

Airdropping Sensor Networks from Drones and Insects

Vikram Iyer*, Maruchi Kim*, Shirley Xue*, Anran Wang, Shyamnath Gollakota

*Co-primary Student Authors

{vsiyer,mkimhj,qxue2,anranw,gshyam}@uw.edu

University of Washington

ABSTRACT

We present the first system that can airdrop wireless sensors from small drones and live insects. In addition to the challenges of achieving low-power consumption and long-range communication, airdropping wireless sensors is difficult because it requires the sensor to survive the impact when dropped in mid-air. Our design takes inspiration from nature: small insects like ants can fall from tall buildings and survive because of their tiny mass and size. Inspired by this, we design insect-scale wireless sensors that come fully integrated with an onboard power supply and a lightweight mechanical actuator to detach from the aerial platform. Our system introduces a first-of-its-kind 37 mg mechanical release mechanism to drop the sensor during flight, using only 450 μ J of energy as well as a wireless communication link that can transmit sensor data at 33 kbps up to 1 km. Once deployed, our 98 mg wireless sensor can run for 1.3-2.5 years when transmitting 10-50 packets per hour on a 68 mg battery. We demonstrate attachment to a small 28 mm wide drone and a moth (*Manduca sexta*) and show that our insect-scale sensors flutter as they fall, suffering no damage on impact onto a tile floor from heights of 22 m.

CCS CONCEPTS

• Computer Systems Organization → Embedded and Cyber-Physical Systems.

KEYWORDS

Bio-inspired sensor systems; IoT; Smart farming; Sensor placement

ACM Reference Format:

Vikram Iyer, Maruchi Kim, Shirley Xue, Anran Wang, Shyamnath Gollakota. 2020. Airdropping Sensor Networks from Drones and Insects. In *MobiCom 2020, September 21–25, 2020, London*. ACM, 14 pages. <https://doi.org/10.1145/3372224.3419981>

1 INTRODUCTION

Sensor networks have the potential to transform multiple domains including industrial sensing [21], smart farming [8] and forest management [42]. More topically, climate scientists are deploying sensors for environmental monitoring across diverse terrain ranging from glacier crevasses [50], to volcanoes [61] and wild forests [59].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom '20, September 21–25, 2020, London, United Kingdom

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7085-1/20/09...\$15.00

<https://doi.org/10.1145/3372224.3419981>




Figure 1: Insect-scale sensor released from a 28 mm wide commercial drone. <https://youtu.be/mzwBvYae740>

The overhead of manual sensor deployment however remains a major barrier for wide-spread adoption of sensor networks [38]. Manual sensor deployment is an expensive, labor intensive process that does not scale well with the number of sensors or for deployments over large geographical areas. Further, deploying sensors in difficult terrain like forests and glaciers as well as hard-to-reach spaces inside pipes or near electrical appliances in factories raises safety concerns. This adds the requirement of significant physical training and expertise for those deploying the sensors.

This paper explores the possibility of using aerial platforms to deploy and airdrop wireless sensors from air. In addition to the typical challenges of range and power faced in designing wireless sensors, aerial deployment requires a mechanism to drop the sensor and perhaps most importantly a strategy for preventing structural damage upon impact with the ground. To achieve this, we take inspiration from nature: small insects like ants can fall from trees and even tall buildings but survive the impact [1]. These insects have such a small mass relative to their surface area causing them to fall slowly and reach a low terminal velocity due to air resistance. Given their low terminal velocity and mass, they have little energy to dissipate when they hit the surface.

Inspired by this, we design insect-scale wireless sensors that come fully integrated with an onboard power supply and an low-weight mechanical actuator to detach from the aerial platform. These insect-scale sensors can be released from aerial platforms at large heights above the ground since they achieve a low terminal velocity due to their low mass. This minimalist design approach of making the sensor small and light-weight eliminates the need for complex parachute or glider-style structures [5] to survive impact. Moreover, given their small footprint, a large number of sensors can be attached to a drone. Thus, on a single charge, one drone can release many sensors across the target area, enabling rapid deployment of wireless sensors across large areas.

Insect-scale detachable sensors also open up applications for state-of-the-art aerial platforms that are payload limited.

• *Small drones.* Unlike their larger counterparts, small drones like the SKEYE Nano 2 Camera Drone and the FY804 miniature quadrap-ture weigh only 16 grams and are as small as $2.2 \times 2.2 \times 2$ cm [4].




Figure 2: Insect-scale sensor with our small size, weight and power (SWaP) constrained release mechanism and communication.

These small drones are compelling as an aerial platform to release sensors since they can go into smaller spaces like pipes, crevices as well as traverse under forest and plant canopies.

- *Live insects.* Insects like moths can nearly invert themselves to maneuver obstacles and traverse through narrow spaces better than any existing man-made drones. Further, they can sustain much longer flights than the 5–30 minutes achieved by man-made drones [58]. Finally, prior work has demonstrated the ability to control flight of insects like moths, beetles, dragonflies and locusts [13, 39, 51]. This combined with their maneuverability, makes them a compelling platform to deploy sensors in hard-to-maneuver spaces that is difficult with man-made drones.

Achieving this is however challenging due to the extreme size, weight and power (SWaP) requirements. Specifically, there are two key challenges in designing such a system:

- *SWaP-constrained release mechanism.* We need a release mechanism on the sensor that can detach from the aerial platform. Since the release mechanism has to operate on small drones as well as work with insect-scale sensors, it needs a small weight and millimeter form factor. For operation with live insects, it also needs low power consumption so as not to drain the battery. While some large industrial drones (e.g., Amazon drones) support mechanical release and capture operations [2], they have orders of magnitude larger size, weight and power footprint than our target requirements.

- *SWaP-constrained wireless communication.* Aerial drones can operate at distances of hundreds of meters and thus after detaching from the drones, the wireless sensors should be able to communicate information back to a base station at comparable distances. The size constraints however make this challenging since the small, light-weight antennas required for this design degrade the signal quality. Further, while LoRa chipsets can achieve long ranges, they come in size and weight form factors that are much larger than our target footprint. Backscatter-based systems like Living IoT [24] can achieve low power but are limited to an uplink range of 1-2 meters.

We present the first system that can release wireless sensors from small drones and live insects. To achieve this goal, we make two key technical contributions. First, we introduce *the first SWaP-constrained release mechanism* that can operate on small drones and live insects (see §4.1). This is challenging because existing designs are targeted for industrial-sized drones that use grippers or strong electromagnets which are both heavy and power consuming. Further, this requires the holding force of the mechanism to be significantly greater than the payload in order to provide resilience to vibrations encountered during flight. We instead take a different



Figure 3: Self-deploying insect-scale sensor attached to the lower thorax of a *Manduca sexta* moth (see §6).

approach in which our payload hangs on a pin that we pull out to drop the payload. We use a small 3.5 mg permanent magnet as the pin and design a small solenoid consisting of a coil of thin wire surrounding the magnet. We apply current to the coil that generates a magnetic field to pull out the magnet and release the sensor. Our design weighs 37 mg and consumes 450 uJ of energy to release the sensor; this is multiple orders of magnitude lower weight and energy than existing designs.

Second, we present *the first SWaP-constrained wireless link* that supports ranges of up to 0.8–1 km at data rates of 33–66 kbps. This is challenging since none of the existing long range solutions support our size and weight requirements. We observe that the prevalence of Bluetooth chips in wearable devices like Apple airpods has led to commercial light weight packages that weigh as little as 6.8 mg and are 3×3 mm in size. However Bluetooth 4.0 operates at 1 Mbps and has a 1-byte preamble that together limit its range.¹ To achieve long ranges, we first create a 14-byte virtual preamble to detect the packets with better receiver sensitivity. We accomplish this by re-purposing a number of fields in a Bluetooth header and use them for preamble detection. Second, we reverse-engineer the Bluetooth protocol and its whitening process to increase the symbol duration. We achieve this by setting the appropriate application-level bits, without any firmware or hardware changes. This allows us to program the symbol rate achieved by Bluetooth and create a virtual payload with a bit rate as low as 33 kbps (see §4.2).

Fig. 2 shows our wireless temperature and humidity sensor that includes a microcontroller and a custom release mechanism fabricated using laser micro-machining. We airdrop our 98 mg wireless sensor and release mechanism powered by a 68 mg battery from both small drones and insects like *Manduca sexta* moths.

Results. We highlight the key results in our evaluation.

- *Release mechanism.* Discharging a 100 uF capacitor through the actuator moves the pin mechanism around 500 um in 6 ms, releasing the sensor. Further, it is resilient to vibrations greater than 0.8 N, making it reliable on aerial platforms.

- *Fall dynamics.* The sensor is designed such that the battery, and therefore center of mass is in a corner. Thus, the sensor autorotating (fluttering) about this axis as it falls. This provides additional drag as the light-weight circuit board acts as a rotating wing and further slows its descent likely due to leading edge vortices [36]. It reaches a terminal velocity of 5 m/s within the first 4 m of its descent and does not accelerate further. Our results show that the impact of dropping from 22 m above a tile floor does not damage the hardware.

¹The smallest Bluetooth 5.0 chips that support low rates do not satisfy our size and weight constraints (see §4.2). Further, the lowest rate in Bluetooth 5.0 is only 125 kbps.

- *Communication.* We achieve a -112 dBm receiver sensitivity at 33 kbps; in comparison the SX1276 LoRa chipset achieves a sensitivity of -111 dBm at 37.5 kbps. This translates to a range of 1 km to a USRP receiver with a 6 dBi antenna in an open outdoor environment. Our low sleep power duty-cycling hardware (see §4.4) can achieve a battery lifetime of 1.3-2.5 years while transmitting 10–50 packets per hour.

Contributions. While prior work has deployed large sensors and objects from army helicopters and industrial drones [2, 5, 49], we are not aware of prior efforts that deploy wireless sensors from small drones and insects. We take inspiration from insects such as ants and design insect-scale wireless sensors that achieve a low terminal velocity and can survive the impact force. We design the first SWaP-constrained release mechanism for small drones as well as a SWaP-constrained wireless link that can support ranges of up to a kilometer. We also show proof-of-concept feasibility of localizing the sensor from these long-range transmissions. Finally, we attach the wireless sensors to small drones and moths to demonstrate that the release can be wirelessly triggered from a base station.

<https://github.com/uw-x/airdrop-sensors>

2 MOTIVATING APPLICATIONS

We outline the motivations scenarios in detail in §6. At a high level, small drones can be used to deploy sensors in hard to reach spaces like pipes and small openings. These and other larger drones can carry multiple insect-scale sensors, enabling multiple sensors to be deployed on a single charge. Since consumer drones can have ranges of up to a kilometer, they can be used for large area deployment. The motivating scenario we envision is that the drone flies to the desired locations and a wireless command is transmitted from the operator to release a wireless sensor at each location. More ambitiously, insects such as moths and beetles have been shown to carry close to a gram of payload [23, 60]; thus they can each carry and deploy multiple of our insect-scale sensors. We describe detailed scenarios for deployment using insects in §6.

3 OUR INSECT-INSPIRED APPROACH

Our goal is to design sensors that can be released from mid-air. To do this, we imitate small animals such as insects that are not harmed even when falling from extreme heights.

To understand why this is, we begin by investigating how much force an object (e.g., a sensor) will experience when dropped. The total kinetic energy at the time of impact is $\frac{1}{2}mv^2$. This energy is transferred in the impact as the dissipation energy. Here, v denotes the velocity of the object and m is its mass. As an object falls it experiences a downward force due to gravity as well as a counteracting force due to the drag caused by air resistance. We denote g as the acceleration due to gravity, C_d as the drag coefficient, ρ as the density of the fluid (air) and A as the projected area of the object. The net force on the object (e.g., sensor) can be written as, $F_{net} = F_{gravity} - F_{drag}$. This can be rewritten as, $ma_{net} = m\frac{dv}{dt} = mg - \frac{1}{2}\rho v^2 AC_d$.

The above equation shows that as an object falls, its velocity increases, but this also causes an increasing drag force in the opposite direction. As a result, the object reaches a terminal velocity

when it experiences zero net force and therefore does not accelerate further as it continues to fall. We integrate the above equation and solve for the velocity, $v(t) = \sqrt{\frac{2mg}{\rho AC_d}} \tanh\left(t\sqrt{\frac{g\rho AC_d}{2m}}\right)$. Taking the limit as t approaches infinity, yields the following equation for the terminal velocity: $V_t = \lim_{t \rightarrow \infty} v(t) = \sqrt{\frac{2mg}{\rho AC_d}}$. Note that since the kinetic energy is $\frac{1}{2}mv^2$, the impact energy when the object reaches its terminal velocity is given by, $F_{impact} \approx \frac{1}{2}mV_t^2 = \frac{m^2g}{\rho AC_d}$ [14].

The above equations show that the terminal velocity is directly proportional to the square root of the object’s mass and inversely proportional to the square root of its area. While prior designs including aircraft take the approach of significantly increasing the surface area to provide gliding and lift, we take the alternate approach of reducing the sensor’s mass. Since the force of impact is proportional to the square of the object’s mass and inversely proportional to its area, reducing its mass can decrease the impact force much faster than increasing its surface area. Thus, despite having a smaller surface area, insects like ants have a small impact force and can survive falls due to their incredibly small mass.

Suppose we have a 200 mg wireless sensor modeled as a flat plate ($C_d \approx 2$) with an area of $8 \text{ mm} \times 10 \text{ mm}$ falling in room temperature air with $\rho = 1.225$. Assuming that it is being dropped from a height of 400 ft (which is the FAA limit for drones), this translates to a terminal velocity of 4.5 m/s and an impact energy of less than 2 mJ. In contrast, an 80 kg adult human with a projected area of 0.7 m^2 , has a 47.3 m/s terminal velocity and 3786 J impact energy.

4 SYSTEM DESIGN

Our system has four components: 1) a release mechanism that drops the sensor from the air and can be carried by small insects and drones, 2) a long range communication link that can send sensor data back to the access point, 3) the ability to localize the sensor if it is dropped from an uncontrolled flying insect and 4) sensor hardware that integrates all these components and achieves over an year of operations and enables a long battery life.

4.1 SWaP-Constrained Release Mechanism

4.1.1 Requirements. Aerial sensor deployment requires a mechanism to drop the sensor payload. This mechanism must satisfy three key requirements. First, in order to fly on the back of insects as small as moths or small drones, the deployment mechanism itself must be light-weight. The mechanism should be designed to be a fraction of the total weight and dimensions of the payload. Second, since the actuation mechanism can be deployed from insects as platforms, they must be self-contained with the sensor and must be able to run on the small, low-drain battery which fits within our weight budget. Third, the mechanism must be able to release the sensor at a controlled time. Doing so requires that it be able to support the weight of the sensor payload, and also be robust to the vibrations of the insect or drone during flight.

4.1.2 Design space. Existing off-the-shelf solutions do not satisfy our requirements. So we explore the design space.

- *Motors and servos.* These are the most common actuation mechanisms. However, it is well known that such actuators are generally




Figure 4: Mechanism with (b) attached sensor, (c) moving magnet and (d) falling sensor. https://youtu.be/dEcDV_k5qT4

heavy and inefficient when scaled to millimeter sizes [27]. For example, even the smallest commercial motors weigh 360 mg and consume over 200 mW of power [41].

- *Electrostatic actuators.* Electrostatic adhesion occurs between surfaces with opposite charge. While this was used to hold a 100 mg robot in [17], it did not include on-board electronics and used an external 1000 V power supply. The challenge is that state of the art lightweight boost converters [26, 27] can only achieve 300 V and capacitors that tolerate 1000 V are large and heavy.

- *Piezoelectric actuators.* These are commonly made of piezo elements layered on one or both sides of another material. Applying an electric field across the piezo element produces a change in strain which causes the whole structure to deflect. Piezos are attractive due to their low power combined with their relatively high output force and small sizes. However, piezos require drive signals of hundreds of volts, requiring a boost converter. Light-weight boost converters weigh 73.7 mg and require more than 50 mW [27]. This would make the mechanism and added electronics prohibitive.

- *Electromagnets and electropermanent magnets.* While an electromagnet could hold the sensor, it would require applying current constantly to maintain a magnetic field which is incompatible with our low-power system. An alternative design is an electropermanent magnet (EPM), a permanent magnet in which the external magnetic field can be switched on and off using a pulsed electrical signal instead of a constant current [33]. The challenge is that switching the magnetic state of an EPM requires a high current pulse and can result in a high total energy for release [33]. High currents incur greater heating losses in the thin wires. In our experiments when building small coils we observe that this can cause melting of adhesives like cyanoacrylate glue.

4.1.3 *Our design.* In addition to the SWaP requirements, our release mechanism satisfies three additional constraints.

- *Resilience to vibrations.* Drones and insects typically use rotating propellers or flapping wings that require periodic motion to generate lift. This causes the body of the flying platform to experience vibrations during flight. As a result, our release mechanism should be designed to be robust to vibrations and avoid accidental release.

- *Support sensor weight.* The release mechanism must support the weight of the sensor. While the battery and electronics only weigh around 100 mg, the actuator must support that weight during flight with vibrations as well as apply sufficient force to release it.

- *Zero static power.* To maximize battery life of the sensor system after deployment, we require zero static power to hold the sensor in place during flight. Doing so will significantly reduce the energy requirements of the system over a design that actively holds the

payload in place during flight. Moreover, if the mechanism does not require static power to operate, the battery life is independent of the drone or insect flight time prior to deployment.

Our key observation is that the approaches above such as electropermanent magnets or electrostatic actuators are designed with the payload hanging from the magnet. This requires the holding force of the mechanism to be significantly greater than the payload in order to provide resilience to vibrations. We instead take a different approach in which our weight hangs on a pin that we pull out to drop the payload as shown in Fig 4. The pin fits through a hole in a piece of material attached to the payload and the release mechanism holds it together while it is in the position on the right. As it slides out of the holes to the left there is no longer any material holding the payload and the sensor falls.

This design has a number of advantages. First, the weight that the pin can support depends on its material properties and by using a metal cylinder, this structure can hold large amounts of weight. Second, the force required to pull out the pin is in a direction orthogonal to the force of gravity on the payload. In contrast to applying a force equivalent to αmg to hold up the payload where m is the mass of the sensor payload, g is the acceleration of gravity, and $\alpha > 1$ is some safety factor chosen for resilience to vibration, applying a force on the orthogonal axis is significantly lower.

Designing this mechanism requires four key components. First, it requires an actuator to pull out the pin. Second, a restoring force that will keep the pin from falling out during flight. Third, a frame or structure to hold the above components together and allow the pin to slide. Finally it requires an electronic trigger circuit to release.

- **Pin and Electromagnetic actuator.** We use a magnet as a pin that can slide through the holes in the frame material. To build an actuator to move the pin, we design a small solenoid that consists of coil of wire surrounding the permanent magnet. Applying current to the coil generates a magnetic field, which in turn applies a force to move the magnet. By generating sufficient current, we can move the magnet to the left towards the coil and as a result detach. At a high level, the force of a solenoid depends on the strength of the permanent magnet used. Specifically, the Lorentz force acting on the magnet can be written as [10, 11]: $F_{Lorentz}(t) = B_{radial}(x(t)) \cdot I_{coil}(t) \cdot l_{coil} \cdot n_{turns}$. Here, $B_{radial}(x(t))$ is the magnetic field seen by the coil which varies with the position of the magnet $x(t)$. Because the magnetic field strength has an inverse square law decay with distance, we minimize the difference between the magnet width and the inner coil diameter down to our fabrication capabilities (125 μ m). I_{coil} is the current through the coil which varies with the magnet's position and the current applied by the capacitor. We choose low series resistance capacitors that can supply a high peak current to produce maximum force. The

remaining terms describe the length of wire in the circumference $l_{coil} = 2\pi r_{coil}$ and $n_{turns} = 60$ is the total number of turns of wire in the coil. We choose the radius of the coil, r_{coil} , to minimize the air gap as explained above and choose to design a coil with multiple concentric windings to increase the number of turns while minimizing the width and therefore size of the actuator. The small diameter minimizes the air gap. While thinner wires could be used, they significantly increase the complexity of fabrication.

- **Holding force mechanism.** Solenoids use springs to provide some restoring force to hold the magnet in a fixed position and be resilient to vibrations. Since designing springs in our form factor is challenging, we take a different approach that leverages magnetic force. Since our actuator uses a permanent magnet, we place a metal plate on the right side of the actuator and leverage the magnetic attraction to hold it in place as shown in Fig. 4. Further, we add a thin carbon fiber stop between the metal and the magnet. This approach has two advantages over a spring. First, moving the carbon fiber stop to increase the distance between the metal plate and magnet allows for fine tuning of the attractive force. In contrast, a spring would have to be redesigned for each target force. The metal plate is also flat and occupies a smaller space than cantilever beam springs. Second, the force exerted by a typical spring changes linearly with distance whereas magnetic forces vary with an inverse square law relationship. This is ideal as it provides a strong attractive force when holding the payload, but when the solenoid is powered and begins to move the magnet away, the force decreases rapidly with distance. We use an 2×2 mm piece of 125 μ m thick stainless steel plate as the metal and place it 620 μ m from the magnet.

- **Frame structure.** We construct a lightweight frame out of 90 μ m thick sheets of carbon fiber and 300 μ m thick carbon fiber rods. The grey structure in Fig. 4 consists of two parallel carbon fiber rods on which another rectangular piece of carbon fiber with a magnet glued in the center can slide back and forth freely. Two parallel pieces of carbon fiber with a hole that accommodates the magnet are mounted on the right side, which form a slot. A separate piece of carbon fiber with a hole attached to the payload can be placed in this slot, allowing the magnet to act as a pin holding it in place. A second small 0.3 mm diameter magnet on the left begins to attract the pin as it slides out. This creates a bi-stable structure that keeps the pin retracted and make this operation robust.

- **Electronic trigger circuit.** The solenoid requires a pulse of current to move the magnet and its force scales with current. The lightweight, low drain battery (Seiko MS412FE) in our system can supply less than 30 mA due to its internal resistance of 100 Ω . To supply a higher current pulse, we charge a capacitor which can quickly dissipate power into the solenoid and trigger the motion. Our circuit consists of a low resistance capacitor connected directly to the microcontroller through a current limiting resistor for charging as shown in Fig. 4. A transistor then connects the capacitor to the solenoid. When the microcontroller turns on the transistor, the capacitor provides the desired impulse.

4.2 SWaP-Constrained Communication

4.2.1 Requirements. Our sensors require a wireless link to transmit data and receive ACKs. The form factor of the electronics should be small and lightweight, around tens of milligrams with all of

the components. Also, minimizing form factor and weight requires using small, limited capacity batteries, thus requiring low power consumption. Further, aerial sensor deployment allows for rapid coverage of large areas, requiring communication links that must operate over large ranges in open spaces. Finally, since our sensors will not send large amounts of data, we should optimize for longer ranges and low power while targeting rates of tens of kbps.

4.2.2 Design space. Prior work uses backscatter [19, 24, 25, 29, 37, 46] however given the antenna losses and the nature of backscatter, its range is limited with our small form factor devices. Here we explore radio-based solutions instead.

- **Custom ASIC radios.** Creating a custom system on a chip can allow for maximum optimization of form factor and power [12]. However, the development process for ASICs tend to be expensive and time-consuming, taking anywhere from 6 to 18 months. In contrast, many commercial off-the-shelf (COTS) radio chips now include microcontrollers and offer many of the advantages of an ASIC, but with significantly greater programming flexibility. Moreover, their existing economies of scale allow for immediate scalability. Thus, we opt to use COTS components to design our system.

- **LPWAN radios.** Low power wide area network (LPWAN) protocols like LoRa are designed for long ranges with data rates of 8-37.5 kbps [20, 53]. A LoRa chip from Semtech however weighs over 40 mg without external components, making it too heavy for our application. It also requires 40 mW to transmit, which is more than double that of Bluetooth. While there does exist a smaller LoRa chip, it also consumes 60-200 mW [40]. Zigbee chips like the CC2530 have similar issues [22].

- **BLE radios.** BLE chips consume 20 mW or less and are embedded in wearable devices. This has led manufacturers to fabricate smaller and light-weight packages that support Bluetooth – the wafer level, chip scale package of the Nordic NRF52832 Bluetooth-enabled microcontroller is 6.8 mg and 3×3 mm in size. The limitation is that Bluetooth 4.0 has a limited sensitivity at 1 Mbps, significantly limiting its range. While Bluetooth 5.0 supports longer range modes with a data rate of 125 kbps, they are limited for two main reasons: 1) the smallest package Bluetooth 5.0 chips that support low data rates do not satisfy our size and weight constraints, and 2) they only support the lowest data rate of 125 kbps. Finally, we note that while academic BLE 4.0 chips are being designed [32, 35], they are not widely available and do not provide the reliability and programmability of commercial SOCs.

We take a different approach: we use the above small, low-power hardware of Bluetooth 4.0 chips and reverse engineer the protocol to piggyback a lower rate modulation scheme on its transmissions that can achieve our target range. Doing so allows us to make the same trade-off as LPWAN protocols that reduce their data rate to achieve a longer range. The “uplink” of our system consists of a small BLE chip, antenna and associated software used to transmit sensor data to a custom base station. The “downlink” consists of the basestation with an amplifier to send ACKs to the sensor device.

4.2.3 SWaP-constrained uplink. We reverse engineer the Bluetooth beacon packet structure. BLE beacons are broadcast packets transmitted on three primary advertising channels: 37 (2402 MHz), 38




Figure 5: SWaP-constrained uplink packet structure.

(2426 MHz), and 39 (2480 MHz). These broadcasting messages do not require the complexity and energy overhead of exchanging packets needed to pair and connect with a receiver. However, in order to decode these packets at higher receiver sensitivity, we need both a preamble and a payload that can support such long ranges.

Virtual preamble. The first step required to extend the range of BLE transmissions is to detect the packets at low SNRs. As shown in Fig. 5, BLE 4.0 uses a 1-byte preamble. To understand how this impacts range, we connect the output of an NRF52832 BLE chip to a series of attenuators in a shielded metal box and correlate for the preamble on a USRP E310 SDR. The resulting packet error rate (PER) indicates that the 1 byte preamble used for standard BLE 4.0 can only perform well down to around -92 to -94 dBm. This is consistent with the sensitivity of typical BLE chips.

To achieve a higher receiver sensitivity, we leverage the BLE beacon packet structure to create a virtual preamble that is longer than the 1 byte used in the Bluetooth 4.0 standard. As shown in Fig. 5, a beacon packet starts with a 1 byte preamble consisting of a sequence of alternating zeros and ones. This is followed by the Access Address field which indicates the packet type and is also fixed for advertising packets. The next eight bytes consist of a field indicating the length of the packet, followed by the advertiser address. This means that the first 13 bytes of the packet are constant for a device. We can use all these bytes as a virtual preamble to create the ability to detect packets at much lower SNRs. Further, we can add an additional byte to the preamble by setting the first byte in the Bluetooth payload (payload length) to a constant value. Our results in §5.2 show that a 14-byte virtual preamble allows decoding packets at -112 dBm.

Virtual payload. The physical layer of Bluetooth 4.0 uses binary Gaussian Frequency Shift Keying (GFSK) to send data at a rate of 1 Mbps. At a high level, if we set the application-layer payload to a sequence of N ones to represent the ‘1’ bit and a sequence of N zeros to represent the ‘0’ bit, we can reduce the bit rate to $\frac{1}{N}$ Mbps. The challenge however is that Bluetooth uses data whitening. Specifically, it uses a 7-bit linear feedback shift register (LFSR) circuit with the polynomial $x^7 + x^4 + 1$. Given an initial seed, the circuit outputs a sequence of bits that are used to whiten the application data by XORing the data bits with the bits output by the circuit.

We reverse this whitening process to create the desired sequence of ones and zeros at the target bit rate in our virtual payload. To do this, we note that the seed for the LFSR is simply the channel number, which makes it possible to reverse the sequence given the channel on which beacon is being transmitted. Thus, we can deterministically compute the whitening sequence used by the chipset. By XORing our sequence of N repeating 1s and 0s with the

whitening sequence, we expand the symbol length for each GFSK symbol, thus reducing bit rate while increasing range.

With 30 bytes in the BLE beacon payload, this translates to 3, 2 and 1 bytes at around 100, 66 and 33 kbps respectively. While this is a small payload, this is sufficient for applications where each sensor reading (e.g., temperature, humidity) can fit into one byte of data. Longer data could also be spread across multiple packets. Given our small payload, we use a lightweight parity virtual bit.

We note the following additional techniques in our design.

- *Operating across BLE channels.* BLE chips re-transmit the advertising packets using the same data on three channels. Since we only reverse the whitening sequence for one frequency, we set a channel mask in our firmware to send the de-whitened data on only the set frequency. To send data on another frequency, we set up an internal software timer which, after firing, updates the channel mask and recalculates the updated whitening sequence for the desired channel. While recalculating the whitening sequence, we are also free to change the data in the packet, allowing us to assign different types of data for each particular frequency. For example, we can send temperature data on channel 37, humidity data on channel 38, and accelerometer data on channel 39.
- *Identifying multiple sensors.* We need an ID embedded in each of the packets to identify the wireless sensor that is transmitting. To avoid the overhead of explicitly sending IDs, we leverage the Bluetooth protocol which embeds a 48-bit advertiser address that corresponds to the MAC hardware address. On our device and other common Bluetooth chips this is set by default to a random value and can also be defined in software. We first register the MAC addresses on all our wireless sensors. For each received packet, the basestation correlates at the signal-level with each of the registered MAC addresses and uses the ID that maximizes correlation.
- *Repetition versus other codes.* Since our de-whitening approach enables us to transmit any bit pattern, in principle, we can use more complex codes (e.g., convolutional codes). The decision to use repetition coding is two fold: 1) it reduces computational overhead on our microcontroller and 2) it results in long sinusoidal segments of the packet that can be used for our localization algorithm in §4.3.
- *Using Bluetooth 5.0.* While Bluetooth 5.0 supports 125 kbps, it is difficult to perform de-whitening on these packets. The long range mode uses advertisement extensions, that first broadcasts an auxiliary pointer on a primary advertising frequency. This pointer contains a randomly selected secondary data channel and time offset at which the actual data of the payload will be transmitted [56]. This auxiliary pointer is abstracted away from developers and is not available in software, making it hard to reverse the whitening sequence. In order to compensate for poor antenna performance on small, lightweight devices we focus this work on using the more flexible Bluetooth 4 standard to achieve further data rate reduction beyond what is available in Bluetooth 5. With future chips or firmware that expose these pointers or the sequence, this technique could be used with Bluetooth 5 as well.

4.2.4 SWaP-constrained downlink. The basestation may trigger the release remotely, send updates to schedule transmissions, or change configuration settings. We leverage the asymmetry in transmit power between the insect-scale sensor and the basestation. On the




Figure 6: Vibration testing. <https://youtu.be/6eFevIMv-YM>

insect-scale devices, due to power constraints, we cannot transmit greater than 4 dBm. In contrast, the basestation can repeatedly transmit advertising packets at 1 Mbps using a high-power amplifier. To stay within FCC limitations we use a transmit power of 30 dBm, effectively increasing our link budget to levels similar to our long range uplink. This allows the insect-scale sensor to decode the 1 Mbps transmissions at a long range using its BLE chipset.

4.2.5 Low-power ACKs and MAC protocol. Running the receiver takes almost as much power as running the transmitter. Thus, we minimize the amount of time the receiver on the wireless sensor is ON. In our scheme, after a sensor wakes up and transmits its data, it transitions to receive mode for a fixed short period of time to receive the ACK. The basestation is ON all the time and transmits the ACK immediately after successfully receiving the uplink packet. At this time, the low power device goes back to sleep.

Since Bluetooth does not perform carrier sense, we use a simple Aloha MAC protocol which is effective for heavily duty-cycled networks. For a network with 1000 insect-scale sensors, the channel occupancy is less than 0.21% even with each of the sensor transmitting 20 packets per hour, with each packet occupying only 376 μ s.

4.2.6 Interference from Wi-Fi and other BLE transmissions. Many of our application scenarios are outdoors (e.g., farms) without power infrastructure and backhaul links required to support Wi-Fi. However, to prevent collisions with Wi-Fi packets we can use the *CTS_to_Self* packet. Because the sensor is programmed to transmit at a known interval, after receiving one transmission the basestation can schedule a *CTS_to_Self* packet to be transmitted by the basestation before the expected Bluetooth packet. The *CTS_to_Self* can reserve the channel for up to 32 ms and prevent other Wi-Fi devices from transmitting concurrently [16, 28]. To filter out other BLE advertisements at the basestation, we leverage the fact that we are always maximizing the length of the payload in our virtual packets. Because the received packets must be the maximum possible length, we can filter away any shorter or longer packets on the same frequency channel. Further, since existing BLE transmissions do not use de-whitening, they cannot be properly decoded.

4.3 SWaP-Constrained Localization

While drones may know the location where they release the insect-scale sensor, applications involving small insects can release the sensor at unknown locations (see §2). We build on prior techniques that localize wireless devices [15, 34, 44, 62]. Recent work localizes Bluetooth devices using time-of-arrival (TOA) techniques [9]. This however combines the phase information across all the 40 Bluetooth channels to compute the time of arrival information. In

our scenario, we only have information on three narrowband advertising channels resulting in phase ambiguity. We instead use Angle-of-Arrival (AoA) with a multi-antenna basestation. By using the angles from multiple basestations, we can triangulate to obtain the sensor’s location. Our application scenarios with insects are in outdoor environments like farms which have reasonable line-of-sight conditions, allowing us to identify a direct path to the deployed sensor from the basestation.

4.3.1 Finding AoA. To address multipath and identify the incident angle corresponding to the direct path, AoA techniques use algorithms such as MUSIC [52]. Since we do not know the number of incident angles a-priori, MUSIC algorithm is sometimes unstable because of a wrong estimate of the number of incident angles. We find that Capon algorithm [18] is most stable while achieving acceptable accuracy in our scenarios. This algorithm tries to beamform into each possible angle using Minimum Variance Distortionless Response (MVDR) beamformer, which minimizes the signal from other angles while keeping the signal from the target frequency intact. The peaks of the total energy among the results of each angle are then identified as the angles of arrival. Additionally, we use spatial smoothing before the processing since the signals from each direction are coherent [54].

BLE uses Gaussian Frequency Shift Keying (GFSK) where bit 0 and 1 are modulated to two different frequencies, f_0 and f_1 . However, Bluetooth also uses Gaussian filtering to smooth the frequency transition in the transmitted signal, which results in unstable channel information at each frequency and thus inaccurate phase measurements. Since our SWaP-constrained uplink described in §4.2.3 uses repetition coding to create a long sequence of bit 0 and 1, the frequency settles at f_0 and f_1 long enough for computing phase.

4.3.2 Combining three BLE channels. The three advertisement Bluetooth channels occupy three spread out frequencies in the ISM band, providing frequency diversity. We can combine the information from all three channels to improve accuracy. Specifically, the received signal at antenna k can be written as a combination of N multiple paths as, $S(k) = \sum_{i=0}^N S(t)A_i \exp(j2\pi f k \frac{d}{c} \cos \alpha_i) + N_k$. Here, A_i is the attenuation factor of path i , f is the frequency of the signal, d is the separation between two antennas, c is the speed of light, N_k is noise, and α is the angle of arrival we are trying to solve. The channels centered at each of the three BLE advertisement frequencies can now be considered as a virtual antenna array where the spacing between each virtual antenna is $\frac{fd}{f_0}$, where f_0 is the frequency of one of the channels used as a reference channel. Thus by combining the channel information across all the three advertising frequencies we can increase the SNR of our AoA measurements.

4.4 Battery & Power Optimizations

For environmental monitoring applications, the battery life of our deployable wireless sensor must last for several months to a year in the field. Simultaneously, the battery itself must also be small and lightweight so that it can be carried on small drones and live insects. Conditions like temperature and humidity don’t change too drastically, allowing power saving by limiting the number of uplink packets per hour. The size, weight, and application requirements of the system drive our decision to pick a 1 mAh lithium ion battery.




Figure 7: Release Mechanism Results. (a) Displaced distance and capacitor voltage versus time, (b) displaced distance versus different capacitor, (c) CDF of release time, and (d) displaced distance versus peak current.

A challenge in achieving year-long operation is that the losses due to sleep current begin dominating the battery lifetime. The sleep mode current of our nRF52832 Bluetooth SOC is around 350nA. On our low-capacity battery, this translates to a battery time of 4 months even without transmitting any Bluetooth packets. To reduce sleep current, we integrate a TPL5111, a hardware timer that has an operating current of 35nA. Integrating TPL5111 results in a 10x reduction in current consumption while sleeping between transmitted packets. The timer weighs 13.5 mg, which while not insignificant, is essential to achieve the years-long operation time.

5 RESULTS

The hardware in Fig. 2 is constructed from pieces of 90 μm thick carbon fiber laminate composed of 3 orthogonal layers to increase strength. The laminate is made by pressing the carbon fiber layers together in a heat press. We then use a 355 nm diode pumped solid state laser to cut pieces of carbon fiber frame and assemble them using cyanoacrylate (CA) glue. The coil is fabricated by wrapping 46 AWG magnet wire around a 1 mm cylindrical support. The inner support is then removed and the coils are held together using glue. The circuit is fabricated on a 12.5 μm thick sheet of polyimide coated with a layer of 12.5 μm thick copper. The substrate is covered in an etch resist that is removed with laser and etched in Ferric chloride.

Our ultra-lightweight insect-scale transmitter uses the Nordic NRF52832 Bluetooth chips. We use a two layer flex PCB consisting of copper, adhesive, and polyimide layers which includes the coverlay total 104 μm in thickness. The PCB is fabricated and assembled using standard flex PCB manufacturing techniques. We connect a chip antenna (Johanson 2450AT14A0100) to the Bluetooth SoC. However this antenna is designed for a thicker substrate PCB with a larger ground plane and an additional trace extending from the end. We instead attach a 41 AWG wire to the end of the chip antenna to improve its performance. We connect the antenna to a VNA and trim the wire until it achieves resonance at 2.45 GHz, a length of around 8 mm. We note that while this thin wire does extend outside the dimensions of the PCB, it is flexible and does not impede the motion of small insects or drones.

We compare it to a 3 dBi monopole antenna in an anechoic chamber. We connect each antenna to a USRP transmitting at 2.45 GHz and measure the received power on a spectrum analyzer placed 3.4 m away. We place the antenna in different orientations and find that aside from a null along the axis of the feed line, the measured power is 3-8 dB lower than the reference antenna. The antenna weighing less than 2 mg. All of the above components are placed under a microscope and soldered using a hot plate. The weight of the components in the resulting assembly is detailed in Table 1.

Component	Weight (mg)
Microcontroller + antenna + sensor	47
Low power timer	13.5
Release Mechanism	37
Magnet	3.5
Coil	4.5
Frame	13
Capacitor + Switch	16
Battery	68

Table 1: Weight of various components in our system.

5.1 Release and Dropping dynamics

5.1.1 Release mechanism. We first measure the motion of the magnet (pin) as it moves. We use a laser distance sensor (Keyence IA-030) that outputs a voltage proportional to distance at a rate of 3 kHz. This allows us to measure the precise motion of the magnet as we apply a pulse of current and determine whether it moves far enough to release the payload. We begin by connecting a capacitor to a mechanical single pole dual throw switch with negligible resistance. The capacitor is first connected to a 3 V DC power supply to charge it, and upon toggling the switch, it is disconnected from the power supply and connected to the release mechanism. Fig. 7(a) shows that the capacitor voltage decreases over time as it supplies power to the actuator which then moves the magnet. As the magnet approaches the far end of the release mechanism it reaches the magnetic stop and remains at the fixed distance of approximately 500 μm .

We repeat this measurement using three different capacitors and show the results in Fig. 7(b). The plots show that a 47 μF capacitor is unable to move the magnet far enough to release, whereas two different 100 μF capacitors with different equivalent series resistance (ESR) are able to. We observe that the length of the pulse provided by a 100 μF capacitor is 3.5 ms. This suggests that the 2.2 ms pulse provided by the 47 μF capacitor is too short to actuate the magnet. To understand the current required to achieve actuation, we perform the same measurements with increasing series resistance to determine the peak current required. Fig. 7(c) shows the maximum distance reached by the magnet indicating that below 2.4 A the mechanism fails to completely release.

Next, to understand the latency of our mechanism we measure the amount of time it takes for the magnet to move from one side to the other. We use the same setup described above to measure the motion with a 100 μF capacitor 10 times and plot the distribution of the release time in Fig. 7(d). This shows that the actuator releases within 7 ms. We note that the actual time it takes the payload to fall may be slightly higher due to friction as it slides out.

Finally, we evaluate robustness to vibrations that could be encountered on aerial platforms. To do this, we attach our mechanism




Figure 8: Dynamics of the drop. (a) CDF of time when sensor starts fluttering when dropped from 8 m. (b) Average velocity of the sensor over the final 1 m before impact when dropped from different heights. The terminal velocity is around 5–6 m/s. (c) Angle of displacement caused by different wind speeds. Fluttering video demonstration: <https://youtu.be/qkD5lh4c0e8>.

to an eccentric rotating mass (ERM) vibration motor (Seeed Technology 1027). Fig. 6 shows the release mechanism secured to a 254 μm thick sheet of FR4 which is glued to the vibration motor held in a pair of tweezers. A test weight of 100 mg is hung from the release mechanism. The video shows the vibration motor first powered at 3 V resulting in a rotation speed of 10,000 RPM and a force of 0.8 N. Next we increase the voltage to 5 V. We observe that this is outside the rating on the datasheet and causes visibly stronger vibrations and produces more sound. However, the weight is still secure to the release mechanism in both cases.

5.1.2 Dropping dynamics. We evaluate aspects of the drop.

- **Impact on hardware.** We begin by dropping the sensor from varying heights up to 22 m off of a sixth storey balcony onto a hard tile surface. At each height, we repeat the release five times. After dropping the device we place a BLE receiver close to the device to ensure good SNR and observe no packet errors.
- **Fluttering behavior.** While performing these experiments we noticed that above some height the sensor begins to spin about its center of mass. The structure of the sensor consists of a thin, flexible PCB with the most dense parts such as the battery and release mechanism mounted at one end. This is similar to the structure of certain seeds which have been observed to autorotate as they fall [49]. We drop the sensor ten times from a height of 8 m and record its descent using a high speed camera. We analyze the videos to measure the time of release and the time at which the sensor begins to rotate and plot a CDF of this distribution in Fig. 8(a).
- **Terminal velocity.** This spinning behavior, which typically occurs within half a second of release seems to cause increased drag and reduce the terminal velocity of the sensor. Next we measure the terminal velocity empirically by dropping the sensor from increasing heights and record its trajectory over the last meter before impact with a high speed camera. We use the video to determine the time at which it reaches 1 m and the time of impact which we use to calculate the velocity over this distance. As shown in Fig. 8(b), the device reaches around 5 m/s terminal velocity above a height of 4 m and does not noticeably increase after that. The results above show that the sensor slows due to spinning, reaches a terminal velocity, and survives impact.
- **Wind effect.** We begin by dropping the sensor from a height of 2 m and placing a large box fan 1 m above the ground. We observe that even when blowing at the max speed of 4.1 m/s the sensor falls too fast, and has too little surface area to be affected, suggesting small wind gusts over a small height will have little impact. To evaluate sustained wind over the descent, we instead hold the

sensor 20 cm in front of the fan when running at different speeds and measure the landing position. The results, shown in Fig 8(c), indicate a maximum angle deviation of 25 from the point of release. This shows the potential for ambiguity in the landing location under windy conditions. In future work, this could be compensated for on drone platforms by measuring the wind strength.

5.2 SWaP-constrained wireless link

5.2.1 Benchtop experiments. To evaluate the communication link, we first perform a benchtop test of packet error rate (PER) with a controlled attenuation increase to empirically determine our receiver sensitivity. We first set the transmit power level of NRF52832 Development kit in software and measure the output on a spectrum analyzer to determine a starting reference level. Next, we place the transmitter in a shielded metal box to prevent weak radiated signals from coupling directly to the receiver, and connect its antenna output through a series of RF attenuators. We record the raw RF signal on a software radio (USRP E310, National Instruments) and decode the data. We divide the number of packets received by the number transmitted to determine PER. Fig. 10(a) shows that our receiver can operate down to -112 dBm with a 33 kbps data rate. The dotted line denotes the sensitivity of a SX1276 LoRa chipset at 37.5 kbps, for comparison. Note however that LoRa operates at 900 MHz which typically has a better signal penetration.

5.2.2 Effect of duty-cycling. Fig. 10(b) shows the projected battery life of our sensor platform at different transmitted packets per hour. We acquired this data by measuring the current consumed during a packet transmission (3.358mA, measured with Fluke 287 Digital Multimeter) and while sleeping (35nA, measured with a Keithley Source Measurement Unit). Doubling the rate from 10 to 20 packets per hour does not halve the battery life. This is because at that rate the sleep current consumed by our hardware timer (35nA) is the dominating factor. As we increase the number of packets per hour, the current consumed by each packet transmission (3.358mA) begins dominating, further reducing battery life.

5.2.3 Long-range experiments. Our transmitter consists of our small form factor PCB prototype connected to the lightweight antenna shown in Fig. 2. We program the device to transmit a packet at 4 dBm every 20 ms. Our receiver consists of the same USRP E310 described above connected to a 6 dBi patch antenna. It is connected to a 30 dB amplifier while transmitting. We move our transmitter and receiver to various locations shown in Fig. 11(a). The plot shows that at 33 kbps our system can operate at ranges up to 0.8-1 km, significantly increasing the range over the 1 Mbps Bluetooth link.




Figure 9: Mobility Experiments. (a) Packet error rate CDF for different flight speeds. (b) Spectrum of a ‘0’ bit transmitted while flying. (c) Spectrum of a ‘1’ bit transmitted while flying.




Figure 10: Benchtop communication and power experiments.




Figure 11: In-air communication experiments.

The graph also shows that across the whole range the downlink with the 30 dBm transmissions were decoded at the BLE chipset. The low PER is because of the lack of other Bluetooth radios in the field. This demonstrates that our SWaP-constrained link can enable bi-directional communication for insect-scale sensors that can be deployed on drones at long distances from the basestation.

5.2.4 Mobility and packet error rate. We also evaluate our wireless link in flight to empirically determine whether factors such as Doppler shifts cause performance degradation. To evaluate this, we attach our payload with the form-factor Bluetooth transmitter and our lightweight antenna to a DJI Phantom 3 SE drone. We opt to use this larger drone not because of payload limitations but rather because the Phantom 3 can record its precise GPS location and flight speed unlike the smaller drones used before.

We set our transmitter to broadcast packets at a fixed rate and attached it to the bottom right corner of the drone. We set up a 6 dBi patch antenna angled 15° above the horizon and connected to a USRP N210 which acts as the receiver. We controlled the drone to lift off approximately 5 m in front of the receiver antenna and rise to an altitude of roughly 8-10 m. From this point we flew the drone in a straight line for a distance of approximately 350 m, and then back along the same path toward the receiver. Fig. 9 shows the packet error rates when the drone was stationary, as well as when moving at 3 m/s and 7 m/s which exceeds the speed of our target insect deployments. The error rate across these conditions remained below 3%. We note the presence of strong wind during some of these experiments which moved the drone off course and potentially contributed to higher variance.

To further confirm that the data was not affected by Doppler shifts we also examine the frequency of the received signal. Fig. 9 shows the representative spectrum of a ‘0’ and ‘1’ bit from the recorded data. The plots show that all three speeds have no noticeable difference in the frequency of each of the bits. Further, upon closer inspection of the Bluetooth specification we observe that it is designed to tolerate frequency drift within 50 kHz to account for the accuracy of commercially available crystal oscillators used as frequency references for Bluetooth chips [55]. For comparison, Doppler shift at 7 m/s and 2.4 GHz is less than 100 Hz. As a result, despite increasing the symbol length, since the frequency separation between the 0 and 1 bits is orders of magnitude more than the Doppler shift, it does not noticeably increase the PER.

5.3 Supporting multiple sensors

To evaluate our receiver’s ability to differentiate between multiple wireless sensors, we set up an experiment with three of our wireless sensors and a single USRP receiver. All of the transmitters on the sensors are programmed with identical firmware, and set to send packets at random intervals between 20-40ms on power-up. The receiver stores three virtual preambles corresponding to the three unique preambles from the transmitters. Even though our wireless sensors are all running the same firmware, they each have unique preambles due to their randomly assigned Bluetooth addresses.

Next, we turn ON all three sensors so that they begin to send wireless packets. Each of these sensors use Bluetooth’s inherent MAC protocol to randomize when they transmit. We record the transmissions from each of these sensors at approximately three different SNR values. Upon receiving a packet, the receiver correlates the received packet’s preamble against all three stored virtual preambles. The virtual preamble that generates the maximum auto-correlation is used to identify the specific sensor (see §4.2).

Fig. 12 shows the difference in correlation with the three virtual preambles at three different SNRs (12dB, 6dB, and 2dB). We plot the normalized correlation values to account to differences as a function of the signal strength. In the i^{th} sub-figure, we correlate with the virtual preamble corresponding to the i^{th} and plot the results for each of the three sensors at each of the three SNR values at which they were received. We can visually observe that the large difference in correlation between the expected sensor and the others. At 12dB, the difference in normalized correlation is approximately 0.7. These correlations begin to converge toward the center as the signal quality degrades. However, even at 2dB SNR, the difference in correlation is above 0.4. This confirms using the built-in MAC addresses of each Bluetooth chip to identify each sensor node.




Figure 12: Normalized cross correlation values. (a) CDF of correlation with the first sensor (tx1). (b) CDF of correlation with the second sensor (tx2). (c) CDF of correlation with the third sensor (tx3).




Figure 13: Angle of arrival error versus range.

5.4 AoA experiments

We use USRP N200s and X300s to create a four antenna basestation. All the USRPs on the basestation are synced by an external clock. To account for the random phase offsets introduced when the USRPs boot up, we have an additional USRP that transmits a pure sine wave signal so that all four USRP receivers can calibrate via cables. These are controlled by four external RF switches. We use our Bluetooth chip with the small form factor antenna as the transmitter. We place the basestation at a fixed location and move the wireless sensor to increasing distances as shown in Fig.13(a). To measure the ground truth angle we first measure the compass direction of the antenna array using a smartphone. We then log the GPS coordinates of the AP as well as the sensor. To minimize the error introduced by GPS, we use satellite imagery from Google Maps to choose locations with visible features or landmarks. We compute the angle and compare it to the ground truth obtained from the GPS coordinates and plot the results in Fig. 13(b). We show the angular errors for both three and four antennas. The plot shows that the average angular error is smaller with four antennas and is around 1.3 degrees when averaged across all locations. AoA when combined with multiple base stations could provide a rough location estimate.

6 USING SMALL DRONES & INSECTS

Drones. We begin by describing the most useful scenario for use with drones. We attach sensors to a drone which flies to the desired locations. A command is sent to the drone using its built-in wireless communication to release a wireless sensor at each location. Thus, RF localization is not required and the wireless communication on the sensor in §4.2 does not need to be used until after release.

When attaching our sensors to a drone we seek to minimize the surface area they occupy to allow a drone to carry the maximum number of sensors. Hanging vertically, each sensor only occupies an 4×4.2 mm area. We observe that the smallest commercial drones measuring 28 mm in width can reliably carry less than 500 mg limiting their capacity to 2 sensors; however even slightly larger drones measuring 92 mm can lift 15 g which would be sufficient for more than 80 sensors. Fig. 1 shows a single sensor suspended from

the bottom of the smallest commercially available drone [4]. A piece of carbon fiber is glued to the bottom of the drone which slides into the release mechanism. In the video link from the figure, the drone is placed on a raised Styrofoam takeoff platform that prevents the sensor from touching the ground. After takeoff, a Bluetooth packet triggers the sensor release. The sensor falls to the ground and received packets are shown on a computer screen.

Insects. While in the future, our technology can be combined with moth flight control [60], here, we describe two scenarios where the insect is freely flying. In the first scenario, the insect is tagged with the sensor(s). The insect flies freely and the sensor is programmed to wake up and drop after a random duration. After dropping, the sensor sends out localization broadcasts and then periodically transmits temperature or humidity sensor data. Considering social insects such as hornets are known to fly directly to their nests after capture [30] while others will wander randomly when introduced to a new environment [63], we can select for a desired behavior by choosing different insect species. Deploying a large number of sensors with multiple insects and random drop times could be used to achieve good coverage, albeit with some redundancy. Here, communication is only used after the release.

In the second scenario, the sensor is dropped based on a command from the base station. Here, the insect is first instrumented with the sensors and released to fly freely. The sensor wakes up either periodically (say every 10 minutes) or when the insect comes to a stop and transmits a message. The basestation uses these transmissions to localize the insect. If the insect is in a new location, the base station immediately sends an ACK (see §4.2.5) to release the sensor. This step can be done either when the insect is stationary or in flight. The sensor then receives the packet and checks the data to see if it is the intended recipient. If so, it triggers the dropping mechanism to drop the sensor. Given the results in §5.2.4, the sensor can transmit and receive wirelessly while in motion. The Doppler shift from motion is orders of magnitude smaller than the 50 kHz frequency drift supported by the Bluetooth radio spec [55]. We can also use a low-power accelerometer to identify when the insect is stationary before transmitting and receiving. This is useful if we would want to release sensors near insect nests; this could help eradicate invasive species like hornets. Finally, since AoA based localization accuracies degrade with distance, the precision with which insects can be used to deploy in specific locations depends on the localization error. This is unlike the drone use case where we could deploy our sensors with precision. An alternative approach is to use recent advancements in low power insect-scale cameras [23] to more precisely target a drop location. We leave integrating cameras and airdropping sensors to future work.

To attach our sensors to live insects, we choose *Manduca sexta*, a species of moth that has been used extensively as a model animal for the study of flight dynamics. These insects have been shown to carry a payload sufficient to support our sensor [60]. We first place the moth in a freezer for approximately 5 min to apply cold anesthesia. Reducing the insect's body temperature slows its metabolism and briefly immobilizes it allowing for handling. We then apply a drop of CA glue to the lower thorax to secure the carbon fiber rod attached to the release mechanism as seen in Fig. 3. We choose this position as most prior work has selected this site for payload attachment and shown it has minimal impact on flight and that the animal cannot remove it. Additionally, it rests against the insect's abdomen which curves down below this point allowing the sensor to fall. Once the release mechanism is open, the periodic wing strokes cause the whole insect to vibrate at 25 Hz [43] which guarantees release.

7 RELATED WORK

Device deployment from large drones. Prior work focuses on deploying sensors and other objects (e.g., packages, fishing bait [6]) from large industrial drones. To the best of our knowledge, we are the first to take the approach of creating insect-scale sensors that can fall and not be affected by their impact due to their small size and mass. Further, in addition to the long range communication, we are the first to design a small size, weight and power release mechanism that can be deployed with small drones as well as insects. Here, we describe four prior works that we build on.

Commercial cargo carrying drones like the Amazon drone package delivery system [2] use grippers to deliver packages. These are designed to be placed on the ground carefully and not released in air. Further they require large industrial scale drones, consume much higher power and have orders of magnitude larger sizes than our target applications. Commercial electropermanent magnets [7] are large and are targeted for much larger payloads.

Samara [48, 49] is a large seed-shaped sensor that can be dropped without damage because of its aerodynamic properties. The sensor weighs 25 grams, is 12 cm wide and requires a large 2 kg industrial-scale drone to deploy the sensor. Further, the release mechanism on the drone is large in size and itself requires an industrial scale drone. Further it uses a motor, which consumes orders of magnitude more energy than is achieved by our release mechanism design.

In 2019, NASA demonstrated a release mechanism for gliding aircraft equipped with sensors [5]. They use the CICADA gliders from US naval research, which are small controllable aircraft that glide in the air. These gliders are 15-20 cm in width and weigh around 65 grams [3]. The release mechanism is a few meters in length and is designed to be dropped from army helicopters and industrial size drones. Along similar lines, recent work uses a helicarrier that releases a large industrial drone in air, which then is controlled to land on the ground and dig holes [47]. Our work is focused on orders of magnitude smaller size, weight and power requirements and is complementary to these works.

Finally, [31, 45] designs flexible sensors that can be 'worn' by plants by mimicking the shape of butterflies and paper wings. Dropping these sensors from drone in the air was not demonstrated. Further, the sensor and assembly weigh around 5 grams and require a release mechanism that uses a servo and a rack and pinion gear

mechanism that is longer than 20 cm. Finally, the sensors have a 5 μ A sleep current; in comparison, we achieve 35 nA sleep current.

Data collection using drones and insects. Drones have been used as a platform to enable mobility and collect data across farms. [57] uses sensors on a drone to collect water temperature information as the drone is flying over a lake. Using drones to collect data however is limiting since it does not allow for temporal sensor information in each location as the information gathering duration is limited by the 5-30 minute operation time of the drone.

Living IoT [24] places a wireless sensor on a living bee that collects data as it flies over the target area. Our work differs from Living IoT in three ways: First, our focus is on releasing wireless sensor from small drones as well as insects. This requires designing a SWaP release mechanism that can fit on our target platforms. Second, Living IoT uses backscatter and has an uplink range of less than a meter; hence it only uploads the collected data once the bee is back to the hive co-located with the basestation. In contrast, we design a SWaP programmable communication system that has a bi-directional communication range of around a kilometer in open outdoor settings. Third, we design a low-power timer based sensor system that can achieve sleep currents as low as 35 nA. This allows the sensor to duty cycle and operate for more than 1.3-2.5 years.

8 DISCUSSION AND CONCLUSION

We present the first system that can airdrop wireless sensors from small drones and live insects. We take inspiration from insects such as ants and design insect-scale wireless sensors that achieve a low terminal velocity and can survive impact. While we believe that this work is an important milestone in enabling automated wireless sensor deployment, we outline limitations and future directions.

- *Deploying multiple sensors.* While we discuss the potential of deploying multiple sensors from a single aerial platform, our implementation is limited to a single sensor. Releasing multiple sensors using a single small drone or insect is an important next step.
- *Controlled sensor deployment using insects.* Prior designs show control of insects such as beetles, dragonflies and Locusts [13, 39, 51]. In contrast to this work, we are the first to use insects as a platform to deploy sensors. A future direction is to control insects to navigate to specific locations and then release sensors.
- *Fully automated deployment.* We manually control the drone using a remote operator. Recent robotics work use deep learning and better sensors to automate navigation. Integrating sensor deployment with better navigation can automate our deployment process.
- *Picking up dead sensors.* Widespread sensor network deployments can lead to electronics waste once the batteries run through their capacity. While this paper focuses on deploying these sensors, we also need techniques to recover our sensors using drones.

Acknowledgments. We thank Jose Jaime, Sawyer Fuller as well as the anonymous MobiCom reviewers for their constructive feedback. The authors are funded in part by NSF awards CNS-1812554, CNS-1452494, CNS-1823148, and Google Faculty Awards.

REFERENCES

- [1] 2013. An ant dropped off the Empire State Building. (2013). <https://van.physics.illinois.edu/qa/listing.php?id=2099&t=an-ant-dropped-off->

- the-empire-state-building.
- [2] 2016. Amazon Prime Air. (2016). <https://www.amazon.com/Amazon-Prime-Air/>.
 - [3] 2017. Naval Research Lab Tests Swarm of Stackable CICADA Microdrones. (2017). <https://spectrum.ieee.org/automaton/robotics/drones/naval-research-lab-tests-swarm-of-stackable-cicada-microdrones>.
 - [4] 2018. Mini Drone FY804 4CH 2.4G 6Axis 360 Degree Roll Mini Helicopter LED Plane Model Toys RC Aircraft. (2018). <https://www.aliexpress.com/i/32983575363.html>.
 - [5] 2019. Drones swarm over Beaver Dam. (2019). <https://www.nasa.gov/feature/langley/drones-swarm-over-beaver-dam>.
 - [6] 2019. Professional Release and Drop Device for DJI Phantom 4 All Models, for Drone Fishing, Bait Release, Payload Delivery, Search & Rescue, Fun Activities. - Free Drop Parachute Included. (2019). <https://www.amazon.com/Professional-Release-Fishing-Delivery-Activities/dp/B07FC467WZ>.
 - [7] 2020. "Nica Drone". (2020). <https://nicadrone.com/products/epm-v3>.
 - [8] M. H. Almarshadi and S. M. Ismail. 2011. 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*.
 - [9] Roshan Ayyalasomayajula, Deepak Vasisht, and Dinesh Bharadia. 2018. BLoc: CSI-Based Accurate Localization for BLE Tags. In *Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT '18)*. Association for Computing Machinery, 126–138.
 - [10] Palak Bhushan. 2019. *Untethered Microrobots of the Rolling, Jumping & Flying kinds*. Ph.D. Dissertation. UC Berkeley.
 - [11] P. Bhushan and C. Tomlin. 2020. An Insect-Scale Self-Sufficient Rolling Micro-robot. *IEEE Robotics and Automation Letters* 5, 1 (2020), 167–172.
 - [12] 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. 2016. Energy-Autonomous Wireless Communication for Millimeter-Scale Internet-of-Things Sensor Nodes. *IEEE Journal on Selected Areas in Communications* 34, 12 (2016), 3962–3977.
 - [13] Evan ckerman. 2017. DragonflyEye Project Wants to Turn Insects Into Cyborg Drones. (2017). <https://spectrum.ieee.org/automaton/robotics/industrial-robots/draper-dragonflyeye-project>.
 - [14] Douglas C Giancoli. 2007. *Physics for Scientists Engineers* (4 ed.). Pearson.
 - [15] Jon Gjengset, Jie Xiong, Graeme McPhillips, and Kyle Jamieson. [n. d.]. Phaser: Enabling Phased Array Signal Processing on Commodity WiFi Access Points. In *MobiCom '14*.
 - [16] Shyamnath Gollakota, Nabeel Ahmed, Nikolai Zeldovich, and Dina Katabi. 2011. Secure In-Band Wireless Pairing. In *Proceedings of the 20th USENIX Conference on Security (SEC'11)*. USENIX Association, USA, 16.
 - [17] Moritz Graule, P. Chirarattananon, Sawyer Fuller, N. Jafferis, K. Ma, Matthew Spenko, Roy Kornbluh, and R. Wood. 2016. Perching and takeoff of a robotic insect on overhangs using switchable electrostatic adhesion. *Science* 352 (05 2016), 978–982. <https://doi.org/10.1126/science.aaf1092>
 - [18] Malcolm Hawkes and Arye Nehorai. 1998. Acoustic vector-sensor beamforming and Capon direction estimation. *IEEE transactions on signal processing* 46, 9 (1998), 2291–2304.
 - [19] Mehrdad Hesar, Ali Najafi, and Shyamnath Gollakota. 2019. NetScatter: Enabling Large-Scale Backscatter Networks. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. USENIX Association, Boston, MA. <https://www.usenix.org/conference/nsdi19/presentation/hesar>
 - [20] Mehrdad Hesar, Ali Najafi, Vikram Iyer, and Shyamnath Gollakota. 2020. TinySDR: Low-Power SDR Platform for Over-the-Air Programmable IoT Testbeds. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*.
 - [21] Marina Indri, Luca Lachello, Ivan Lazzero, Fiorella Sibona, and Stefano Trapani. 2019. Smart Sensors Applications for a New Paradigm of a Production Line. *Sensors* 19, 3 (2019).
 - [22] Texas Instrument. 2020. CC2530. (2020). Retrieved March 25, 2020 from <http://www.ti.com/product/CC2530>
 - [23] Vikram Iyer, Ali Najafi, Johannes James, Sawyer Fuller, and Shyamnath Gollakota. 2020. Wireless steerable vision for live insects and insect-scale robots. *Science Robotics* 5, 44 (2020). <https://doi.org/10.1126/scirobotics.abb0839>
 - [24] Vikram Iyer, Rajalakshmi Nandakumar, Anran Wang, Sawyer B. Fuller, and Shyamnath Gollakota. 2019. Living IoT: A Flying Wireless Platform on Live Insects. In *The 25th Annual International Conference on Mobile Computing and Networking (MobiCom '19)*. Association for Computing Machinery, New York, NY, USA.
 - [25] Vikram Iyer, Vamsi Talla, Bryce Kellogg, Shyamnath Gollakota, and Joshua Smith. 2016. Inter-Technology Backscatter: Towards Internet Connectivity for Implanted Devices. In *SIGCOMM '16*.
 - [26] Noah Jafferis, E. Helbling, Michael Karpelson, and Robert Wood. 2019. Untethered flight of an insect-sized flapping-wing microscale aerial vehicle. *Nature* 570 (06 2019), 491–495. <https://doi.org/10.1038/s41586-019-1322-0>
 - [27] J. James, V. Iyer, Y. Chukewad, S. Gollakota, and S. B. Fuller. 2018. Liftoff of a 190 mg Laser-Powered Aerial Vehicle: The Lightest Wireless Robot to Fly. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. 3587–3594.
 - [28] Bryce Kellogg, Aaron Parks, Shyamnath Gollakota, Joshua R. Smith, and David Wetherall. 2014. Wi-Fi Backscatter: Internet Connectivity for RF-Powered Devices. In *Proceedings of the 2014 ACM Conference on SIGCOMM (SIGCOMM '14)*. Association for Computing Machinery, New York, NY, USA, 607–618. <https://doi.org/10.1145/2619239.2626319>
 - [29] Bryce Kellogg, Vamsi Talla, Shyamnath Gollakota, and Joshua R. Smith. 2016. Passive Wi-Fi: Bringing Low Power to Wi-Fi Transmissions. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*. USENIX Association, Santa Clara, CA, 151–164.
 - [30] Peter J Kennedy, Scott M Ford, Juliette Poidatz, Denis Thiéry, and Juliet L Osborne. 2018. Searching for nests of the invasive Asian hornet (*Vespa velutina*) using radio-telemetry. *Communications biology* 1, 1 (2018), 1–8.
 - [31] S. Khan and M. M. Hussain. 2019. IoT enabled Plant Sensing Systems for Small and Large Scale Automated Horticultural Monitoring. In *2019 IEEE 5th World Forum on Internet of Things (WF-IoT)*. 303–308.
 - [32] N. Kim. 2020. A Digital-Intensive Extended-Range Dual-Mode BLE5.0 and IEEE802.15.4 Transceiver SoC. *IEEE Transactions on Microwave Theory and Techniques* 68, 6 (2020), 2020–2029.
 - [33] Ara Nerses Knaian. 2010. *Electropermanent magnetic connectors and actuators: devices and their application in programmable matter*. Ph.D. Dissertation. Massachusetts Institute of Technology.
 - [34] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. 2015. SpotFi: Decimeter Level Localization Using WiFi. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (SIGCOMM '15)*. Association for Computing Machinery, 269–282.
 - [35] F. Kuo, S. Binsfeld Ferreira, H. R. Chen, L. Cho, C. Jou, F. Hsueh, I. Madadi, M. Tohidian, M. Shahmohammadi, M. Bahaie, and R. B. Staszewski. 2017. A Bluetooth Low-Energy Transceiver With 3.7-mW All-Digital Transmitter, 2.75-mW High-IF Discrete-Time Receiver, and TX/RX Switchable On-Chip Matching Network. *IEEE Journal of Solid-State Circuits* 52, 4 (2017), 1144–1162.
 - [36] David Lentink, William B Dickson, Johan L Van Leeuwen, and Michael H Dickinson. 2009. Leading-edge vortices elevate lift of autorotating plant seeds. *Science* 324, 5933 (2009), 1438–1440.
 - [37] Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R. Smith. 2013. Ambient Backscatter: Wireless Communication out of Thin Air. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM (SIGCOMM '13)*. Association for Computing Machinery, New York, NY, USA, 39–50. <https://doi.org/10.1145/2486001.2486015>
 - [38] Jess Lowenberg-DeBoer. 2015. *The Precision Agriculture Revolution: Making the modern farmer*. (2015).
 - [39] K Mann, TL Massey, S Guha, JP van Kleef, and MM Maharbiz. [n. d.]. A wearable wireless platform for visually stimulating small flying insects. In *Engineering in Medicine and Biology Society (IEEE), 2014*.
 - [40] Microchip. 2020. Small Lora Chip: ATSAMR34J16BT-I/7JX. (2020). Retrieved March 25, 2020 from <https://www.digikey.com/product-detail/en/microchip-technology/ATSAMR34J16BT-I-7JX/ATSAMR34J16BT-I-7JXCT-ND/9739172>
 - [41] MicronWings. 2020. smallest commercial motor. (2020). Retrieved March 25, 2020 from <https://micronwings.com/Products/MotorBrushedNanoSize3mmSurfaceMount/index.shtml>
 - [42] Antonio Molina-Pico, David Cuesta-Frau, Alvaro Araujo, Javier Alejandre, and Alba Rozas. 2016. Forest Monitoring and Wildland Early Fire Detection by a Hierarchical Wireless Sensor Network. *Journal of Sensors* 2016 (02 2016), 1–8. <https://doi.org/10.1155/2016/8325845>
 - [43] KC Moses, SC Michaels, M Willis, and RD Quinn. 2017. Artificial Manduca sexta forewings for flapping-wing micro aerial vehicles: how wing structure affects performance. *Bioinspiration & biomimetics* 12, 5 (2017), 055003.
 - [44] Rajalakshmi Nandakumar, Vikram Iyer, and Shyamnath Gollakota. 2018. 3D Localization for Sub-Centimeter Sized Devices. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems (SenSys '18)*. Association for Computing Machinery, New York, NY, USA, 108–119. <https://doi.org/10.1145/3274783.3274851>
 - [45] Joanna Nassar, Sherjeel Khan, Diego Villalva, Maha Nour, and Muhammad Hussain. 2018. Compliant plant wearables for localized microclimate and plant growth monitoring. *npj Flexible Electronics* 2 (12 2018). <https://doi.org/10.1038/s41528-018-0039-8>
 - [46] Aaron N. Parks, Angli Liu, Shyamnath Gollakota, and Joshua R. Smith. 2014. Turbocharging Ambient Backscatter Communication. In *Proceedings of the 2014 ACM Conference on SIGCOMM (SIGCOMM '14)*. Association for Computing Machinery, New York, NY, USA, 619–630. <https://doi.org/10.1145/2619239.2626312>
 - [47] Adam Plowcha, Yue Sun, Carrick Detweiler, and Justin Bradley. 2018. Predicting Digging Success for Unmanned Aircraft System Sensor Employment. In *International Symposium on Experimental Robotics*. Springer, 153–164.
 - [48] Paul Pounds, Timothy Potie, Farid Kendoul, Surya Singh, Raja Jurdak, and Jonathan Roberts. 2014. Automatic Distribution of Disposable Self-Deploying Sensor Modules, Vol. 109. https://doi.org/10.1007/978-3-319-23778-7_35
 - [49] P. Pounds and S. Singh. 2015. Samara: Biologically Inspired Self-Deploying Sensor Networks. *IEEE Potentials* 34, 2 (March 2015), 10–14. <https://doi.org/10.1109/MPOT.2014.2359034>

- [50] Ludovic Ravel, Emmanuel Malet, and Philippe Batoux. 2017. Multi-parameter monitoring of the construction and evolution of a snowbridge over a crevasse on an Alpine glacier. *Geophysical Research Abstracts* (2017).
- [51] Hiroataka Sato, Christopher W Berry, Yoav Peeri, Emen Baghoomian, Brendan E Casey, Gabriel Lavella, John M VandenBrooks, Jon Harrison, and Michel M Maharbiz. 2009. Remote radio control of insect flight. *Frontiers in integrative neuroscience* 3 (2009), 24.
- [52] Ralph Schmidt. 1986. Multiple emitter location and signal parameter estimation. *IEEE transactions on antennas and propagation* 34, 3 (1986), 276–280.
- [53] SEMTECH. 2020. sx1276. (2020). Retrieved March 25, 2020 from <https://www.semtech.com/products/wireless-rf/lora-transceivers/sx1276>
- [54] Tie-Jun Shan, Mati Wax, and Thomas Kailath. 1985. On spatial smoothing for direction-of-arrival estimation of coherent signals. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 33, 4 (1985), 806–811.
- [55] Bluetooth SIG. 2014. Bluetooth Specification Version 4.2. (2014). <https://www.bluetooth.com/specifications/bluetooth-core-specification/bluetooth5>
- [56] Bluetooth SIG. 2020. Bluetooth 5 advertising. (2020). Retrieved March 25, 2020 from <https://www.bluetooth.com/blog/exploring-bluetooth5-whats-new-in-advertising/>
- [57] L. Sørensen, L. Jacobsen, and J. Hansen. 2017. Low Cost and Flexible UAV Deployment of Sensors. *Sensors* (2017).
- [58] Michael M Thackeray, Christopher Wolverton, and Eric D Isaacs. 2012. Electrical energy storage for transportation—approaching the limits of, and going beyond, lithium-ion batteries. *Energy & Environmental Science* 5, 7 (2012), 7854–7863.
- [59] Josué Toledo-Castro, Pino Caballero-Gil, Nayra Rodríguez Pérez, I. Santos-González, Candelaria Hernández-Goya, and Ricardo Aguiasca. 2018. Forest Fire Prevention, Detection, and Fighting Based on Fuzzy Logic and Wireless Sensor Networks. *Complexity* 2018 (12 2018), 1–17. <https://doi.org/10.1155/2018/1639715>
- [60] W. Tsang, Alice Stone, David Otten, Zane Aldworth, Tom Daniel, John Hildebrand, Richard Levine, and Joel Voldman. 2011. Insect-machine interface: A carbon nanotube-enhanced flexible neural probe. *Journal of neuroscience methods* 204 (11 2011), 355–65. <https://doi.org/10.1016/j.jneumeth.2011.11.026>
- [61] G. Werner-Allen, J. Johnson, M. Ruiz, J. Lees, and M. Welsh. 2005. Monitoring volcanic eruptions with a wireless sensor network. In *Proceedings of the Second European Workshop on Wireless Sensor Networks, 2005*. 108–120.
- [62] Jie Xiong, Karthikeyan Sundaresan, and Kyle Jamieson. 2015. ToneTrack: Leveraging Frequency-Agile Radios for Time-Based Indoor Wireless Localization. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (MobiCom '15)*. Association for Computing Machinery, New York, NY, USA, 537–549.
- [63] Michael Zimmerman. 1979. Optimal foraging: a case for random movement. *Oecologia* 43, 3 (1979), 261–267.