ink that the laser rastered away in regions around the desired pads and traces. The exposed copper was then etched in a solution of Copper(II) Chloride, and the remaining ink was washed away using isopropanol. The circuit was then manually populated with components and reflow soldered at 285°C. The combined boost converter and switch network weigh a total of 96 mg, dominated by the weight of the 26 mg coupled inductor.

To steer the head we charge up the piezo capacitance to our desired voltage using the boost converter and immediately turn it off. Because of the leakages of the transistors and diodes connected to the boost converter output, the piezo capacitance discharges over time. To increase the discharge time, we used ultra-low leakage diodes and transistors. This enables the piezo to hold its position without being actively charged. We identified the ON-resistance of the switching transistor in the boost converter circuit as the factor limiting efficiency, and we use a low resistance Gallium Nitride transistor (GaN-FET) to reduce the energy consumption at high voltage outputs. The resulting circuit is able to move the head over an angle of 60° (see Video 4.5), which is within the range observed for live insects such as hawkmoths and blowflies [152, 49]. The actuator can also move at angular velocities of over 1000°/s without damage which is similar to a blowfly which can move its neck at 1000°/s [14].

The time it takes for the piezo to return to its initial angle depends on the leakage of the components used and the total capacitance composed of piezo capacitance and external capacitance. To increase this time constant, we place a large external capacitor at the output. Because this element would be able to store more charge it could hold the angle for a longer time; a large capacitor which extends the fall time however also increases the rise time requiring longer to charge it up to the maximum voltage and therefore consuming more energy. For our application we choose a capacitance of 660 pF to minimize energy while still taking over 1 min to return to zero degrees after being powered off. This capacitance could be increased for applications that require holding the angle longer, for example a 100 nF capacitor increases the time for the angle to return to zero

after powering off to more than 5 min.

4.2.4 *Low-power wireless camera*

When designing the wireless camera we again consider the weight constraints discussed above to fit within the payload capacity of insects. This is challenging because while commodity image sensors used in smartphones are as small as 5 mm in width [156], the batteries needed to support them are on the order of 5 cm or more and weigh over 100 g. Small batteries have significantly lower capacity, which also has a nonlinear dependence on the amount of current it supplies. For example, a sub-gram LiPo battery only lasts 5 min with a load of 50 mA but can last for over an hour with a load of 10 mA [138].

Existing wireless camera solutions require significant amounts of power that would limit battery life. For example, Wi-Fi based streaming cameras consume hundreds of milliwatts [96]. Similarly, the small vision systems developed for wireless capsule endoscopy robots [22, 31, 39] use small image sensors that are 3.4 mm in width but still draw 34-50 mA of current for the image sensor alone (see Fig 4.6). When combined with wireless communication, this leads to a total current of 65-115 mA. Because endoscopes do not have weight constraints, they can use heavier, higher capacity batteries. For example, the most recent untethered capsule endoscopy systems use larger 50 mAh batteries (Varta CP1254) weighing 1.6 g [31, 39], with which they achieve a 26 min battery life. Similarly the commercial PillCam uses 2 Energizer 399 batteries weighing a total of 1.6 g [78]. Both would only last 5 min on light-weight batteries.

Recent works have also used low-power backscatter to reduce the power consumption of wireless streaming cameras [95, 118]. These systems however can only achieve limited wireless ranges, which would limit their usability on freely moving insects or robots. To allow for robust operation outdoors in the field we therefore target a line-of-sight range of 100 m. The power and range limitations of existing cameras highlights the need to design a custom wireless camera solution. Ideally, in order for this custom design to easily scale and adapt to a variety of different robotic applications we use commercially available chips rather than custom designed application specific integrated circuits (ASICs).

Image Sensor	Active Pixels	Color	SNRmax (dB)	Responsivity (V/Lux-sec)	Power Consumption (mW)	Dynamic Range (dB)
ON Semi MT9M114	1296x976	RGB	37	2.24	146	70.8
OmniVision OV6922	328x250	RGB	42	1	66	66
OmniVision OV5640	2592x1944	RGB	36	0.6	364	68
OmniVision OV6930	400x400	RGB	38	3	48	68
Misumi MO-S588	510x492	RGB	48	-X-	350	-X-
AWAIBA NanEye GS	320x320	RGB	-X-	6.4	83	60
Himax HM01B0	320x320	RGB, Mono*	38.7	5.6	2	70

* This image sensor comes in two versions of RGB and Mono both of which have the same characteristics.

Figure 4.6: Comparison table for image sensors. Comparison of commercial image sensor specifications including their resolution, power consumption and other characteristics.

We take a clean slate design approach and determine the minimum necessary components for a streaming camera system. We then systematically identify the smallest size and lowest power commercial off-the-shelf (COTS) components that meet these requirements. We begin by selecting an ultra-low-power 320×320 pixel image sensor (Himax HM01B0) that measures less than 2.3 mm wide and weighs 6.7 mg (see Table 4.6).

Our target applications scenarios in which the vision system will ride aboard small insects and robots introduces another requirement not found in biological systems: a bi-directional wireless communication link to record data from this image sensor and control the actuator. The primary design considerations for this component are an uplink bit rate of 1-2 Mbps for video streaming, a small form factor and long-range. Instead of designing a custom radio and communication protocol, we leverage the Nordic NRF52832 Bluetooth 5.0 chipset which supports up to 2 Mbps bit rates and is available in small (3×3 mm), highly integrated packages, weighing only 6.8 mg.

Building this system however requires addressing the following three systems challenges. First, it requires interfacing the Bluetooth chip, which has limited memory and computing resources, with the image sensor in a manner that allows it to operate at the maximum frame rate. Second, the system requires a light-weight lens and connector to the image sensor, and third it requires a small light-weight 2.4 GHz antenna for the Bluetooth chip.

To avoid using external components like field programmable gate arrays (FPGAs) that add

weight, we connect the camera directly to the Bluetooth chip (see Fig. 4.3). The microcontroller sends control commands to the camera using a standard Inter-Integrated Circuit (I2C) bus to set parameters such as amplifier gain, operating mode, etc. Upon triggering an image capture using I2C, the camera begins sending data and provides line-valid and frame-valid signals that indicate when one row of the image or full frame has been sampled and can be read.

To allow the bluetooth chip to read the image sensor data directly at the highest frame rate requires using the direct memory access (DMA) feature which bypasses the central processing unit (CPU) to maximize the speed of memory reads and writes. The DMA feature however is only exposed for certain common communication protocols, none of which are supported by the camera. The camera only outputs data and clock signals. In order to read the camera data we repurpose the Serial Peripheral Interface Slave (SPIS) interface which does have access to DMA. Reading data over the SPIS interface requires a data signal, clock signal and a chip select (CS) signal. We can directly use the camera's data and clock outputs for the first two. To provide the missing CS signal required by the protocol we leverage the line-valid output of the camera. This signal is configured to trigger an interrupt on the Bluetooth chip which toggles an output pin to spoof the missing CS signal.

This method is however limited by the 256 B DMA buffer, and data greater than this size again requires using the CPU. This was addressed by using the line valid signal to start a timer to increment memory addresses after reading each line. The timer is set so that after all pixels of a line have been read, the GPIO CS signal is reset to indicate to the SPI peripheral that the transfer was completed, and the memory addresses are incremented. When the line-valid signal goes low again, indicating new camera data is available, the process is restarted to read the next line of the image; this sequence repeats until all lines in the image have been read. The data was transmitted to a Bluetooth transceiver (e.g., smartphone) using the built-in library functions used for transmitting a file via Bluetooth.

The additional components required to complete the camera assembly are a lens and connector. While there exist commercial modules that integrate all of these components[122], they weigh 158 mg and we therefore build a custom camera assembly weighing only 24 mg. One approach

to build a custom ultra-lightweight lens is to use a pinhole design; however the small aperture required restricts the system to very bright light on the order of 100,000 lux[9].

We instead use a plano-convex lens placed over this bare image sensor. We constructed carbon fiber structures to maintain the correct distance between the image sensor and the lens (see fig. 4.1C). The smaller 1.5 mm lens requires a distance of 570 μm to the sensing element. To achieve this, we placed two square pieces of 90 μm thick carbon fiber laminates with 1.4 mm circular cutouts directly on top of the sensor's protective covering and centered over the active area. A carbon fiber box was then assembled around the sides of the image sensor and lens assembly. For the larger lens, we assembled a 1.49 mm tall, 5-sided carbon fiber box to hold the lens at the correct height, which was placed over the image sensor. For both lenses, a final square piece of carbon fiber with an 0.5 mm aperture was placed over the lens to prevent saturation in brighter outdoor light conditions (>1000 lux). The aperture size was found experimentally to allow for good balance between performance indoors and outdoors (see Video 4.11); however, it could be removed or modified for operation in very low light conditions (<10 lux). All parts for the lens assembly were attached together using CA glue.

We captured images using these two lightweight lenses, a 1.5 mm diameter 1 mm focal length lens (Edmund Optics 43394) weighing 4.8 mg (sample raw images in Fig. 4.4C) and a larger 3.8 mm diameter 2.33 mm focal length lens (Panasonic EYLGUDM128) weighing 20.3 mg (sample raw images in Fig. 4.4D and Fig. 4.7). The lenses were placed on an optics table, and a printed test pattern image (1951 USAF resolution test chart, MIL-STD-150A) was moved to progressively farther distances with dimmer light conditions. We note that in addition to the properties of the lens, the quality of focusing could have been affected by fabrication error during manual assembly. The smaller lens is ideal for capturing images up close to the camera. The larger lens allows this camera to zoom in on large objects farther away, such as buildings or humans. Because many insects live in dark environments, we tested the camera and found that it could capture discernable images down to luminance levels of 5 lux or below (see Fig. 4.7 and Video 4.6).

In addition to a lens we also design a custom transmit antenna to reduce weight. While small lightweight chip antennas are available for use at 2.4 GHz, they often require large ground planes




Figure 4.7: Sample images comparing lenses. Sample images comparing 1.5 mm diameter and 3.8 mm diameter plano-convex lenses imaging a standard test pattern (1951 USAF resolution test chart, MIL-STD-150A). The test pattern was printed and held at the specified distance away from the camera. Images were captured indoors at a resolution of 160×120 at illuminance levels 200 lux and 5 lux. The comparable image quality at 5 lux demonstrates the system can operate in low light conditions.

of 3 cm or more and clearance to other components to achieve good performance. Considering the required clearance and ground plane exceed our target form factor dimensions, we instead design a custom antenna using an ≈ 5 mm long segment of 43 AWG wire. In order to achieve resonance with this short structure we incorporate a 3 turn helix at the base held in shape by cyanoacrylate adhesive.

Using this antenna we evaluated wireless communication range versus video frame rate by placing our form factor Bluetooth radio at a fixed location at ground level in a grassy field and moving a receiver to increasing distances. The small form-factor Bluetooth circuit was held approximately 1 cm above the ground in a grassy field to simulate the height of the back of an insect. A second remote Bluetooth transceiver held at a height of around 1 m was placed along a straight line at increasing distances from the first device. Two remote devices were tested: a Galaxy S9 smartphone and an nRF52832 Bluetooth development board connected to an 8 dBi patch antenna (L-Com RE09P). Adding the patch antenna at the remote Bluetooth transceiver compensates for

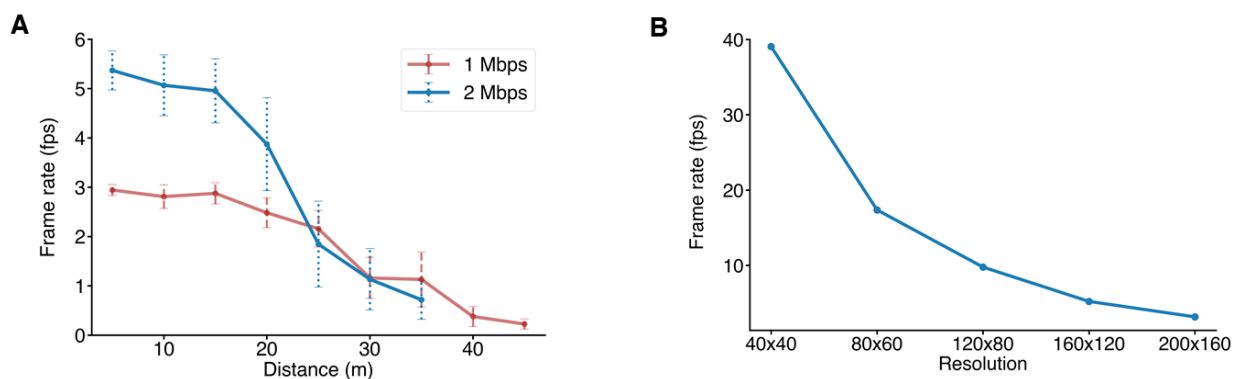


Figure 4.8: Frame rate versus distance and resolution using Bluetooth. **A)** Frame rate measured versus range between the small Bluetooth chip and smartphone at two data rates. Error bars indicate mean $\pm 1\sigma$ ($N=50$ frames). **B)** Maximum frame rate at different resolutions.

the loss of the small wire antenna and extends the operating range. Each test was initiated by a command sent from the remote Bluetooth transceiver, after which it would receive data for a 160×120 image from our millimeter-scale Bluetooth device. The time between the first packet of the image and its last packet was logged and used to compute the data rate, and this sequence was repeated over 10 times at each distance.

The system was able to transmit uncompressed images up to 120 m to the Bluetooth chipset connected to an patch antenna (see Fig. 4.4E), and 45 m to the smartphone (see Fig. 4.8A). The rate remained approximately constant for some range of distances after which it began to degrade, with the lower Bluetooth signal strength available at longer ranges. We note that the maximum range in these plots represents the range at which the Bluetooth devices could no longer connect. We implemented JPEG compression on the microcontroller but found that the time and power required to compress video was significantly longer than that needed to stream the raw data.

The total frame rate for receiving data on the phone is affected by the fixed time needed to read the image from the camera, as well as the time required to send the image over the Bluetooth link. We found the latter to be the primary factor limiting the frame rate since the Bluetooth link provides a maximum of 1.4 Mbps of end-to-end throughput using Bluetooth 5 at high signal-to-noise ratio (SNR). Fig. 4.8B shows the tradeoff between increasing image size and frame rate at a Bluetooth

throughput of 1 Mbps. Because of the limitations of the interface between the microcontroller and image sensor, as well as the resolution settings available on the image sensor, the frame rate increases nonlinearly for lower resolution images.

Fully wireless operation requires a power source in addition to a radio link. We used small LiPo batteries to power the system since they can provide sufficient current for both Bluetooth and the boost converter that dominate the system's power consumption. We evaluated battery life when continuously streaming 160×120 images with the boost converter off (camera at 0° angle), which achieved battery lifetimes of 130 min and 307 min for 5 and 1 fps, respectively (Fig. 4.4F). Turning on the boost converter to sweep the arm over 60° in discrete 15° steps for each frame at 1 fps results in a battery life of 270 min (Fig. 4.4F).

4.2.5 *Video streaming from live insects*

We demonstrate that the form factor of this wireless vision system is compatible with live insects. Cameras have been used to study a variety of larger animals giving biologists insights into their behavior with a first person view. This miniaturized camera system extends that same capability to significantly smaller animals like insects. In addition to the form factor, the wireless range of 100 m enables experiments in real natural habitats outside the lab. While prior work [28, 119, 145] can control insect motion and our recent work [60] instruments small insects with sensors, none of these works have been able to demonstrate a streaming video from small insects.

The natural motion patterns of insects present an opportunity for additional power savings. For scenarios in which an insect may be asleep and inactive for long periods of time, we can also put the camera system into a low power sleep mode. We include a low power accelerometer in our design, capable of detecting when the insect moves, that wakes up the system and begins capturing images.

To evaluate whether the accelerometer-triggered operation is a viable strategy for power saving and to demonstrate that the system can be used with live insects in the field, we attached our camera, accelerometer, Bluetooth radio, and battery (PGEB201212C) to the thorax of a death feigning beetle (*Asbolus laevis*, Fig. 4.9A). We choose species of darkling beetles such as this and

the Pinacate beetle (*Eleodes nigrina*, Video 4.2) due to availability and prior work which has shown they can carry the required payload [145]. We note that other species like moths and spiders could also be used as long as they can carry the camera system payload.

To attach the camera, insects were gripped in a pair of tweezers by the thorax and more active specimens were placed briefly in a freezer for cold anaesthesia when necessary. Using CA glue, a small piece of carbon fiber was first adhered to the desired attachment site on the insect's abdomen or thorax to create a flat surface on which the battery could be easily placed and that could be easily gripped by tweezers. Additional CA glue was then used to attach the battery to the carbon fiber platform. The insects were housed indoors in plastic terrariums with rocks and dry soil. They were fed leafy vegetables such as lettuce, or fruits such as apples every few days. The two systems mounted on the insects, shown in Fig. 4.9D and Fig. 4.9G, remained functional as long as the batteries were recharged, and the same systems were operated for repeated experiments over the course of more than 8 weeks. The electronics were removed after all experiments were completed. At the time of writing both beetles were alive and active. No noticeable adverse effects on behavior were observed.

On an overcast day, beginning at dusk and ending well into the night with low light conditions, we allowed the beetle to walk freely in different outdoor environments, including an overgrown parking lot, a fallen log in a grove of trees, a dry stream bed and on a gravel road (Fig. 4.9B). The system was programmed to wake up from a low-power sleep mode and transmit 5 images to a smartphone receiver each time the accelerometer detected motion (Fig. 4.9C). After operating for approximately one hour in sleep mode during experimental setup, the system ran for an additional 363 min in its accelerometer-triggered mode. The number of accelerometer triggers recorded on the smartphone, which indicates the activity of the insect, is plotted versus time in Fig. 4.9D. Additionally, the top 25 intervals between movements, showing periods of inactivity totaling 189 min, are plotted in Fig. 4.9E; these explain the significant battery life improvement compared to continuous streaming results shown earlier in Fig. 4.4F.

We similarly evaluated the full system, including a scenario where the robotic arm pans the camera horizontally over a roughly 60° range. We attached the full system with the battery

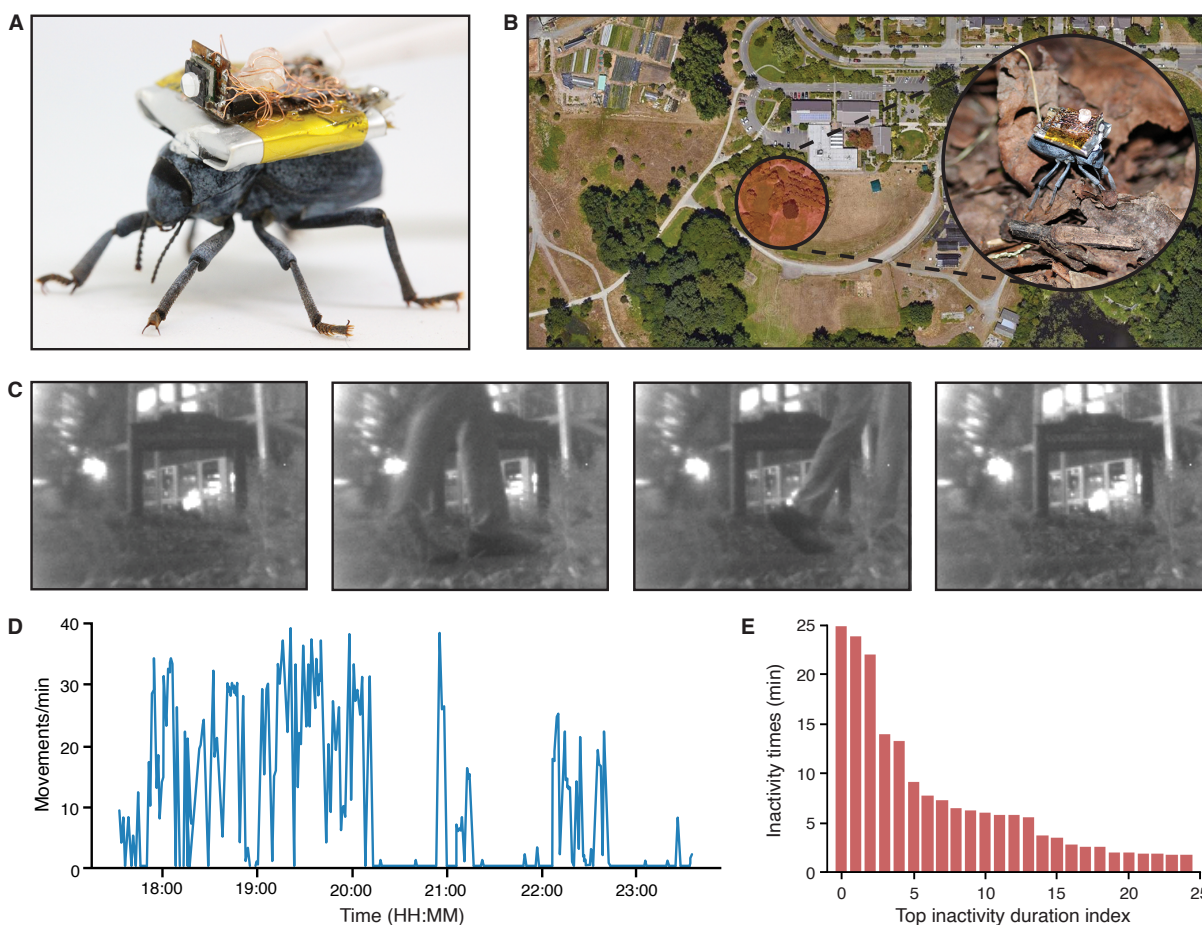


Figure 4.9: Field evaluation of accelerometer triggered camera. **A)** Close-up image of the wireless camera without a micro-robotic arm on the back of a death feigning beetle. **B)** Experiment site with a grove of trees, a dry stream bed, and gravel paths. The beetle walked freely in 4 different locations. **C)** Images from the camera showing a person walking. **D)** Beetle motions detected by the accelerometer per minute over the 363 min experiment. The inactivity explains the improvement in battery life over continuous streaming. **E)** Top 25 intervals between accelerometer triggers when the system is in sleep mode.

(PGE201212C) to the abdomen of a Pinacate beetle (*E. nigrina*, 1.13 g, Fig. 4.10A, Video 4.7) rather than its thorax since this is the largest flat region of the body for this species. We then placed the beetle on the ground and let it wander freely in a large gravel parking lot until it reached an edge, at which point it was placed back in the center (Fig. 4.10B). The system was programmed to activate upon an accelerometer trigger and capture a sequence of 5 images (see Video 4.8). The

robotic arm was set to move in discrete steps of 15° per image. The actuator moved to the new angle and captured an image within 690 ms, achieving a new panoramic view every 3.45 s. This delay value between frames included the total time required to move the camera, read the image data, and transmit the image over Bluetooth, which included a conservative margin added to account for the higher latency observed at longer ranges with low Bluetooth signal strength. Fig. 4.10C shows a sample panoramic view. Images are combined to create a panorama offline using standard panorama stitching algorithms. Unlike the death feigning beetle in the previous experiment, this beetle was active almost continuously, as illustrated in the plot of accelerometer activity in Fig. 4.10D,E.

We observe that in both of these experiments the insect's motion does not cause much image distortion. This can be seen in Video 4.9 in which the beetle walks on a flat desk surface where the images follow the gait of the insect. Additionally, as seen in Video 4.10, even when the beetle is traversing uneven terrain with rocks approximately equal to its height, the objects in the individual frames are still identifiable. These evaluations also indicate that the attached payload does not significantly impact the insect's ability to navigate complex terrain.

4.2.6 *Video streaming from insect-scale robots*

The ultra-lightweight form factor of our steerable wireless camera can enable numerous applications in robotics. By reducing the size of the camera system, we can enable the development of more complex centimeter-scale robots that support vision. Additionally, minimizing the weight of our steerable vision system reduces the robot's payload requirements resulting in longer battery life or faster motion.

We design a small power-autonomous, terrestrial robot to demonstrate the potential applications of our vision system. Fig 4.11 shows prior power-autonomous insect-scale robots [132, 51, 115, 42, 13, 66], none of which support wireless vision and are all larger in size than our robot. To design our insect-scale robot we begin by examining the options for achieving locomotion at this scale. One option is to use piezo actuators similar to the one used to move our camera. While piezos are attractive due to their low weight, operating our boost converter to produce a contin-

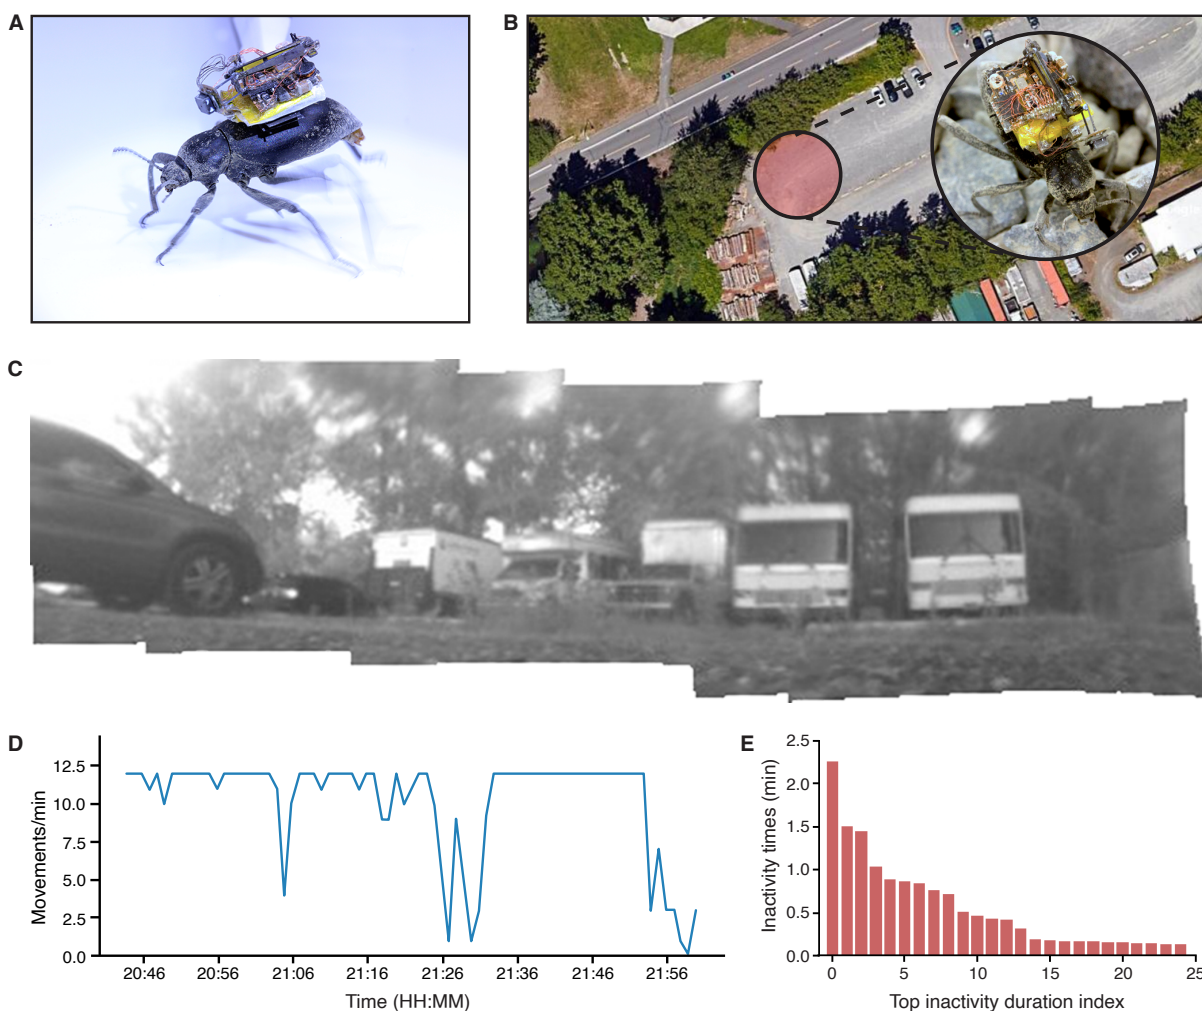


Figure 4.10: Field evaluation of micro-robotic arm on darkling beetle capturing panorama images.

A) Wireless camera system with the micro-robotic arm attached to a darkling beetle. The camera was set to capture five images as it rotated in 15° steps to capture a panorama. **B)** Aerial view of the experiment site in a gravel parking lot showing buses and trucks in the Southwest corner. **C)** Panorama showing the trucks and buses composed of 5 images captured by the insect-mounted camera while rotating 60° . **D)** Beetle motions detected by the accelerometer per minute. **E)** Top 25 intervals of beetle inactivity.

uous high voltage sinusoidal waveform consumes over 150 mW. This high power draw requires compromising on battery size since many small batteries do not support high peak currents, battery lifetime due to the limits of battery energy density, and speed as the robot must instead carry

Power Autonomous Robots	Camera	Length (cm)	Motion Power Consumption (mW)	Weight (g)	Wireless Comm.
Kilobot [36]	X	3.3	-X-	>3*	Infrared
HAMR [37]	X	4.5	395	2.8	RF
RoACH [35]	X	3	-X-	2.4	Infrared
DASH [38]	X	10	277	16.2	RF
This work	✓	1.6	33	2.8	RF

* This design does not specify the weight but uses a CR2032 coin cell battery which itself weighs more than 3g.

Figure 4.11: Comparison with prior power autonomous robots. Comparison of prior small, power autonomous robots including their size scale, mass, power consumption and wireless interface.

a heavier battery payload. Alternative actuators such as eccentric rotating mass (ERM) vibration motors have been used to achieve locomotion at this scale [115, 73, 143] and do not incur the inefficiencies of producing a high voltage output. Although these motors weigh significantly more than piezos (≈ 1 g), we observe that they can produce enough force for forward locomotion of a robot using tens of milliwatts, after our power optimizations, as opposed to the hundreds of milliwatts consumed by state-of-the-art piezo systems[42].

We design an insect-scale robot using two vertically oriented ERM vibration motors (Seed Technology 1027) as shown in Fig. 4.12A,B. The motors are held 13 mm apart in a chassis constructed using two parallel sheets of 254 μ m thick FR4 and three 7 mm long metal legs. These dimensions are chosen such that they are the minimum required to accommodate the 12 mm wide lithium polymer battery mounted between the FR4 layers. All of the electronics including Bluetooth, the boost converter, and voltage regulators are mounted on the top FR4 layer. The piezo actuator and camera assembly are mounted above the electronics using carbon fiber rods. The robot chassis measures 1.6 cm in length and 2 cm in width.

The robot is capable of moving at speeds up to 3.5 cm/s (2.2 times the body length) with no additional payload (see Video 4.11). Fig. 4.12C shows the effect of running the vibration motors at a constant power while carrying additional payload beyond the steerable wireless camera system and the battery. We note that prior vibration motor-based designs use heavier and larger coin cell batteries and have a speed of 1 cm/s (0.33 times the body length)[115]. Thus, in comparison to prior work, our design achieves a 6.67 speedup when normalized to body length.

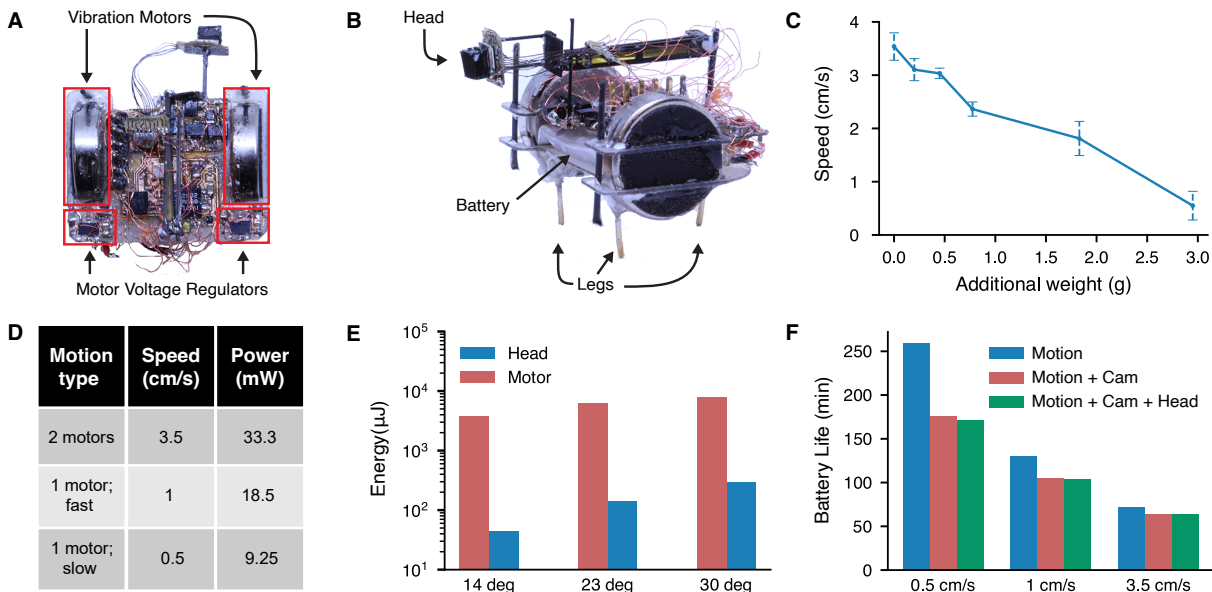


Figure 4.12: Evaluation of the insect scale wireless robot. **A,B)** Close-up images of the robot with vibration motors, three legs, a battery, Bluetooth chip, camera and robotic head. **C)** Effect of carrying payload greater than the camera system and the battery on robot speed, while running the vibration motors at a constant power. Error bars indicate mean $\pm 1\sigma$ ($N=5$ trials). **D)** Robot power consumption for different motion types and speeds. **E)** Energy consumption for different angles when using the motors versus head to turn the camera. **F)** Battery life for different robot speeds without video streaming, with 1 fps video streaming, and 1 fps video streaming while panning the camera.

Our design also employs additional techniques to further decrease the robot's power consumption. We employ multiple strategies to reduce the power required by our robot. First, we observe that the motors can be run at voltages as low as 0.7 V while still producing enough force to move the robot. However, we note that if directly powered at 0.7 V at startup, they do not produce the required motion. To address this issue, we apply a ≈ 20 ms higher power pulse at 2.8 V to begin rotating the internal mass and then transition to a low-power steady state of 0.7 V. At 0.7 V, each motor draw ≈ 14 mA of current resulting in power consumption of 10 mW. Second, to generate this 0.7 V low voltage output from the 3.7 V battery, we connect the motor to an H-Bridge circuit controlled by a pulse width modulated (PWM) signal. Setting the duty cycle of the PWM signal results in the desired average voltage, however this method has a high peak current. This results in

lower efficiency as some power is lost to conduction losses. To reduce this loss we instead use a dedicated high efficiency switching voltage regulator to achieve a total power of 16 mW when running one motor, the voltage regulator and the microcontroller. Under these conditions the robot can move forward at speeds of up to 3.5 cm/s when running both motors simultaneously. Running only one motor at a time allows the robot to turn right or left. Third, we further optimize to increase the battery lifetime to account for the nonlinear relation to power consumption. Since reducing power results in a nonlinear increase in operation time, we further optimize for battery life by having the option to operate only one vibration motor at a time. Doing this causes the robot to pivot off of one leg and move forward. Alternating operation of the motors produces a forward walking gait. Operating in this mode however trades lower power for slower speed.

Fig. 4.12D shows the results of these strategies such as the power required for motion after reducing the motor power and running only one motor at a time. These optimization in the robot power raises the question of whether the piezo based head mechanism described above is still a more efficient mode of adjusting the camera's field of view. While the head mechanism presents the benefit of the ability to move the camera independently of the robot, Fig. 4.12E demonstrates that the piezo mechanism requires less total energy. Specifically, it takes between 26-84 times less energy to rotate the head than to rotate the robot body by the same angle. Because of its actuation time of a few milliseconds and ability to hold an angle without continuous power input, the energy required to move the piezo is lower than rotating the whole robot. The difference in the energy savings across angles is because of the nonlinear behavior of the boost converter which takes longer to move the piezo mechanism to larger angles.

The battery life of running the whole robotic system is shown next in Fig. 4.12F. Components are added cumulatively to show the battery life reduction when adding each additional component. In these experiments, the robot is moving continuously, the camera is streaming at 1 fps and the arm moves 30° once a second. We observe that at the lower speed of 0.5 cm/s, the power of the robot and camera are the same order of magnitude resulting in a noticeable difference in battery life. In contrast at its maximum speed, the robot requires 10x more power to move its body than to capture images at 1 fps. We also note that in all cases, operating the piezo actuator requires

minimal energy and therefore has little impact on battery life.

Running the streaming camera while the robot is in motion however introduces an additional challenge of addressing vibration. We primarily observe oscillation about the horizontal axis of the image due to the hinge of the head mechanism. The hinge and connectors are designed to provide minimal resistance to the piezo in order to achieve a large angle. However the result of the low stiffness is that the vibration of the robot induces oscillation. Because the camera uses a rolling shutter, if the sampling time is not significantly faster than the period of the vibrations caused when moving, the robot images appear distorted (see Fig. 4.13). One approach for minimizing distortion is to reduce the image resolution and as a result increase the frame rate. This however significantly compromises image quality. A second approach would be to use the piezo mechanism to stabilize the image. This however would add complexity as it requires a feedback mechanism, additional computation, and potentially more power. We implement an alternative solution of duty cycling the robot motion, effectively trading off speed to maintain image quality. To achieve an effective frame rate of 1 fps, we set a timer to run the robot for 640 ms and trigger image capture after a few hundred millisecond pause to allow any remaining vibrations to settle. This results in a trade-off between speed and frame rate. At 1 fps, the average speed of the robot is reduced to 2.24 cm/s due to the intermittent pauses and results in a battery life of 93 mins. In contrast, a lower frame rate of 0.5 fps allows for an average speed of 2.87 cm/s and a battery life of 81 mins. Despite a reduction from the peak speed of 3.5 cm/s, these speeds are still faster than the 1 cm/s speed achieved by prior vibration motor designs[115].

Fig. 4.14A,B and Video 4.3 demonstrate the potential for using the robot's new onboard vision capability for navigation. The robot is placed on a desk with obstacles and steered remotely by a human operator sending commands from a smartphone showing the live video stream. The live video is used by the teleoperator, in real-time, to send navigation commands from the smartphone to the robot. Fig. 4.14A shows an overhead view illustrating the path taken by the robot and Fig. 4.14B shows the sample images from the robot's camera as it travels.

Fig. 4.14C,D,E and Video 4.4 demonstrate an additional application of using the piezo driven head to focus on a moving object. In this video the robot itself remains in a fixed position, but




Figure 4.13: Distortion during motion. A) Image showing the bottom corner of a Rubik's cube from the robot's perspective. B,C) Images captured during robot motion. The side-to-side distortion results from the rolling shutter of the camera which samples data line by line while the vibration motors shake the robot.

another similar sized robot moves in a straight line in front of it. As the moving robot exits the field of view of the camera, a command is sent by a human operator to move the piezo to the left to track the moving robot. The background of the image includes a series of numbers that increment from right to left to provide an additional reference showing the camera has moved. This shows the feasibility of using our steerable wireless camera to maintain focus on moving objects and robots, without the need to move the whole robot body.

4.3 Discussion

Here, we describe a fully wireless mechanically steerable vision system in the form factor compatible with small robots and live insects. Our design draws parallels to biological insects in multiple ways. Our design uses a commercial image sensor with over 100,000 pixels in a 2×2 mm area compared to the 4000-5000 imaging elements in a blowfly [134]. Additionally, like many insects that have limited color perception [79], we use a monochrome camera in order to reduce communication bandwidth and optimize for power consumption. Many insect vision systems have a flicker fusion frequency of 200 Hz or less [94]. While our current frame rate is limited by the data rate of our wireless link, our microcontroller can read an 80×60 image from the camera at 200 Hz. This opens the potential for using the vision data on-device for control. Insects of the Diptera family




Figure 4.14: Navigation and focusing on another moving robot. **A)** Overhead view of robot's path when a human operator uses the camera to navigate. **B)** Sample images from the robot's camera during navigation. **C)** Side view showing our wireless robot is stationary with the camera looking at another robot marked with an arrow moving left. Background shows numbers incrementing from right to left to show motion. **D)** Top view of scene shows the moving robot going across the field of view of the camera. The camera rotates in discrete steps to maintain focus on the moving robot. **E)** Images captured at each angle.

(flies) use between 1-13% of their mass for vision including the cornea, retina and optic lobe of the brain [8]. In a blowfly for example, the visual system consumes 2.5% of the animal's total mass. For comparison the image sensor and Bluetooth-enabled microcontroller consume 4.8% of the total mass of our robot.

Our wireless video transmission could also provide precise position feedback, for example, by combining it with accelerometer data to perform vision-based simultaneous localization and mapping (SLAM). While we demonstrate integration on a vibration motor based robot, given its low mass, our steerable camera system could also be integrated on other insect-scale robots [42]. Given the low weight and power requirements for our Bluetooth video streaming subsystem, it could also enable wireless video streaming applications for aerial micro-robots, which currently lack this capability[89, 63]. This work also complements efforts to steer the gross muscle motions of insects through neural input[28, 119, 145]: it adds the capability of controlled actuation along application-specific degrees of freedom. In addition to steering a camera for a wider field of view, incorporating actuators could enable a much richer interaction with the environment.

The primary bottlenecks to further scaling down our vision system are the energy density of batteries and power consumption. One approach to building a system small enough to fit on smaller insects like bumblebees would require an improved battery technology with higher energy density, however this could also be achieved by reducing the power required for the vision system by an order of magnitude. The addition of a lightweight solar cell to supplement the battery could also make this feasible. Scaling down further to the weight that insects like blowflies could carry would require more significant improvement of all system components. At this scale even the camera and lens weight of 24 mg become significant. The weight of the optics could potentially be reduced using emerging metasurface technologies [11]. Additionally the size of the image sensor itself could be reduced by developing a custom image sensor with smaller pixel size or lower resolution. Reducing resolution would also be advantageous from a power perspective as this reduces the time the radio transmits. The steering mechanism can be scaled down by further optimizing the high voltage piezo driver. This requires scaling down the weight of the inductor and its power consumption which is challenging. Another approach is to explore improved actuation mechanisms that could operate at low voltage.

Vision serves as an important sensory input for both animals and robotic systems. While vision has traditionally been reserved for physically larger robots or animal monitoring systems, we built a steerable wireless vision system that is integrated on live insects and insect scale terrestrial robots.

Further, imitating nature, we demonstrated the ability to move the visual system of insect-scale robots independent of their body allowing them to gather more visual information in an energy efficient manner. Given the importance of vision for navigation and communication, this work represents an important milestone in realizing autonomy for insect-scale robotics.

4.4 Videos

Video links:

https://homes.cs.washington.edu/~vsiyer/thesis/beetle_videos.html

Video 4.1: Accelerometer-triggered camera. A darkling beetle (*E. nigrina*) is held facing a metric ruler. The insect's motion triggers the camera to turn and stream images to a smartphone.

Video 4.2: Sample real-time video streams. The smartphone sets the resolution and starts streaming video. The inset video shows a reference view of the scene (a person walking).

Video 4.3: Navigation of an insect-scale robot using wireless vision. An insect-scale robot uses its camera to navigate around obstacles. The camera streams video to a smartphone allowing a human to steer the robot.

Video 4.4: Focusing on moving objects. A stationary robot streams video of another robot moving across its field of view. A human operator steers the camera to maintain focus on the moving robot.

Video 4.5: Camera motion angle measurement. The full camera and microrobotic arm assembly without the beetle is held at a fixed position above a printed protractor. An 28 mm carbon fiber rod is attached to the camera to show the angle. The smartphone moves the camera across its roughly 60° of motion.

Video 4.6: Automatic light level adjustment. The beetle is held in the same position as Videos S1 and S2 at night. A flashlight illuminates the camera from above and is then turned off. The inset video of a digital light meter at the top left shows the change in light level. The camera automatically adjusts to the lower light level.

Video 4.7 Smartphone-controlled camera steering. The mechanically steerable wireless camera is mounted on the back of a live darkling beetle held in tweezers. The camera streams monochrome 160×120 resolution video to a smartphone at 5 fps. The phone sends commands to move the microbotic arm and directs the camera to different angles.

Video 4.8: Capturing a panorama. The beetle is held on tweezers in a gravel parking lot facing parked trucks. The camera is operated in accelerometer-triggered mode as in Video S2. The inset image at the top left shows a reference image of the captured scene.

Video 4.9: Video from beetle walking on flat surface. Video stream from a beetle walking on a flat table.

Video 4.10: Free walking beetle in the field. The beetle walks freely in a gravel parking lot while streaming video to a smartphone. The video shows both the rocky ground and the truck from the insect's perspective.

Video 4.11: Robot speed when carrying payloads. The robot moves on a steel surface at its maximum speed of 3.5 cm/s carrying no payload. Its speed decreases with progressively heavier payloads.

Chapter 5

LESSONS LEARNED & GENERAL DESIGN PRINCIPLES

This chapter gives an overview of the lessons learned from building the three systems discussed in this dissertation and outlines general design principles and methodologies for building millimeter scale wireless systems. Millimeter scale systems are constrained by size, weight, and power (SWaP). These properties are often interconnected, as scaling down in size leads to weight reduction based on the density of materials. Similarly, the finite energy density of batteries or availability of harvested power necessitates that small and light-weight systems be low power.

Scaling down in size also raises practical concerns as it constrains the commercial off-the-shelf (COTS) components that can be used. While this is especially true for actuators and power sources which are limited by battery chemistry, COTS electronic components for computing, communication and sensing are more readily available. The confluence of Moore's law scaling, advances in MEMs sensor manufacturing and packaging, high sensitivity radio protocols, and highly integrated programmable microcontrollers developed for wearable and IoT devices provides a robust set of primitives for developing novel wireless systems. As explained above, building off of programmable computing components can enable rapid innovation in new domains without the need for time consuming and resource intensive custom silicon. This thesis demonstrates how combining that approach with novel backscatter technologies, power harvesting, and bio-inspired designs or bio-integration can achieve mobile, millimeter scale wireless sensing.

The remainder of this section discusses trade-offs and solutions to consider for power, communication, and computation, in SWaP constrained systems.

5.0.1 Computing

Computing platforms are the core component of each of the systems in this dissertation as they are required to interface with sensors, control communication schemes, and perform power management. Computing chips span a variety of capabilities and resource requirements. The smallest microcontroller tested in these works is the ATtiny20 measuring only 1.540×1.388 mm. Although this chip is small and low power, it also has a number of limitations. It runs at a maximum of 12 MHz and 2 kB of program memory. Additionally it does not include hardware support for floating point or protocols like the I2C interface used by many sensors. Because the chip is programmable, these functions can be implemented in software to enable this smallest form-factor class of devices, however the memory limit gives an upper bound on the complexity of tasks it can perform. The system architecture described in chapter 2 uses this chip to do the bare minimum of reading raw data from the device and then leverages its ability to communicate using backscatter to offload floating point processing and other operations onto the receiver.

This example presents a system architecture that can extend to larger computing devices as well. At a high level, this is the same idea of using edge or cloud computing resources to augment the capabilities of resource limited IoT systems. For systems that seek to minimize SWaP constraints, it is often advantageous to perform only the bare minimum computation on the device itself. In situations with high bandwidth sensor data such as streaming cameras this may shift the majority of power consumption into communication. The streaming camera presented in chapter 4 shows an example of this; however in this case streaming uncompressed data was still faster than performing the computation needed for JPEG compression onboard with the chip's main CPU. These limitations highlight an opportunity to use other computing devices like a second co-processor or field programmable gate arrays (FPGAs) to accelerate some of these computations. Similar to the use of graphics processing units to accelerate machine learning computations, these architectures may soon become commonplace for miniaturized devices with high bandwidth sensors or tight latency requirements that necessitate onboard processing.

5.0.2 Communication

This section discusses three options for wireless communication including backscatter, Bluetooth, and Low power Wide area Network (LPWAN) protocols like LoRa.

- *Backscatter.* Backscatter provides a powerful tool to decouple power and size expensive radio components onto another device. Beyond the modulation techniques presented in chapter 2, backscatter can be used to produce any arbitrary modulation and implement a variety of communication schemes [72, 147, 61, 135]. Moreover, backscatter can even be used to communicate with physical objects that lack electronic interfaces [57, 56]. This comes at a cost however and backscatter is not the solution to every problem. The combination of self interference and high path loss limit the range of these systems compared to radios. While chapters 2 presents a method to extend wireless range it comes at the expense of greater infrastructure requirements for access points. In particular, for systems such as robots where the power required for motion may be an order of magnitude larger than radio communication, using backscatter adds complexity with only marginal benefit. In contrast there are other systems where SWaP constraints make backscatter the only viable options. Backscatter and radio communication are not exclusive however, and there is significant potential for hybrid systems that leverage both.
- *Bluetooth radios.* While backscatter enables operation in the smallest power and weight regimes, slightly larger systems (> 100 mg) can leverage highly integrated chipsets which combine a general purpose microcontroller with a Bluetooth radio such as the Nordic nRF52 series. The extra size and weight increases communication range without the need for a signal source as seen in chapter 4. Bluetooth also comes with the advantage of ubiquitous receivers on smartphones, laptops, cars, and many other devices. This enables applications such as crowdsourced wildlife tracking where anyone with a smartphone app can receive transmissions. For many scenarios Bluetooth provides a balance between relatively small size and low power, line of sight ranges up to 1 km [58], along with bandwidth sufficient

for supporting many kinds of sensors including low resolution video streaming. The convenience of receiver compatibility makes it an ideal choice for systems that can afford a larger power source such as a LiPo battery or very low duty cycle.

- *LPWAN protocols.* In addition to physical SWaP constraints, communication range is also an important aspect of any wireless system. Backscatter minimizes power and size but is limited to shorter ranges (< 100 m). Bluetooth requires additional power and size resources but enables ranges up to 1 km. LPWAN protocols such as LoRa enable even longer ranges in line of sight conditions, or robust operation in non-line-of-sight conditions. These protocols leverage coding techniques such as chirp spread spectrum, lower data rates, and higher transmit powers to achieve very long ranges. This classic trade-off of communication rate versus range is also tightly coupled to SWaP requirements: lowering data rate increases transmission time and therefore power and size. These techniques can however be adapted to small backscatter devices to improve their range, and example is discussed in chapter 2.

5.0.3 Power sources

All electronic systems require a power source, and as explained above the energy density of a power source like a battery directly couples it to the system's size and weight. Identifying the power requirements and the mass constraints of the system can narrow the design space and suggest the optimal power source for an application. For example, wind dispersal requires ultra-light payloads of < 70 mg or less but environmental sensors can operate at low power with long sampling intervals making this system amenable to power harvesting. In contrast, insects such as bumblebees can carry approximately 100 mg, while larger insects like moths, asian giant hornets, and beetles can support larger payloads. The larger payload allowed by these animals enables greater capabilities like higher bandwidth sensor and actuators.

- *RF Power harvesting.* The two methods of RF power transfer include near field magnetic induction and far field power transfer. Near field power transfer is highly efficient and can

deliver large amounts of power (commonly used for smartphone charging and other commercial applications) but limited to centimeter ranges [149, 68]. This makes it useful for mobile systems with a battery that can go perform a task and return to a charging pad, or for devices moving within a confined area, but cannot be used to deliver power during truly remote operation. Far field power transfer (distances $d \gg 2\lambda$) can operate at longer ranges, however due to the inverse square power scaling with distance, these systems can often deliver microwatts of power or less beyond a few meters and suffer poor end to end efficiencies less than 1% [136].

- *Solar harvesting.* Solar power harvesting is a very scalable technique with small cells that can provide milliwatts of power to run a microcontroller to megawatt solar farms. The trade-off for solar of course is that light availability often varies in real world conditions and is not present at night. Variability in light levels can be addressed using energy storage elements like capacitors and system design techniques like duty cycled operations or dynamic sampling rates depending on power availability. Applications which cannot tolerate overnight outages should be combined with a battery or alternative harvesting element.
- *Laser power delivery.* Lasers provide another method for optical power delivery which is scalable and can provide significant amounts of power [63, 55]. Using an optical link however comes at the expense of requiring line of sight, dust-free conditions, and precise alignment [55]. This approach is best for harvesting applications that require high power such as continuously running actuators or high bandwidth sensors like cameras. These systems can also be used indoors [55], in environments without sunlight, or with multirobot systems where a robot with a large battery could deliver power to smaller ones.
- *Vibration harvesting.* Power can be harvested from vibrations as well using piezo materials. While not explored in this thesis, prior work has shown the use of resonant structures to harvest energy from flapping wing flight of moths [112, 110]. A major challenge of this approach is that the length scale required for integration on insects and the material properties

of piezos make designing a structure with resonance at the flapping frequency challenging. Vibration harvesting could be used in other less size constrained applications such as on structures, roads or machinery that experience vibrations.

- *Batteries.* While power harvesting is attractive and often the only option at the smallest size scales, batteries allow for significant simplification in system design and reliable operation during their lifetime. Pairing batteries with harvesting, and system level techniques to reduce power consumption provides an ideal balance of reliability and longevity. Batteries are available in a range of sizes and have their own limitations. Small or thin film batteries often cannot provide sustained, high current outputs to support radios making them preferable for heavily duty cycled applications. In contrast larger lithium polymer (LiPo) batteries such as those used in consumer devices ranging from smartphones to electric vehicles can provide large amounts of power but at the cost of higher weight.

Chapter 6

CONCLUSION AND FUTURE DIRECTIONS

This dissertation demonstrates how taking an interdisciplinary approach combining programmable general purpose computing devices with low power communication techniques like backscatter, energy harvesting solutions, and microfabrication techniques to build small actuators can enable a variety of innovative, tiny wireless sensing systems. Bringing this together with inspiration from nature and biological systems add new capabilities such as mobility to these systems to push the boundaries of technology even further.

This approach has led to the following three innovations described in detail above:

- *Bio-inspired wind dispersal of battery free wireless devices.* This system demonstrates millimeter scale battery free wireless sensors that take inspiration from natural structures like dandelion seeds to disperse in the wind for automated deployment of sensor networks.
- *The Internet of Biological Things: A Flying Wireless Platform on Live Insects.* This work demonstrates how small flying insects like bumblebees can carry wireless sensors. Unlike drones, bees can fly for hours which enables long term wide area sensing for applications like smart agriculture.
- *Wireless steerable vision for live insects and insect-scale robots.* This work takes inspiration from eye and head motion in animals to create a wireless steerable vision system that can efficiently capture panoramic views at low power and small size. This work also enables vision, previously only available on large robots, to power autonomous insect scale robotics, which can be a foundation for building numerous future applications.

While biology does not have the answer to every problem, there are so many things nature can

do better than electronics that we can learn from. This interdisciplinary approach has the potential to enable many more innovations in the next 10 years for wireless and cyberphysical systems.

For example many insects can see in color and at high framerates to be able to navigate the world. Further they can process all that sensory information onboard in their brains. A natural next step for this work it so continue pushing the limits of tiny vision systems. and there is impactful future work in designing custom cameras and combine them with machine learning to make sense of the images onboard which can have applications across IoT, AR, VR and robotics.

Another space where miniaturized wireless sensors could have significant impact is for medical devices. For example, bionic eyes, or visual prostheses that could enable blind people to see. Imagine, if we can make a programmable wireless platform for these applications, it would enable lots of researchers across different domains to get involved in designing technologies for bionic eyes and retina. This would also enable rapid innovation.

There are also many opportunities to go further along the route of bio-integration to design better sensors. For example, building a small, highly sensitive gas sensor is very challenging, but animals like hawkmoths (*Manduca sexta*) have antennae that do exactly that. Researchers have shown that electrical signals from antennae that have been cut off of a moth can be read out, allowing them to be used as sensors. A further step beyond this would be to genetically engineer the antenna to make it more sensitive to a specific gas.

Looking beyond building single devices, these systems can be scaled up to large numbers like the examples in nature. Dandelions produces hundreds of seeds and bees live in colonies with hundreds or thousands of animals. By leveraging the ability to build these systems with commodity hardware, they can be quickly and cost efficiently scaled up to start building large networks of devices and exploring new problems like swarm communication algorithms.

These technologies also have lots of potential for wildlife conservation of small animals and insects. For example, studying hummingbirds in the wild is challenging as these animals can move at speeds of up to 60 mph, and can only carry about as much weight as a bee. Long term tracking to understand their energy consumption, as well as migratory patterns across changing seasons and climates requires innovation in wireless technologies while advancing science and conservation

efforts as well.

The future of wireless and cyber-physical systems hold lots of exciting potential to build new technologies for the internet of biological and bioinspired things and reduce the gap between programmable IoT systems and biology.

BIBLIOGRAPHY

- [1] Evan Ackerman. Dragonfleye project wants to turn insects into cyborg drones, 2017. <https://spectrum.ieee.org/autoton/robotics/industrial-robots/draper-dragonfleye-project>.
- [2] Ethem Erkan Aktakka, Hanseup Kim, and Khalil Najafi. Energy scavenging from insect flight. *Journal of Micromechanics and Microengineering*, 21(9):095016, 2011.
- [3] M. H. Almarshadi and S. M. Ismail. Effects of precision irrigation on productivity and water use efficiency of alfalfa under different irrigation methods in arid climates. In *Journal of Applied Sciences Research*, 2011.
- [4] Amscope. 144 led Four-zone microscope ring light with adapter.
- [5] Mark C. Andersen. Diaspore morphology and seed dispersal in several wind-dispersed asteraceae. *American Journal of Botany*, 80(5):487–492, 1993.
- [6] Pulse Larsen Antennas. W3320 ISM868/915, ISM2.4g, 2017.
- [7] Carol K Augspurger and Susan E Franson. Wind dispersal of artificial fruits varying in mass, area, and morphology. *Ecology*, 68(1):27–42, 1987.
- [8] Roland Baddeley, Peter Hancock, and Peter Földiák. *Information theory and the brain*. Cambridge University Press, 2008.
- [9] Sivakumar Balasubramanian, Yogesh M Chukewad, Johannes M James, Geoffrey L Barrows, and Sawyer B Fuller. An insect-sized robot that uses a custom-built onboard camera and a neural network to classify and respond to visual input. *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*, pages 1297–1302, 2018.
- [10] Austin A Bauer, Murray K Clayton, and Johanne Brunet. Floral traits influencing plant attractiveness to three bee species: Consequences for plant reproductive success. *American journal of botany*, 104(5):772–781, 2017.
- [11] Elyas Bayati, Raphael Pestourie, Shane Colburn, Zin Lin, Steven G Johnson, and Arka Majumdar. Inverse designed metalenses with extended depth of focus. *ACS Photonics*, 7(4):873–878, 2020.

- [12] P. Bhushan and C. J. Tomlin. Milligram-scale micro aerial vehicle design for low-voltage operation. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–9, 2018.
- [13] Paul Birkmeyer and Ronald S Fearing. Dash: A resilient high-speed 15g hexapedal robot. *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 418–419, 2009.
- [14] G Blaj and JH Van Hateren. Saccadic head and thorax movements in freely walking blowflies. *Journal of Comparative Physiology A*, 190(11):861–868, 2004.
- [15] Norbert Boeddeker, Laura Dittmar, Wolfgang Stürzl, and Martin Egelhaaf. The fine structure of honeybee head and body yaw movements in a homing task. *Proceedings of the Royal Society B: Biological Sciences*, 277(1689):1899–1906, 2010.
- [16] Rodney A. Brooks and Anita M. Flynn. Fast, cheap, and out of control: A robot invasion of the solar system. *Journal of the British Interplanetary Society*, 42:478–485, 1989.
- [17] Johanne Brunet and Zainulabeuddin Syed. Enhancing pollination by attracting & retaining leaf cutting bees (megachile rotundata) in alfalfa seed production fields. In *Meeting Proceedings*, pages 67–73, 2017.
- [18] S C Chang, F M Yaul, A Dominguez-Garcia, F O ’sullivan, D M Otten, and J H Lang. Harvesting energy from moth vibrations during flight. *Proc. PowerMEMS*, 2009.
- [19] Vincent Casseau, Guido De Croon, Dario Izzo, and Camilla Pandolfi. Morphologic and aerodynamic considerations regarding the plumed seeds of tragopogon pratensis and their implications for seed dispersal. *PloS one*, 10(5):e0125040, 2015.
- [20] Y. Chen, N. Chiotellis, L. Chuo, C. Pfeiffer, Y. Shi, R. G. Dreslinski, A. Grbic, T. Mudge, D. D. Wentzloff, D. Blaauw, and H. S. Kim. Energy-autonomous wireless communication for millimeter-scale internet-of-things sensor nodes. *IEEE Journal on Selected Areas in Communications*, 34(12):3962–3977, 2016.
- [21] L. Chuo, Y. Shi, Z. Luo, N. Chiotellis, Z. Foo, G. Kim, Y. Kim, A. Grbic, D. Wentzloff, H. Kim, and D. Blaauw. A 915mhz asymmetric radio using q-enhanced amplifier for a fully integrated 3x3x3mm³wireless sensor node with 20m non-line-of-sight communication. *IEEE International Solid-State Circuits Conference*, 2017.
- [22] G Ciuti, R Donlin, P Valdastrì, Alberto Arezzo, A Menciassi, Mario Morino, and P Dario. Robotic versus manual control in magnetic steering of an endoscopic capsule. *Endoscopy*, 42(02):148–152, 2010.

- [23] Donald D Clarke. Circulation and energy metabolism of the brain. *Basic neurochemistry: Molecular, cellular, and medical aspects*, 1999.
- [24] Coherent. Matrix 355 Solid-State Q-switched Laser.
- [25] Green Methods Biological Pest Control and Integrated Pest Management. Bumblebees for pollination.
- [26] Cathal Cummins, Madeleine Seale, Alice Macente, Daniele Certini, Enrico Mastropaolo, and Naomi Nakayama. A separated vortex ring underlies the flight of the dandelion. *Nature*, 562, 10 2018.
- [27] Milanesio D, Sacconi M, Maggiora R, Laurino D, and Porporato M. Design of an harmonic radar for the tracking of the asian yellow-legged hornet. *Journal of Ecology and Evolution*, 2016.
- [28] D. C. Daly, P. P. Mercier, M. Bhardwaj, A. L. Stone, Z. N. Aldworth, T. L. Daniel, J. Voldman, J. G. Hildebrand, and A. P. Chandrakasan. A pulsed UWB receiver SoC for insect motion control. *IEEE Journal of Solid-State Circuits*, 45(1):153–166, Jan 2010.
- [29] W. Daniel Kissling, David E. Pattemore, and Melanie Hagen. Challenges and prospects in the telemetry of insects. *Biological Reviews*, 2014.
- [30] G. C. H. E. de Croon, K. M. E. de Clercq, R. Ruijsink, B. Remes, and C. de Wagter. Design, aerodynamics, and vision-based control of the delfly. *International Journal of Micro Air Vehicles*, 1(2):71–97, 2009.
- [31] Iris De Falco, Giuseppe Tortora, Paolo Dario, and Arianna Menciassi. An integrated system for wireless capsule endoscopy in a liquid-distended stomach. *IEEE Transactions on Biomedical Engineering*, 61(3):794–804, 2013.
- [32] Vincent Gaston Dethier. The hungry fly: A physiological study of the behavior associated with feeding. *Harvard University Press*, 1976.
- [33] Michael H. Dickinson, Fritz-Olaf Lehmann, and Sanjay P. Sane. Wing rotation and the aerodynamic basis of insect flight. *Science*, 284(5422):1954–1960, 1999.
- [34] Robert Dudley. *The biomechanics of insect flight: form, function, evolution*. Princeton University Press, Princeton, NJ, 2002.
- [35] Reuven Dukas and P Kirk Visscher. Lifetime learning by foraging honey bees. *Animal behaviour*, 48(5):1007–1012, 1994.

- [36] Hector D. Escobar-Alvarez, Neil Johnson, Tom Hebble, Karl Klingebiel, Steven A. P. Quintero, Jacob Regenstein, and N. Andrew Browning. R-advance: Rapid adaptive prediction for vision-based autonomous navigation, control, and evasion. *Journal of Field Robotics*, 35(1):91–100, 2018.
- [37] R.S. Fearing, K.H. Chiang, M.H. Dickinson, DL Pick, M. Sitti, and J. Yan. Wing transmission for a micromechanical flying insect. In *Robotics and Automation (ICRA), 2000 IEEE Int. Conf. IEEE*, 2000.
- [38] M. Fojtik, D. Kim, G. Chen, Y. Lin, D. Fick, J. Park, M. Seok, M. Chen, Z. Foo, D. Blaauw, and D. Sylvester. A millimeter-scale energy-autonomous sensor system with stacked battery and solar cells. *IEEE Journal of Solid-State Circuits*, 2013.
- [39] Rossella Fontana, Francesca Mulana, Carmen Cavallotti, Giuseppe Tortora, Mario Vigliar, Monica Vatteroni, and Arianna Menciassi. An innovative wireless endoscopic capsule with spherical shape. *IEEE Transactions on Biomedical Circuits and systems*, 11(1):143–152, 2016.
- [40] Sawyer B Fuller, Michael Karpelson, Andrea Censi, Kevin Y Ma, and Robert J Wood. Controlling free flight of a robotic fly using an onboard vision sensor inspired by insect ocelli. *Journal of The Royal Society Interface*, 11(97):20140281, 2014.
- [41] Jon Gjengset, Jie Xiong, Graeme McPhillips, and Kyle Jamieson. Phaser: Enabling phased array signal processing on commodity wifi access points. In *MobiCom '14*, 2014.
- [42] B. Goldberg, R. Zufferey, N. Doshi, E. F. Helbling, G. Whittredge, M. Kovac, and R. J. Wood. Power and control autonomy for high-speed locomotion with an insect-scale legged robot. *IEEE Robotics and Automation Letters*, 3(2):987–993, April 2018.
- [43] Denise Grady. The vision thing: Mainly in the brain. *Discover*, 1993.
- [44] M. A. Graule, P. Chirarattananon, S. B. Fuller, N. T. Jafferis, K. Y. Ma, M. Spenko, R. Kornbluh, and R. J. Wood. Perching and takeoff of a robotic insect on overhangs using switchable electrostatic adhesion. *Science*, 352(6288):978–982, 2016.
- [45] David Greene. The role of abscission in long-distance seed dispersal by the wind. *Ecology*, 86:3105–3110, 11 2005.
- [46] DF Greene and EA Johnson. The aerodynamics of plumed seeds. *Functional Ecology*, pages 117–125, 1990.

- [47] Melanie Hagen, Martin Wikelski, and Daniel Kissling. Space use of bumblebees (*bombus* spp.) revealed by radio-tracking. *PLoS ONE* 6(5): e19997, 2011.
- [48] Ben J Hardcastle and Holger G Krapp. Evolution of biological image stabilization. *Current Biology*, 26(20):R1010–R1021, 2016.
- [49] R Hengstenberg, DC Sandeman, and B Hengstenberg. Compensatory head roll in the blowfly calliphora during flight. *Proceedings of the Royal society of London. Series B. Biological sciences*, 227(1249):455–482, 1986.
- [50] Mehrdad Hesar, Ali Najafi, and Shyamnath Gollakota. Netscatter: Enabling large-scale backscatter networks. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, Boston, MA, 2019. USENIX Association.
- [51] Aaron M Hoover, Erik Steltz, and Ronald S Fearing. Roach: An autonomous 2.4 g crawling hexapod robot. *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 26–33, 2008.
- [52] Luke Hurst. Tiny microchips glued to bees to track decline, 2015.
- [53] Suk-Won Hwang, Hu Tao, Dae-Hyeong Kim, Huanyu Cheng, Jun-Kyul Song, Elliott Rill, Mark A Brenckle, Bruce Panilaitis, Sang Min Won, Yun-Soung Kim, et al. A physically transient form of silicon electronics. *Science*, 337(6102):1640–1644, 2012.
- [54] Seiko Instruments. MS Lithium Rechargeable Battery MS412FE. <https://www.sii.co.jp/en/me/datasheets/ms-rechargeable/ms412fe-5/>.
- [55] Vikram Iyer, Elyas Bayati, Rajalakshmi Nandakumar, Arka Majumdar, and Shyamnath Gollakota. Charging a smartphone across a room using lasers. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(4):1–21, 2018.
- [56] Vikram Iyer, Justin Chan, Ian Culhane, Jennifer Mankoff, and Shyamnath Gollakota. Wireless analytics for 3d printed objects. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, pages 141–152, 2018.
- [57] Vikram Iyer, Justin Chan, and Shyamnath Gollakota. 3d printing wireless connected objects. *ACM Transactions on Graphics (TOG)*, 36(6):1–13, 2017.
- [58] Vikram Iyer, Maruchi Kim, Shirley Xue, Anran Wang, and Shyamnath Gollakota. Airdropping sensor networks from drones and insects. *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, 2020.

- [59] Vikram Iyer, Ali Najafi, Johannes James, Sawyer Fuller, and Shyamnath Gollakota. Wireless steerable vision for live insects and insect-scale robots. *Science Robotics*, 5(44), 2020.
- [60] Vikram Iyer, Rajalakshmi Nandakumar, Anran Wang, Sawyer B. Fuller, and Shyamnath Gollakota. Living IoT: A flying wireless platform on live insects. *The 25th Annual International Conference on Mobile Computing and Networking*, pages 1–15, 2019.
- [61] Vikram Iyer, Vamsi Talla, Bryce Kellogg, Shyamnath Gollakota, and Joshua Smith. Inter-technology backscatter: Towards internet connectivity for implanted devices. In *SIGCOMM '16*, 2016.
- [62] Jamie J Matic, Michelle Rasmussen, Alan Pollack, William Weeman, Irene Lee, Mark Willis, Steve Garverick, Roy Ritzmann, and Daniel A Scherson. Wireless communication by an autonomous self-powered cyborg insect. *Journal of the Electrochemical Society*, 161, 10 2013.
- [63] Johannes James, Vikram Iyer, Yogesh Chukewad, Shyamnath Gollakota, and Sawyer B Fuller. Liftoff of a 190 mg laser-powered aerial vehicle: The lightest wireless robot to fly. *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8, 2018.
- [64] T. Jang, G. Kim, B. Kempke, M. B. Henry, N. Chiotellis, C. Pfeiffer, D. Kim, Y. Kim, Z. Foo, H. Kim, A. Grbic, D. Sylvester, H. Kim, D. D. Wentzloff, and D. Blaauw. Circuit and system designs of ultra-low power sensor nodes with illustration in a miniaturized gnss logger for position tracking: Part ii—data communication, energy harvesting, power management, and digital circuits. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 2017.
- [65] Gáspár Jékely, Julien Colombelli, Harald Hausen, Keren Guy, Ernst Stelzer, François Nédélec, and Detlev Arendt. Mechanism of phototaxis in marine zooplankton. *Nature*, 456(7220):395–399, 2008.
- [66] Xiaobin Ji, Xinchang Liu, Vito Cacucciolo, Matthias Imboden, Yoan Civet, Alae El Haitami, Sophie Cantin, Yves Perriard, and Herbert Shea. An autonomous untethered fast soft robotic insect driven by low-voltage dielectric elastomer actuators. *Science Robotics*, 4(37), 2019.
- [67] J. M. Kahn, R. H. Katz, and K. S. J. Pister. Next Century Challenges: Mobile Networking for “Smart Dust”. *ACM MobiCom*, page 271–278, 1999.
- [68] Michael Karpelson, Benjamin H Waters, Benjamin Goldberg, Brody Mahoney, Onur Ozcan, Andrew Baisch, Pierre-Marie Meyitang, Joshua R Smith, and Robert J Wood. A wirelessly powered, biologically inspired ambulatory microrobot. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2384–2391. IEEE, 2014.

- [69] Mohamad Katanbaf, Anthony Weinand, and Vamsi Talla. Simplifying backscatter deployment: Full-duplex lora backscatter. *18th Symposium on Networked Systems Design and Implementation*, 2021.
- [70] Carol Ann Kearns, David William Inouye, et al. *Techniques for pollination biologists*. University Press of Colorado, 1993.
- [71] Bryce Kellogg, Vamsi Talla, and Shyamnath Gollakota. Bringing gesture recognition to all devices. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, pages 303–316, 2014.
- [72] Bryce Kellogg, Vamsi Talla, Shyamnath Gollakota, and Joshua R. Smith. Passive Wi-Fi: Bringing low power to Wi-Fi transmissions. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, pages 151–164, Santa Clara, CA, March 2016. USENIX Association.
- [73] John Klingner, Anshul Kanakia, Nicholas Farrow, Dustin Reishus, and Nikolaus Correll. A stick-slip omnidirectional powertrain for low-cost swarm robotics: mechanism, calibration, and control. *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 846–851, 2014.
- [74] Je-Sung Koh, Eunjin Yang, Gwang-Pil Jung, Sun-Pill Jung, Jae Hak Son, Sang-Im Lee, Piotr G Jablonski, Robert J Wood, Ho-Young Kim, and Kyu-Jin Cho. Jumping on water: Surface tension-dominated jumping of water striders and robotic insects. *Science*, 349(6247):517–521, 2015.
- [75] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. Spotfi: Decimeter level localization using wifi. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM '15*, page 269–282. Association for Computing Machinery, 2015.
- [76] Karl Kral. Side-to-side head movements to obtain motion depth cues: A short review of research on the praying mantis. *Behavioural Processes*, 43(1):71–77, 1998.
- [77] Michael F Land and Dan-Eric Nilsson. *Animal eyes*. Oxford University Press, 2012.
- [78] David Le Gac. Reverse costing analysis: Pillcam colon 2. Technical report, System Plus Consulting, September 2013.
- [79] Fleur Lehardt and Claude Desplan. Retinal perception and ecological significance of color vision in insects. *Current opinion in insect science*, 24:75–83, 2017.

- [80] W. Lee, V. Alchanatis, C. Yang, M. Hirafuji, D. Moshou, and C. Li. Sensing technologies for precision specialty crop production. In *Computers and Electronics in Agriculture*, 2010.
- [81] Y. Lee, S. Bang, I. Lee, Y. Kim, G. Kim, M. H. Ghaed, P. Pannuto, P. Dutta, D. Sylvester, and D. Blaauw. A modular 1 mm³ die-stacked sensing platform with low power i²c inter-die communication and multi-modal energy harvesting. *IEEE Journal of Solid-State Circuits*, 48(1):229–243, 2013.
- [82] David Lentink, William B Dickson, Johan L Van Leeuwen, and Michael H Dickinson. Leading-edge vortices elevate lift of autorotating plant seeds. *Science*, 324(5933):1438–1440, 2009.
- [83] Friedrich Lenz, Aleksei V Chechkin, and Rainer Klages. Constructing a stochastic model of bumblebee flights from experimental data. *PloS one*, 8(3):e59036, 2013.
- [84] Discover Life. *Bombus impatiens*. <http://www.discoverlife.org/20/q?search=Bombus+impatiens>.
- [85] Jie Liu, Bodhi Priyantha, Ted Hart, Heitor S Ramos, Antonio AF Loureiro, and Qiang Wang. Energy efficient gps sensing with cloud offloading. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, pages 85–98. ACM, 2012.
- [86] Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R. Smith. Ambient backscatter: Wireless communication out of thin air. *ACM SIGCOMM*, page 39–50, 2013.
- [87] M. Lok, E. F. Helbling, X. Zhang, R. Wood, D. Brooks, and G. Y. Wei. A low mass power electronics unit to drive piezoelectric actuators for flying microrobots. *IEEE Transactions on Power Electronics*, PP(99):1–1, 2017.
- [88] Jess Lowenberg-DeBoer. *The precision agriculture revolution: Making the modern farmer*, 2015.
- [89] Kevin Y Ma, Pakpong Chirarattananon, Sawyer B Fuller, and Robert J Wood. Controlled flight of a biologically inspired, insect-scale robot. *Science*, 340(6132):603–607, 2013.
- [90] Sylvain Mange, E Farrell Helbling, Nick Gravish, and Robert J Wood. An actuated gaze stabilization platform for a flapping-wing microrobot. *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5409–5414, 2017.
- [91] K Mann, TL Massey, S Guha, JP van Kleef, and MM Maharbiz. A wearable wireless platform for visually stimulating small flying insects. In *Engineering in Medicine and Biology Society (IEEE)*, 2014.

- [92] Bjorn H Merker. Insects join the consciousness fray. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 1(9):4, 2016.
- [93] MHGoPower. MiH VMJ PV Cell. http://www.mhgopower.com/cell_description.html.
- [94] RC Miall. The flicker fusion frequencies of six laboratory insects, and the response of the compound eye to mains fluorescent ‘ripple’. *Physiological Entomology*, 3(2):99–106, 1978.
- [95] Saman Naderiparizi, Mehrdad Hesar, Vamsi Talla, Shyamnath Gollakota, and Joshua R Smith. Towards battery-free HD video streaming. *15th USENIX Symposium on Networked Systems Design and Implementation*, pages 233–247, April 2018.
- [96] Saman Naderiparizi, Pengyu Zhang, Matthai Philipose, Bodhi Priyantha, Jie Liu, and Deepak Ganesan. Glimpse: A programmable early-discard camera architecture for continuous mobile vision. *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*, pages 292–305, 2017.
- [97] Rajalakshmi Nandakumar, Vikram Iyer, and Shyamnath Gollakota. 3d localization for sub-centimeter sized devices. *SenSys*, 2018.
- [98] NASA. Drones swarm over beaver dam, 2019. <https://www.nasa.gov/feature/langley/drones-swarm-over-beaver-dam>.
- [99] Natupol. Koppert Biological Systems. <https://www.koppert.com/products/pollination/natupol/>.
- [100] Jeremy E Niven. Neural Energetics: Hungry flies turn down the visual gain. *Current Biology*, 24(8):R313–R315, 2014.
- [101] Claudia Nuñez-Penichet, Luis Osorio-Olvera, Victor H Gonzalez, Marlon E Cobos, Laura Jiménez, Devon A DeRaad, Abdelghafar Alkische, Rusby G Contreras-Díaz, Angela Nava-Bolaños, Kaera Utsumi, et al. Geographic potential of the world’s largest hornet, vespa mandarinia smith (hymenoptera: Vespidae), worldwide and particularly in north america. *PeerJ*, 9:e10690, 2021.
- [102] JL Osborne, IH Williams, NL Carreck, GM Poppy, JR Riley, AD Smith, DR Reynolds, and AS Edwards. Harmonic radar: a new technique for investigating bumblebee and honey bee foraging flight. In *VII International Symposium on Pollination 437*, pages 159–164, 1996.
- [103] A. Ozkurt, R. Gilmour, D. Stern, and A. Lal. Mems based bioelectronic neuromuscular interfaces for insect cyborg flight control. In *IEEE Conference on Micro Electro Mechanical Systems*, 2008.

- [104] Aaron N Parks and Joshua R Smith. Sifting through the airwaves: Efficient and scalable multiband rf harvesting. In *IEEE RFID*, 2014.
- [105] Neal Patwari, Joshua N Ash, Spyros Kyperountas, Alfred O Hero, Randolph L Moses, and Neiyer S Correal. Locating the nodes: cooperative localization in wireless sensor networks. *IEEE Signal processing magazine*, 22(4):54–69, 2005.
- [106] Ted Pella. Gel-Pak. https://www.tedpella.com/storage-boxes-bags_html/gel-pak.htm.
- [107] Matthew Piccoli and Mark Yim. Piccolissimo: The smallest micro aerial vehicle. In *ICRA 2017*, 2017.
- [108] P. Pounds and S. Singh. Samara: Biologically inspired self-deploying sensor networks. *IEEE Potentials*, 34(2):10–14, March 2015.
- [109] Corinne Reichert. Intel and CSIRO create RFID bee backpacks with Edison, 2015. <http://www.zdnet.com/article/intel-and-csiro-create-rfid-bee-backpacks-with-edison/>.
- [110] Timothy Reissman. Mesoscale transduction systems and applications to energy harvesting onboard flying insects. *Cornell University*, 2013.
- [111] Timothy Reissman and Ephraim Garcia. Surgically implanted energy harvesting devices for renewable power sources in insect cyborgs. In *ASME 2008 International Mechanical Engineering Congress and Exposition*, pages 645–653. American Society of Mechanical Engineers, 2008.
- [112] Timothy Reissman, Robert B MacCurdy, and Ephraim Garcia. Electrical power generation from insect flight. In *Active and Passive Smart Structures and Integrated Systems 2011*, volume 7977, page 797702. International Society for Optics and Photonics, 2011.
- [113] J.R. Riley, A.D. Smith, Don Reynolds, A S. Edwards, Juliet Osborne, Ingrid Williams, Norman Carreck, and Guy Poppy. Tracking bee with harmonic radar. *Nature*, 379:29–30, 01 1996.
- [114] Mohammad Rostami, Karthik Sundaresan, Eugene Chai, Sampath Rangarajan, and Deepak Ganesan. Redefining passive in backscattering with commodity devices. *ACM MobiCom*, 2020.
- [115] Michael Rubenstein, Christian Ahler, and Radhika Nagpal. Kilobot: A low cost scalable robot system for collective behaviors. *2012 IEEE International Conference on Robotics and Automation*, pages 3293–3298, 2012.

- [116] William Moy Stratton Russell, Rex Leonard Burch, and Charles Westley Hume. *The principles of humane experimental technique*, volume 238. Methuen London, 1959.
- [117] Jihoon Ryoo, Yasha Karimi, Akshay Athalye, Milutin Stanaćević, Samir R. Das, and Petar Djurić. Barnet: Towards activity recognition using passive backscattering tag-to-tag network. In *MobiSys '18*, 2018.
- [118] A. Saffari, M. Hessar, S. Naderiparizi, and J. R. Smith. Battery-free wireless video streaming camera system. *2019 IEEE International Conference on RFID (RFID)*, pages 1–8, April 2019.
- [119] H. Sato, Y. Peeri, E. Baghoomian, C. W. Berry, and M. M. Maharbiz. Radio-controlled cyborg beetles: A radio-frequency system for insect neural flight control. *2009 IEEE 22nd International Conference on Micro Electro Mechanical Systems*, pages 216–219, Jan 2009.
- [120] Hirotaka Sato, Christopher W Berry, Yoav Peeri, Emen Baghoomian, Brendan E Casey, Gabriel Lavella, John M VandenBrooks, Jon Harrison, and Michel M Maharbiz. Remote radio control of insect flight. *Frontiers in integrative neuroscience*, 3:24, 2009.
- [121] Madeleine Seale, Oleksandr Zhdanov, Cathal Cummins, Erika Kroll, Mike Blatt, Hossein Zare-Behtash, Angela Busse, Enrico Mastropaolo, Ignazio Maria Viola, and Naomi Nakayama. Informed dispersal of the dandelion. *bioRxiv*, 2020.
- [122] Lattice Semiconductor. HM01B0 upduino shield. Technical report, Lattice Semiconductor, 2019.
- [123] NXP Semiconductors. Kinetis ultra low power microcontroller. <https://www.nxp.com/docs/en/user-guide/FRDMKL03ZUG.pdf>.
- [124] Dongjin Seo, Ryan M Neely, Konlin Shen, Utkarsh Singhal, Elad Alon, J.M. Rabaey, Jose M Carmena, and Michel M Maharbiz. Wireless recording in the peripheral nervous system with ultrasonic neural dust. *Neuron*, 91:529–539, 08 2016.
- [125] Michael W Shafer, Robert MacCurdy, J Ryan Shipley, David Winkler, Christopher G Guglielmo, and Ephraim Garcia. The case for energy harvesting on wildlife in flight. *Smart Materials and Structures*, 24(2):025031, 2015.
- [126] Longfei Shangguan and Kyle Jamieson. The design and implementation of a mobile RFID tag sorting robot. In *MobiSys '16*, 2016.
- [127] J. C. Sheldon and F. M. Burrows. The dispersal effectiveness of the achene-pappus units of selected compositae in steady winds with convection. *The New Phytologist*, 72(3):665–675, 1973.

- [128] Sanmin Shen, Zhihong Fan, Jiahao Deng, Xiaowei Guo, Lei Zhang, Guanyu Liu, Qiulin Tan, and Jijun Xiong. An LC passive wireless gas sensor based on PANI/CNT composite. *Sensors*, 18(9):3022, 2018.
- [129] A. J. Shum, J. Crest, G. Schubiger, and B. A. Parviz. Drosophila as a live substrate for solid-state microfabrication. *Advanced Materials*, 19(21):3608–3612, 2007.
- [130] Angela J Shum and Babak A Parviz. Vacuum microfabrication on live fruit fly. In *IEEE MEMS, 2007*, 2007.
- [131] Skyworks. SMS7630-061: Surface Mount, 0201 Zero Bias Silicon Schottky Detector Diode. <https://www.skyworksinc.com/Products/Diodes/SMS7630-Series>.
- [132] Ryan St. Pierre and Sarah Bergbreiter. Toward autonomy in sub-gram terrestrial robots. *Annual Review of Control, Robotics, and Autonomous Systems*, 2(1):231–252, 2019.
- [133] Nicholas J Strausfeld, HS Seyan, and JJ Milde. The neck motor system of the fly *Calliphora erythrocephala*-i. muscles and motor neurons. *Journal of Comparative Physiology A*, 160(2):205–224, 1987.
- [134] Kabkaew L Sukontason, Tarinee Chaiwong, Somsak Piangjai, Sorawit Upakut, Kittikhun Moophayak, and Kom Sukontason. Ommatidia of blow fly, house fly, and flesh fly: implication of their vision efficiency. *Parasitology research*, 103(1):123–131, 2008.
- [135] Vamsi Talla, Mehrdad Hesar, Bryce Kellogg, Ali Najafi, Joshua R. Smith, and Shyamnath Gollakota. Lora backscatter: Enabling the vision of ubiquitous connectivity. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2017.
- [136] Vamsi Talla, Bryce Kellogg, Benjamin Ransford, Saman Naderiparizi, Shyamnath Gollakota, and Joshua R. Smith. Powering the next billion devices with wi-fi. *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, pages 4:1–4:13, 2015.
- [137] Hu Tao, Suk-Won Hwang, Benedetto Marelli, Bo An, Jodie E Moreau, Miaomiao Yang, Mark A Brenckle, Stanley Kim, David L Kaplan, John A Rogers, et al. Silk-based resorbable electronic devices for remotely controlled therapy and in vivo infection abatement. *Proceedings of the National Academy of Sciences*, 111(49):17385–17389, 2014.
- [138] PowerStream Technology. Ultra low weight lithium ion batteries, September 2013. <https://www.powerstream.com/ultra-light.htm>.

- [139] Paul Tenczar, Claudia C. Lutz, Vikyath D. Rao, Nigel Goldenfeld, and Gene Robinson. Automated monitoring reveals extreme interindividual variation and plasticity in honeybee foraging activity levels. In *Animal Behaviour*, volume 95, 2014.
- [140] Michael M Thackeray, Christopher Wolverton, and Eric D Isaacs. Electrical energy storage for transportation—approaching the limits of, and going beyond, lithium-ion batteries. *Energy & Environmental Science*, 5(7):7854–7863, 2012.
- [141] S. J. Thomas, R. R. Harrison, A. Leonardo, and M. S. Reynolds. A battery-free multichannel digital neural/emg telemetry system for flying insects. *IEEE Transactions on Biomedical Circuits and Systems*, 2012.
- [142] John Thorson. Dynamics of motion perception in the desert locust. *Science*, 145(3627):69–71, 1964.
- [143] Panagiotis Vartholomeos and Evangelos Papadopoulos. Dynamics, Design and Simulation of a Novel Microrobotic Platform Employing Vibration Microactuators. *Journal of Dynamic Systems, Measurement, and Control*, 128(1):122–133, 11 2005.
- [144] Deepak Vasisht, Swarun Kumar, and Dina Katabi. Decimeter-level localization with a single wifi access point. In *NSDI 16*, 2016.
- [145] Tat Thang Vo Doan, Melvin YW Tan, Xuan Hien Bui, and Hiroataka Sato. An ultra-lightweight and living legged robot. *Soft robotics*, 5(1):17–23, 2018.
- [146] Steven Vogel. *Life in Moving Fluids: The Physical Biology of Flow-Revised and Expanded Second Edition*. Princeton University Press, 2020.
- [147] Anran Wang, Vikram Iyer, Vamsi Talla, Joshua R. Smith, and Shyamnath Gollakota. FM backscatter: Enabling connected cities and smart fabrics. In *NSDI 17*, 2017.
- [148] Jue Wang, Deepak Vasisht, and Dina Katabi. RF-IDraw: Virtual touch screen in the air using rf signals. In *SIGCOMM '14*, 2014.
- [149] Benjamin H Waters, Alanson P Sample, Pramod Bonde, and Joshua R Smith. Powering a ventricular assist device (VAD) with the free-range resonant electrical energy delivery (FREE-D) system. *Proceedings of the IEEE*, 100(1):138–149, 2011.
- [150] Catrin Westphal, INGOLF STEFFAN-DEWENTER, and Teja Tschardt. Foraging trip duration of bumblebees in relation to landscape-wide resource availability. *Ecological Entomology*, 31(4):389–394, 2006.

- [151] Telissa M Wilson, Junichi Takahashi, Sven-Erik Spichiger, Iksoo Kim, and Paul Van Westendorp. First reports of *vespa mandarinia* (hymenoptera: Vespidae) in north america represent two separate maternal lineages in washington state, united states, and british columbia, canada. *Annals of the Entomological Society of America*, 113(6):468–472, 2020.
- [152] Shane P Windsor and Graham K Taylor. Head movements quadruple the range of speeds encoded by the insect motion vision system in hawkmoths. *Proceedings of the Royal Society B: Biological Sciences*, 284(1864):20171622, 2017.
- [153] R. J. Wood. The first takeoff of a biologically inspired at-scale robotic insect. *IEEE Trans. Robotics*, 24(2):341–347, 2008.
- [154] Yichuan Wu, Justin K. Yim, Jiaming Liang, Zhichun Shao, Mingjing Qi, Junwen Zhong, Zihao Luo, Xiaojun Yan, Min Zhang, Xiaohao Wang, Ronald S. Fearing, Robert J. Full, and Liwei Lin. Insect-scale fast moving and ultrarobust soft robot. *Science Robotics*, 4(32), 2019.
- [155] Wei Xi, Yuan He, Yunhao Liu, Jizhong Zhao, Lufeng Mo, Zheng Yang, Jiliang Wang, and Xiangyang Li. Locating sensors in the wild: pursuit of ranging quality. In *Proceedings of the 8th ACM conference on Embedded Networked Sensor Systems*, pages 295–308. ACM, 2010.
- [156] Daniel Yang and Stacy Wegner. Apple iPhone XS Max teardown. Technical report, TechInsights Inc., October 2018. <https://www.techinsights.com/blog/apple-iphone-xs-max-teardown>.