Sensor Networks for Experience and Ecology

Brian D. Mayton

B.S., University of Washington, Seattle (2008)
S.M., Massachusetts Institute of Technology (2013)

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Media Arts and Sciences

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2020

Author ______________________________

Program in Media Arts and Sciences
August 17, 2020

Certified by ______________________________

Joseph A. Paradiso
Alexander W. Dreyfoos (1954) Professor
Program in Media Arts and Sciences
Thesis Supervisor

Accepted by ______________________________

Prof. Tod Machover
Academic Head
Program in Media Arts and Sciences

Sensor Networks for Experience and Ecology

Brian D. Mayton

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning, on August 17, 2020,
in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Media Arts and Sciences

Abstract

Wetlands are critically important ecosystems, providing numerous benefits to the global environment. As cranberry farms in southeastern Massachusetts come out of production, many are undergoing active restoration with the goal of returning them to functional wetlands. This is a developing practice, and we are still learning about techniques that lead to the most favorable outcomes. It is also a disruptive process that brings significant and very visible changes to the landscape. But the process of wetland restoration and the restored wetlands themselves present fantastic opportunities for learning and enjoyment.

In this dissertation, I present a custom sensor network installed at the Tidmarsh Wildlife Sanctuary, a former cranberry farm in Plymouth, Massachusetts that underwent active restoration in 2016. This network combines hundreds of custom low-power environmental sensor nodes with high-bandwidth continuous audio and video streams and the required supporting communications and data storage infrastructure.

As a permanent fixture of the site, the network was designed to serve multiple functions, making Tidmarsh a testbed for many ideas. Long-term continuous monitoring, both before and after the restoration, allows us to observe changes that take years to decades and answer questions about restoration techniques and outcomes. Broad sensing capabilities allow us to make observations and gather data about questions we may not have thought to ask. Real-time data streaming and access protocols allow us to build novel ways of exploring and experiencing the site, both while physically present and remotely. Rich media streams enhance these experiences and allow us to assess complex factors such as biodiversity.

I describe the design, implementation, and deployment of two generations of custom wireless sensor hardware and the supporting network infrastructure; a multi-channel audio streaming installation; and a setup for video streaming and timelapse recording. I demonstrate how the network is used for scientific research through an experiment to determine the impact of microtopography (a restoration technique) on soil hydrology. Finally, I enumerate the many projects that have made use of the network to learn from the data and connect people to restored wetlands through novel experiences and creative expressions.

Thesis Supervisor: Joseph A. Paradiso
Title: Alexander W. Dreyfoos (1954) Professor of Media Arts and Sciences

Sensor Networks for Experience and Ecology

Brian D. Mayton

This dissertation has been reviewed and approved by the following committee members:

Thesis Supervisor ______________________

Joseph A. Paradiso
Alexander W. Dreyfoos (1954) Professor
of Media Arts and Sciences
Massachusetts Institute of Technology

Member, Thesis Committee ______________________

Kate Ballantine
Marjorie Fisher Associate Professor
of Environmental Studies
Mount Holyoke College

Member, Thesis Committee ______________________

Alex Hackman
Ecological Restoration Specialist
Division of Ecological Restoration
Massachusetts Department of Fish and Game

Acknowledgments

There are many people without whom this work would not have happened.

Thank you to Joe Paradiso for advising, supporting, and respecting me for the past decade and providing an environment where this kind of work is possible and encouraged.

To Gershon Dublon, who has been one of my closest friends and collaborators from almost the moment I joined the Media Lab. I don't think either of us had any idea when we went down to Plymouth and did a few RF tests at a place called Tidmarsh just how deep in the muck we would both literally and figuratively end up together there in the years that followed. It's pretty wild, 8 years later, to finally be writing this document. I never would have made it halfway here without you.

To Spencer Russell, who joined the Tidmarsh crew and helped so much to install and maintain the sensor network and found fantastic new ways to use it. Thanks for your support and friendship and for always being willing to have deep conversations about the obscure details of audio compression formats and signal processing.

To Ben Bloomberg. We've been co-conspirators in many different things both inside and outside of the Lab. Thanks for all of the distractions and side projects when I needed them and for being an inspiration in doing big things. And for being an amazing friend and keeping me from falling apart in the final stages of writing this document and preparing for my defense.

To Glorianna Davenport, who introduced me to Tidmarsh and guided the project to become something so much more than I ever thought it was going to be.

To Alex Hackman, whose passion for ecological restoration led me into an entirely new intellectual direction, and whose excitement for this work and its potential as a new tool in restoration has been a constant source of inspiration and guidance.

To Kate Ballantine, who helped me on my journey into ecology and shaped the design of the sensor network to be able to measure the changes set in motion by restoration.

To Christine Hatch, who taught me most of what I know about hydrology, helped me make sense of all of my soil moisture data, and was always excited to try out anything I built.

To Clément Duhart, who took our microphone installations so much further than I imagined with his work on Tidzam, and who helped out with so many things, including assembling a good majority of the v2 sensor nodes.

To Mass Audubon, who has allowed and encouraged this work to continue after taking ownership of the Tidmarsh Wildlife Sanctuary.

To everyone in the Responsive Environments Group, my family at the Lab. I am honored to work with such an interesting, talented, creative, and nice group of people who are exploring frontiers in many different fields. And to Opera of the Future for being my second unofficial family at the Lab.

To Josh Smith, who was my mentor during my time at Intel Labs Seattle, taught me how to be a researcher, and led me down the road to the Media Lab. And to Louis Legrand, with whom I built robots and traveled the world.

To Gaetano Borriello, who introduced me to the world of embedded systems and low-power wireless senor networks, and encouraged and supported me in pursuing graduate school.

To my parents, Nora and Gary Mayton, who encouraged my curiosity and interests in every way possible and continue to support me in everything that I do.

To everyone who has ever helped build or install a Tidmarsh sensor node or find exciting new ways to use the data from the network. And thank you to many others whom I've most certainly forgotten to include on this list.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In this dissertation, I present a new and unique sensor network designed with geographic, temporal, and social factors in mind, to be used by scientists, managers, educators, and artists. The network aims to reliably and practically provide real-time and archived data combined with rich audio and video streams, and supports applications for archiving, viewing, and experiencing long-term change, supporting research, stewardship, outreach, expression, and education in the context of restored wetlands.

1.1 Wetlands, Cranberry Farming, and Restoration

Wetlands are a critical part of our global ecology. These ecosystems, found in almost all parts of the world but occupying a small percentage of the earth's surface, contribute disproportionately to the functioning of the planet. Defined, as their name implies, by their wetness, they are the interfaces between land and water [49]. They provide important ecosystem services, such as filtering pollutants from water, sequestering carbon, and are one of the most diverse ecosystems found on the planet [78].

Cranberries have been cultivated in southeastern Massachusetts since the early nineteenth century, and have become a significant and iconic part of the economy (Figure 1.1). In Massachusetts, cranberry farms currently account for over

13,000 acres of land [11], with the majority located on peatlands. Many of these farms are constructed in riparian fens, with a large, continuously flowing stream channel, and are known as "flowthrough" farms [31].

Though these farms may legally be considered wetlands, intense management of the hydrology and soil degrade their natural functions. Water control structures, used to flood the cranberry beds during the winter and at harvest, alter the natural flow of water through the site. Waterways are straightened or disconnected and drainage ditches are added, reducing the amount of time that water spends on the land. Sand raises the ground surface above the water table, effectively making the ground much drier. Farming enforces a monoculture; pesticides and herbicides are used to deliberately exclude plant and insect species, and the wide-open, artificially flattened beds make an inhospitable habitat for wildlife.

While constructing cranberry farms on natural wetlands was historically advantageous, artificial irrigation has enabled farming on upland sites at lower cost than on former wetlands. New hybrid cultivars of cranberries produce higher yields and larger fruit than the selections of wild varieties present on most Massachusetts farms [12]. As farming on natural wetlands in Massachusetts becomes less economically viable, many of these properties are coming out of production, with this trend expected to continue over the coming decade [31].

As farms are removed from production, questions arise as to the future of the land. One compelling option is to restore these farms to functioning wetlands, with the hope of regaining the beneficial ecosystem services that they once provided. While restoration is a noble goal, it comes with several challenges.

First, wetland restoration is still a developing practice. The history of restoration projects is short compared to the time scale of the ecological processes and cycles involved. It can take years to decades for the effects of decisions made during the planning and construction phases of a project to fully play out. Much remains to be learned about which techniques in restoration practice can effectively and quickly steer restored sites towards desired outcomes.

Figure 1.1: Cranberry farming: an iconic Massachusetts industry.

Figure 1.2: Draining a reservoir: prior to, immediately after, and years following the opening of the dam.

Second, restoration brings abrupt and significant changes to the landscape, which are felt by people as much as they are by nature. These changes might even initially appear undesirable. A reservoir, once a beautiful lake, can appear to become a mud pit (Figure 1.2) when dams that have existed for a century are removed to restore the unobstructed passage of water. Often, years go by between the cessation of farming and the start of construction, during which new growth (often by invasive exotic species) takes place. The construction tears this all down, and for a moment in time heavy machinery reigns over a muddy and barren expanse. Once the heavy equipment leaves, nature seemingly explodes. The natural seed bank takes root and within months the site is awash with new growth. Gradually, wildlife returns as creatures find safe habitats. This significant transition presents a fantastic opportunity for outreach and teaching about restored wetlands and the environment, facilitating a deep connection to nature and ecological processes.

1.2 Tidmarsh and Sensor Networks

Tidmarsh is a large former cranberry farm near Plymouth, Massachusetts. It operated for around a century and at its peak produced 1% of Ocean Spray's annual cranberry harvest. In the 2000s, Glorianna Davenport and Evan Schulman, who had purchased the property in the 1980s, made the decision to stop farming and begin the process of transitioning Tidmarsh back to wetland. In 2010, they placed a conservation and restoration easement on the property under the USDA NRCS Wetland Reserve program. In 2011, the restoration became a priority project for the Massachusetts Department of Fish and Game's Division of Ecological Restoration (DER), in what would be the largest freshwater restoration to date in Massachusetts [43].

Davenport, who was a founding professor of the Media Lab and was familiar with prior work in the Responsive Environments group around sensor networks and sensor data browsers, approached us in 2012 and presented Tidmarsh as an opportunity to explore these tools to address some of the challenges of wetland restoration. She had recently founded a non-profit organization called Living Observatory, the mission of which is ”to tell the long-term story of ecological wetland restorations on retired cranberry farms, and to advance scientific knowledge and public understanding of wetland ecology.” Through the work described in this thesis, I demonstrate that a *sensor network* is well-suited to furthering these goals by designing, deploying, and evaluating such a network at Tidmarsh.

Sensor networks instrument a geographic area through distributed nodes that sense their environment and communicate data back to a central location. Since the mid-2000s, many take the form of *wireless* sensor networks (WSNs), which were enabled by the proliferation of low-power, low-cost microcontrollers and digital radio transceivers. These provide the distributed computation and communication resources needed to automate the collection of sensor readings and transmit them back to where they can be useful. Low cost and low-power operation enable large numbers of sensor nodes to be deployed for long-term monitoring. This makes them particularly well-suited to documenting the process of ecological restoration, capturing the initial state of the landscape and recording the changes that occur during the construction phase of the restoration and as nature continues to transform the site over the course of many years.

Unlike many sensor networks, whose primary objective is to collect data to test a specific research hypothesis or process goal for a single stakeholder, I designed the Tidmarsh sensor network to be a platform suitable for a broad range of applications and users. These users include scientists who are conducting research and furthering knowledge and the practice of wetland restoration; managers assessing restoration outcomes and engaging on long-term stewardship of the site; artists creating works and expressions that integrate sensor data to tell the story of restored wetlands; and educators using sensor data to teach students about ecological processes and change.

To support all of these users with an interest in long-term change, the sensor network provides access to both real-time and archived data. Real-time data can create a strong sense of presence, allowing visitors to access Tidmarsh remotely, or can enhance the experience of being on site by augmenting what a visitor perceives. Real-time data also helps drive management decisions that leverage current conditions. Archived data tells the story over time, and by analyzing and presenting this data we can learn from and teach about the longer-term processes that are occurring in a restored wetland.

The network augments the environmental sensor data from the low-power network with rich media streams: audio and video can capture an enormous amount of information, and hearing and seeing what goes on can be powerful and compelling. Listening to high-fidelity spatial audio can transport a remote listener to the middle of the bog. Timelapse video can show the return of wildlife and condense the gradual transformation so that it can be experienced in one sitting. Audio and video can be processed using machine learning techniques to derive high-level features and facilitate scientific research and assessments.

To accomplish these goals, the sensor network needed to be practical to install, configure, and maintain. And researchers who want to use the data to explore and test hypotheses and developers and artists who want to use the sensor network to build new experiences need to be able to access the data and metadata, both in real-time and from the years of archives.

This work comprises an in-depth investigation of an entire sensor network *system*, including the sensor hardware itself, network and database infrastructure for transmitting and storing the data, tools for documenting, monitoring, and maintaining the network, application programming interfaces (APIs) and tools for integrating data into various works, and the curation of the many projects that have successfully been built upon it.

1.3 Research Questions

This thesis aims to address three central questions:

1. To what extent can a sensor network serve the multiple roles of providing long-term data for ecological research to improve the practice of wetland restoration, as well as a tool for artistic expression and the creation of novel experiences to effectively engage visitors and the public?
2. What form does this multi-purpose sensor network take?
3. What are some of the ways that various stakeholders can beneficially use data from the sensor network, incorporating it into their works, explorations, and learning?

1.4 Contributions

The key contributions provided by this work are summarized as follows.

1.4.1 Sensor Network

The first contribution is the sensor network itself, including:

- The design, engineering, testing, deployment, and evaluation of two generations of custom sensor node hardware and base station hardware, firmware, and protocols.
- The design and implementation of network infrastructure to connect and power sensor sites across the Tidmarsh property.
- A dataset containing over 7 years of ongoing environmental and soil sensor data.
- Software for configuring, managing, monitoring, and documenting a large sensor network.
- Multiple audio capture installations supporting real-time live streaming and 24/7 recording, including a large 20-channel network of microphones, with a dataset now approaching 5 years.

- Multiple camera installations capable of live streaming, motion-triggered recording, and timelapse.

- Software tools for exploring the data.

The sensor network has been installed at Tidmarsh beginning in 2013 and has been iteratively developed to the present. It remains in place and will continue operating to support further research under the umbrella of Living Observatory.

1.4.2 Soil Moisture Study

The second contribution is a study of soil hydrology using the sensor network at Tidmarsh. I designed this study in collaboration with wetland restoration practitioners to explore how the Tidmarsh sensor network can be used to learn about the outcomes of restoration techniques and how a restored wetland continues to change in the years following construction. The outcome of the study is a set of observations about the hydrology of Tidmarsh, and insights into how these sensors might be usefully deployed at future restoration sites.

1.4.3 Experiences

The third contribution of this thesis is a catalog of the many creative works, experiences, and projects that have been enabled by the sensor network at Tidmarsh. These works demonstrate the interest and value of the sensor network in creating new expressions, and explore the various ways that people of different backgrounds, from ecologists to members of the public, can experience restored wetlands through sensors.

1.5 Navigating this Document

Chapter 2 describes related work and relevant background about the Tidmarsh restoration project. Chapters 3 and 4 cover the design, implementation, deployment, and evaluation of the sensor hardware and network. Chapter 5 describes the soil moisture experiments and data analysis. Chapter 6 catalogs the numerous applications, creative expressions, and public outreach projects that utilize data from the sensor network. Finally, Chapter 7 concludes the thesis and discusses potential future directions for this work.

As the installations and experiments described in this document are deeply tied to the changing physical landscape of Tidmarsh, and not all readers will be familiar with the geography of Tidmarsh, an atlas has been provided in Appendix A. These maps show Tidmarsh both before and after the restoration construction and include the names I have used to refer to various locations throughout the document.

Chapter 2

Background and Related Work

2.1 Sensor Networks

2.1.1 Early History and Remote Sensing

We have long paid attention to our environment as it affects every aspect of our lives. Even in our modern world, we are at the mercy of the environment, from the weather's simple effects on day-to-day activities and what clothes one might wear to important larger questions about the future habitability of our planet.

The development of early instruments, such as the thermometer, hygrometer, and barometer, enabled quantitative observations about the environment to be made. Communication via telegraph enabled weather reports to travel faster than the weather itself, and weather observation stations using communications about conditions nearby to predict future local conditions could be considered one of the first sensing networks [24]. Another early observation network was the U.S. Coast and Geodetic Survey's network of pendulum base stations for establishing precise differences in gravitational acceleration across the US [50].

Communicating about observations taken over a distributed area enabled scientists to draw conclusions that could not be made from any of the observations individually. Another motivation for measurements taken over a distance, however, is the ability to make observations where it is not possible or practical for people to go. The field of remote sensing includes sensors that can automatically transmit data, enabling them to send back observations from remote locations. The term "remote sensing" is most frequently applied specifically to imaging technologies [8], but can include other sensors as well. Imaging satellites have been instrumental in modern weather forecasting and learning about the earth and its atmosphere [60]. They have also have found applications in espionage and the corporate space. Probes and robots with batteries of scientific instruments and sensors have been sent to many planets and to the outer reaches of the solar system, and have produced many of the images and taught us most of what we know about the space beyond our own atmosphere.

2.1.2 Wireless Sensor Networks

The modern field of wireless sensor networks largely gained momentum in a 1997 research proposal by Pister et al. [55] to DARPA to develop Smart Dust, a network of tiny wireless sensor nodes [75]. This proposal laid the groundwork for modern sensor networks, and identified many of the challenges and desirable properties of a wireless sensor network. The work on Smart Dust also led to the development of the Mica platform [30], an open design for a low-power wireless sensor node, and TinyOS [39], both of which formed the basis for many groups' research in sensor networks throughout most of the 2000s.

Estrin's classic 2001 paper [22] summarizes some of the key challenges of the field of sensor networks. Wireless sensor nodes are almost always power-limited, having non-renewable energy sources or limited means to collect energy from the environment. Communication is almost always wireless and is often the most expensive operation in terms of power, which motivates distributed processing to reduce the amount of data that needs to be transmitted. Many applications, especially those deriving from the Smart Dust concept, benefit from networks that are ad-hoc, self-configuring, and dynamic. Nodes work together to optimize the collection, processing, and transmission of data, exploiting heterogeneous tiered architectures to maximize the usefulness of the network given limited power.

Many researchers have focused on the routing and networking aspect [3, 54] of wireless sensor networks. Routing encompasses the formation of networks that connect all of the nodes and do so in an energy-efficient manner that is also robust to failures and changes within the network. Data-centric routing approaches, such as Directed Diffusion [33] make sensor nodes aware of their applications and tasks, conserving energy by propagating information through the network and processing in a distributed manner that enables the network to answer useful questions without transmitting all of the observations back to a central node.

Another important theme is the dynamic management of power and resources. Often, sensor nodes are deployed at a density where events can be observed by multiple nodes at the same time [32]; this redundancy enables some of the nodes to be cycled off part of the time while still maintaining nearly complete coverage. The density of operating nodes (and parameters such as sampling and transmission rates) might be varied based on the detection of events [44, 32] or the applications needed at a particular time [36].

2.1.3 Sensor Networks in Practice

Sensor networks have been applied to solve a wide variety of problems. This section lists a few, and highlights characteristics that are particularly relevant to the work done at Tidmarsh.

One of the first successful in-situ wireless sensor networks for monitoring of a natural environment was the collaboration between Intel Research Berkeley and UC Berkeley on Great Duck Island [71, 70]. The research took place on a small island off the coast of Maine, and aimed to answer several questions about nesting petrels, an endangered bird species. The sensor network's ability to make in-situ measurements was critical to the research, as the petrels would change their behavior or leave their nests entirely if disturbed by researchers.

The deployment used a hierarchical architecture, with a collection of sensor nodes forming a "patch". Each patch has its own gateway, which collects the data from the patch and relays it to a central base station.

Another Intel Research project conducted studies with a sensor network installed in a vineyard in Oregon [7]. The goal of the sensor network was to identify conditions and microclimates that affected the quality of the grapes and could lead to diseases and crop loss. The work carefully considered the way that people interact with the data from the sensor network. The researchers produced a graphical interface for interacting with the data, recognizing that different people involved in the process needed to interact with the system and the data in different ways.

On the networking front, the project experimented with the use of *data mules* to collect data from nodes that were too distant to reliably form a connected network. Dogs roaming the vineyard wore motes that would collect data from proximate sensors and bring it back for aggregation and processing.

Researchers at Harvard in collaboration with UNC, UNH, and Instituto Geofísico in Ecuador employed a sensor network to study Volcán Reventador, an active volcano in northern Ecuador [77], using seismic and acoustic sensors. The network was required to cover a large area, resulting in a linear network topology with as many as six hops. The high sample rates required exceeded the bandwidth of the radio links, so the sensor nodes were programmed to detect discrete events to sample and transmit. Millisecond-level time synchronization between the nodes was implemented [46] so that observations of the same event from multiple nodes could be correlated.

In addition to monitoring native species, sensor networks have also been applied to tracking invasive ones. Researchers at CSIRO and UNSW implemented a sensor network to track cane toads [32], an invasive species introduced to Australia for pest control in sugar cane crops that are now causing great problems. The toads were identified and tracked by recording sampled audio and detecting toad vocalizations with an FFT and machine learning.

One of the key developments was the ability to perform high-bandwidth data processing on the nodes themselves, rather than transmitting all of the data back to a central server and processing it there, alleviating network bandwidth and power requirements. They also used a hybrid architecture, with both *resource-rich* ARM-based nodes, capable of both collecting and processing audio data, and less powerful Mica2 nodes capable of capturing audio but not processing. The lower-power, lower-cost nodes enabled much greater coverage, capturing data from more locations that was then sent to a nearby resource-rich node for processing. Spatial redundancy of the network was exploited to interleave sampling and transmission between nodes that are close enough to detect the same acoustic signals.

Work by Watras et al. at the University of Wisconsin-Madison [76] on wireless sensor networks for remote monitoring of wetland sites is extremely relevant to the work being done at Tidmarsh. They present an evaluation of two different sensor networks, one built from the Mica2 platform and the other from commercially available data loggers with wireless capabilities (now becoming more readily available). They present data that show the value of long-term in-situ monitoring for wetland sites. Notably, they cite the management of the huge volume of data produced by such networks as one of the next major challenges that needs to be addressed in sensor network research, particularly when it comes to dissemination of the data.

Many examples of sensor networks are installed at high density over a limited geographic area. The National Science Foundation's National Ecological Observatory Network (NEON) [35, 53] is a sensor network at continental scale, consisting of a network of high-end research stations distributed across North America. NEON aims to collect long-term data to further understand how the environment is changing. The project operates under an open-access data model.

Sensor networks have also been used successfully in citizen science projects. In the months following the 2011 Tōhoku earthquakes and subsequent Fukushima Daiichi nuclear accident, Safecast [5, 66] organized volunteers using home-made sensors to map and monitor radiation across Japan, collecting and publishing the data on their web site. The project has since expanded to include air quality and other environmental data worldwide.

2.1.4 Commercial Sensor Platforms and Smart Agriculture

Towards the end of the 2000s, wireless sensor technology began becoming increasingly available outside of academic research. Initially, most devices on the market were hardware that originated in research, such as the Mica mote, made commercially available by companies such as Crossbow (later acquired by Memsic [48]). These were generally available as bare circuit boards (PCBs) providing the core processing and networking functionality of a sensor node but requiring substantial integration work to add sensors and make a fully functional field-ready device.

The early platforms evolved into more complete systems. A good example of these transitional devices is the sensor node ecosystem offered by Libelium [41]. While the core platform heavily resembles early sensor node hardware (and is still available in bare-board form), their standard sensor node, which costs about $500[1], is packaged in a waterproof enclosure. It does not include any sensors (aside from an accelerometer), which are added by plugging up to six probes into waterproof connectors on the bottom of the enclosure. A wide range of different sensors are available, with price varying depending on sensor type. The nodes also do not contain a power source and must be connected to an external battery and solar panel, which adds about $40 to the per-node cost (more for higher-power options). The nodes can be ordered with sev-

[1]Prices are from the 2016 Libelium Product Catalog, converted from EUR.

eral different radios, including 802.15.4, ZigBee, Wi-Fi, LoRa, and cellular. A gateway allows bridging between one of the low-power radio protocols to Wi-Fi (about $675) or cellular (about $960, plus cellular service fees). They provide software that aggregates the data from the gateway and stores it in a local database as well as having the ability to send data to many cloud-based data platforms.

While the Libelium nodes and similar offerings are "plug-and-play," their flexibility means that a fair bit of integration work is required to assemble nodes, sensing, power, networking, and data management to arrive at a complete system. Towards the later half of the 2010s, many new products have appeared that are truly ready-to-use, trading flexibility and control over data for units that are easily installed and configured by someone without specialized knowledge. This has enabled smart agriculture, allowing farmers to install sensor nodes that report data that may be used to control irrigation and fertilization, determine when to harvest, and so on. This is a rapidly developing area and an exhaustive review of the available products would be difficult, but a few representative samples are included below.

Sensoterra [68] sells a product targeted at monitoring soil moisture. The sensor node comes in a $300 single-depth version, or a $500[2] multi-depth version, which takes six measurements along the 90 cm stake that also supports the electronics above-ground. They monitor soil moisture only, though the multi-depth unit has a single soil temperature sensor midway along its length. Sensoterra does not offer a gateway, but the nodes, which communicate via the LoRa wireless protocol, can join any existing LoRaWAN network or use a gateway from another vendor. There is no subscription fee for their cloud service that aggregates and displays the data, and they provide an open application programming interface (API) for accessing the data through other software. The Sensoterra node does have a limited lifespan: it contains a non-rechargeable, non-replaceable battery that they estimate will last about three years.

[2]Pricing for quantities greater than 10 units.

Teralytic [72] produces a sensor node capable of very comprehensive soil monitoring as well as basic environmental sensing. It measures soil moisture, salinity, nitrogen, phosphorus, and potassium at three depths, as well as aeration, respiration, air temperature, light, and humidity. It is significantly more expensive, with nodes leased at $14,500 for a 3-year lease on a 10-node kit which includes the gateway, or $1,500 per-node. The gateway bridges between the LoRa network used by the nodes and cellular, reporting data to their cloud service which includes extensive analytics and farming recommendations and does provide an open API.

Pycno Agriculture [58] makes a mid-range node for basic soil and environmental monitoring. It measures solar radiation, air temperature, humidity, soil temperature, and soil moisture. The soil sensing can be extended to multiple depths by plugging on additional segments. The nodes are solar-powered with a rechargeable battery and communicate over a LoRa network. Rather than using a dedicated gateway, the nodes communicate via a "master" sensor, which includes a cellular radio but is otherwise identical. A four-node kit with three regular nodes and one master costs $2,200. Data are sent to a cloud service, which has a monthly cost of $12 plus $2 per additional node beyond the first four, and includes the cellular plan. The cloud service displays data and provides an API for external access.

Prior to wireless sensor networks, the primary technology for long-term in-situ monitoring was the datalogger. As dataloggers have evolved the capability to upload data, they have essentially become wireless sensor nodes. Campbell Scientific [9], one of the big players in the datalogger space, produces a wide range of loggers that can be extended to automatically transmit data via networks including cellular, wi-fi, and satellite, or point-to-point VHF or UHF links. These systems are extremely modular, providing significant flexibility but require significant integration work to create a complete field-ready system. Their cost and complexity makes them better suited to a small number of high-cost, high-quality sensing installations rather than large distributed networks. Campbell provides software for programming and communicating with their loggers, but network connectivity and data management are generally left up to the user. They do offer a cloud-based service that can aggregate data, which is another option in the modular system.

METER Group recently introduced their ZL6 datalogger [13], which sits somewhere in between Campbell's offerings and the complete commercial sensor nodes described earlier. It focuses on soil monitoring and is compatible with METER's line of probes. The approximately $700 [10] logger communicates over the cellular network, with an annual subscription cost of about $180 per year per logger. Each logger can connect to 6 probes, with some of the available probes capable of measuring multiple parameters.

Davis Instruments [2] approaches wireless sensors from yet another direction. Known for making prosumer-grade weather stations targeted at small businesses and farms, they have extended into general-purpose customizable sensor nodes with their EnviroMonitor platform. The $400 node accepts up to four sensors and communicates with an $895 gateway, which can network with up to 32 nodes and bridges data via cellular to their cloud service. Subscription pricing is based on the transmission interval, and runs about $280/year per gateway for data reported every 5 minutes.

2.1.5 Sensor(y) Experiences

Sensing the landscape is only the first half of the story. Devices that simply record data are of limited value by themselves—at some point people need to interact with the data to explore and find meaningful results. In the most basic example, measurements are stored in a database, which can be queried to produce plots, tables, and spreadsheets. And indeed, this may be sufficient for many experiments. But working with raw data requires intimate familiarity with the sensors and experiments, as well as skills and additional tools to perform meaningful analysis.

The value of data can be greatly enhanced if it is made accessible to a broader audience. Even within the scientific community, exchange of data between research groups can be a point of significant friction. Communication of scientific data to the public is another significant challenge—and one of vital importance, especially considering issues that affect our global climate.

In this dissertation, I explore ways that people with many different backgrounds and levels of expertise can interact with, or *experience* data from a sensor network. This section summarizes some of the prior work in this area.

The Vineyard Computing work in [7] is an early example of using a graphical user interface to present data from a sensor network, presenting data to vineyard managers and workers so that it could be interpreted to make actionable decisions. Modern sensor products in the smart agriculture space take a similar approach, usually providing an online "dashboard" or smartphone app with graphical visualizations and plots. More advanced products, such as the Teralytic nodes [72], go a step further and incorporate machine learning and analytics to provide recommendations alongside the data. These interfaces are optimized for management, turning sensor networks into a readily applicable tool for more effective farming.

Sensor networks, particularly those that produce real-time data, can also be used for telepresence [14], permitting an observer to remotely experience another location. Traditional telepresence generally uses microphones and cameras to allow the remote location to be heard and seen. Virtual reality creates a similar experience for synthesized worlds rather than physical locations. *Cross-reality* [42] combines these concepts, creating a virtual environment that goes beyond a static representation of a real place by integrating sensor data and media streams.

In *Doppellab* [17], we built a cross-reality browser for sensor networks within the Media Lab building. Beginning with architectural models, we populated the 3D environment with data from multiple sensor networks, including thermostats and our own sensor nodes (Figure 2.1. Real-time and archived data were rendered with colorful visualizations. The virtual building could be nav-

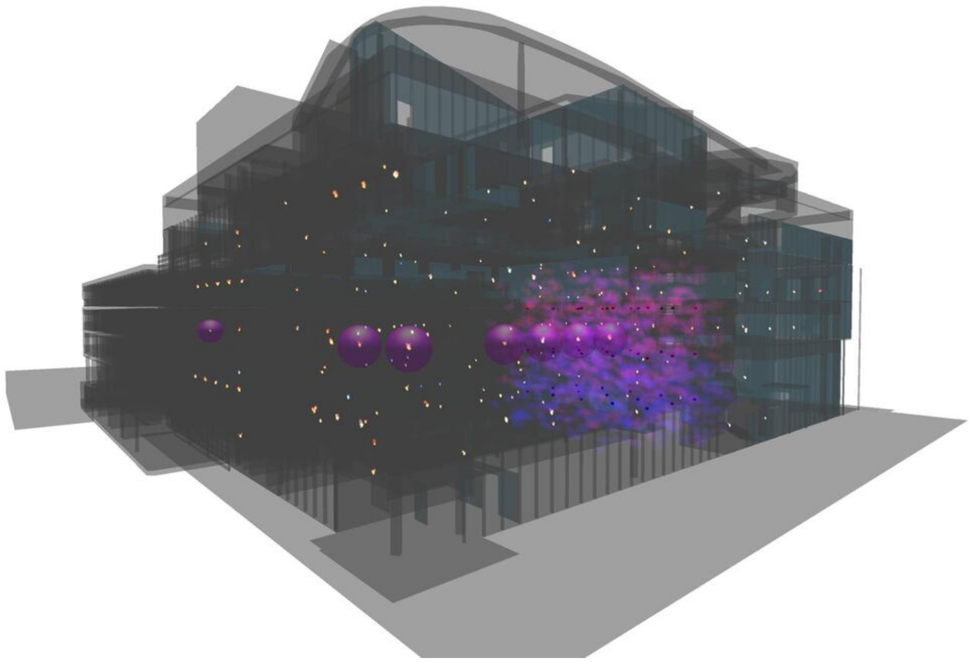

Figure 2.1: The Doppellab cross-reality sensor network browser.

igated by walking around the corridors as if one were physically there, or by flying around and zooming out for a big picture view. Real-time audio feeds, scrambled to preserve the privacy of conversations, were piped in and spatialized, giving the user a sense of the activity within the space.

Sensor networks can also be used in art, driving works that are based on physical phenomena that can even react and change in real-time. In *Forest Symphony* [67], composer Ryuichi Sakamoto instrumented trees around the world (including one that we helped set up at Tidmarsh) with bioelectric potential sensors. He streamed the data back to the installation in Japan where the data from the sensors drove a generative musical piece and visuals. The idea of this kind of work is compelling—art can be a powerful tool for provoking thought and expressing ideas, and integrating real-time data can be a way to create an even deeper connection with the natural world. But actually building and installing the sensors and maintaining the flow of data from the sensors to the piece can take significant effort and specific expertise, placing this kind of work outside the reach of many artists. Seeing this influenced the design of the Tidmarsh network: if real-time environmental data were already available, could we enable more artists to create works like *Forest Symphony*?

Chapter 3

The Tidmarsh Sensor Network

The design of the Tidmarsh sensor network was a heavily iterative process. When I started the project, I envisioned a small network with a few tens of sensor nodes and a handful of microphones streaming data to a Doppellab-like cross-reality application. At the time I had very little ecological background and knew nothing about wetland restoration. It was in this context that I designed the first sensor hardware and parts of the infrastructure. I approached it as a personal challenge to push my knowledge of designing sensors and electronics along a few dimensions. The Tidmarsh network would be larger than anything I had designed before. It would require very low-power operation. It would need to operate in a dynamic outdoor environment for a length of time without significant maintenance. And to support the cross-reality demo we wanted to build, it would need to send data in real-time.

Over the years, the scope of the project grew significantly. As I learned more about ecology and wetland restoration, and saw growing interest and excitement from artists and scientists alike, I realized that the project was going to be a lot more than a single Media Lab demo. The idea of using the network for research that could actually contribute to improving the practice of wetland restoration went from an abstract possibility to a significant driver later in the design and development of the network, especially as I started seeing traction from wetland scientists. What started as a way to remotely experience a little bit of nature through technology became a powerful set of tools for interpret-ing and telling the *stories* of a wetland across geographical and temporal scales and from many different perspectives.

This work also spans nearly a decade, over which there have been advances in technology. Some of the early design decisions, such as designing a sensor node from scratch or the choice of wireless protocols, may have been made differently given what is now commercially available. While some of the specifics in this chapter may be unique to Tidmarsh, the insights and lessons learned will hopefully apply broadly.

This chapter and the one that follows describe the design, implementation, and deployment of the entire system that has become a feature of the Tidmarsh landscape. The iterative nature of this design and the interconnection between different components of the system present a challenge in organizing this information: many facets of the design depend on other pieces and make less sense in isolation. While a chronological history of the development of the system as a whole would perhaps be most straightforward to write, here I have chosen to break it into its key components to provide a clearer overview of how these pieces fit together. I have included pointers to parts of the system described in other sections where relevant.

To provide a high-level view before diving into the details of individual subsystems, Figure 3.1 shows a schematic diagram of the hardware components of the network and the physical interconnections between them as they exist in the summer of 2020. At the bottom of the diagram are the inputs to the network: sensor nodes, microphones, and cameras (shown here as part of the base stations). These feed data up through two levels of on-site network infrastructure. The sensor nodes are divided into geographically distinct groupings ("sensor sites"), each with a "base station" that collects data from the individual inputs and relays them via high-bandwidth links to a single on-site "head end." At the head end, some data are stored and processed locally, while others are transmitted via the internet to the Media Lab for archiving and real-time streaming to many different end applications.

In addition to the physical network, data flows through several software components, which are mostly encapsulated within the "Tidmarsh Servers" box at the top of the physical system diagram. Figure 3.2 presents more of a logical flow of *data* through these software components, in contrast to Figure 3.1's *physical* interconnections. Data enter the software infrastructure at the left side of the diagram, and flow to the right through a chain of *services*, which generally store data and make it available to clients. The applications to the right are broken into two groups. *Middleware applications* take this raw data and transform it in some way, and while their outputs can be useful as-is, they primarily process and store it in ways that are useful for other applications. At the right hand side of the diagram are the end *applications*, which primarily provide a means for experiencing or viewing the data.

This chapter is divided into a few sections. In Section 3.1, I discuss several of the design considerations that played a role in the development of the network. Sections 3.2 and 3.3 describe the two generations of wireless sensor nodes. Section 3.4 describes the base stations. Section 3.5 covers the head end and the fiber and wireless links that connect the base stations to the head end.

After presenting these key system components, Section 3.8 describes how they have been deployed across multiple sensor sites at Tidmarsh.

3.1 Designing The Sensor Node

The low-power wireless sensor node is one of the key data inputs in the network. It consists of four main components. *Sensors* measure the environment, scheduled and coordinated by a *microcontroller* (MCU) that acts as the node's "brain." A low-power digital *radio* transmits those measurements, sending them on to the next part of the network infrastructure. And finally, a *power source* provides the energy.

To date, there have been two significant versions of the hardware. The "first-generation" or "v1" node dates to the very beginning of the project. The "second-

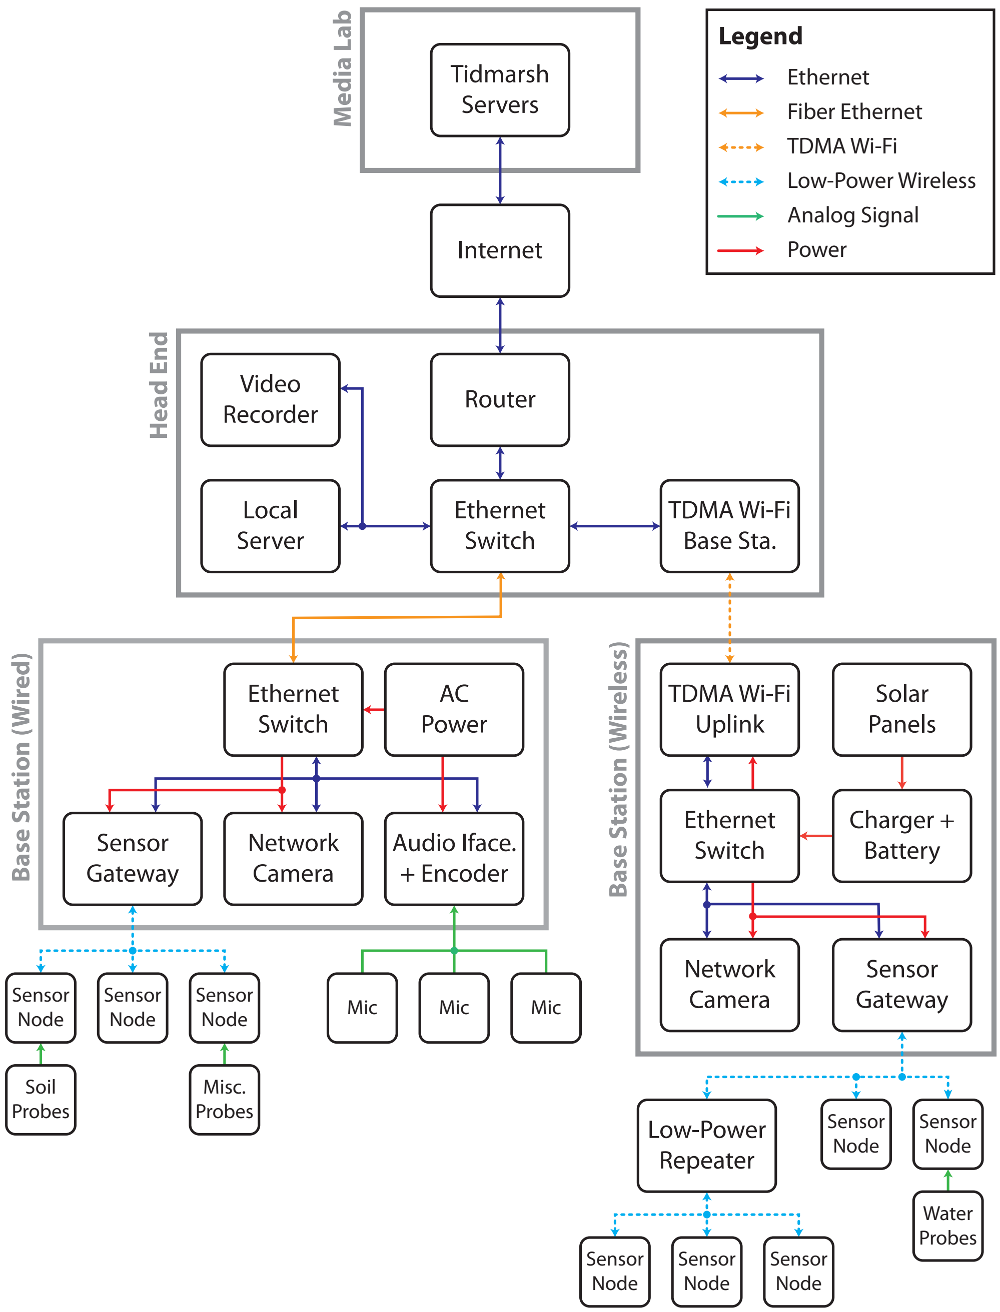

Figure 3.1: Schematic diagram showing the physical structure of the Tidmarsh sensor network as of the summer of 2020. Note that there are more base stations, sensor nodes, cameras, etc. than represented here.

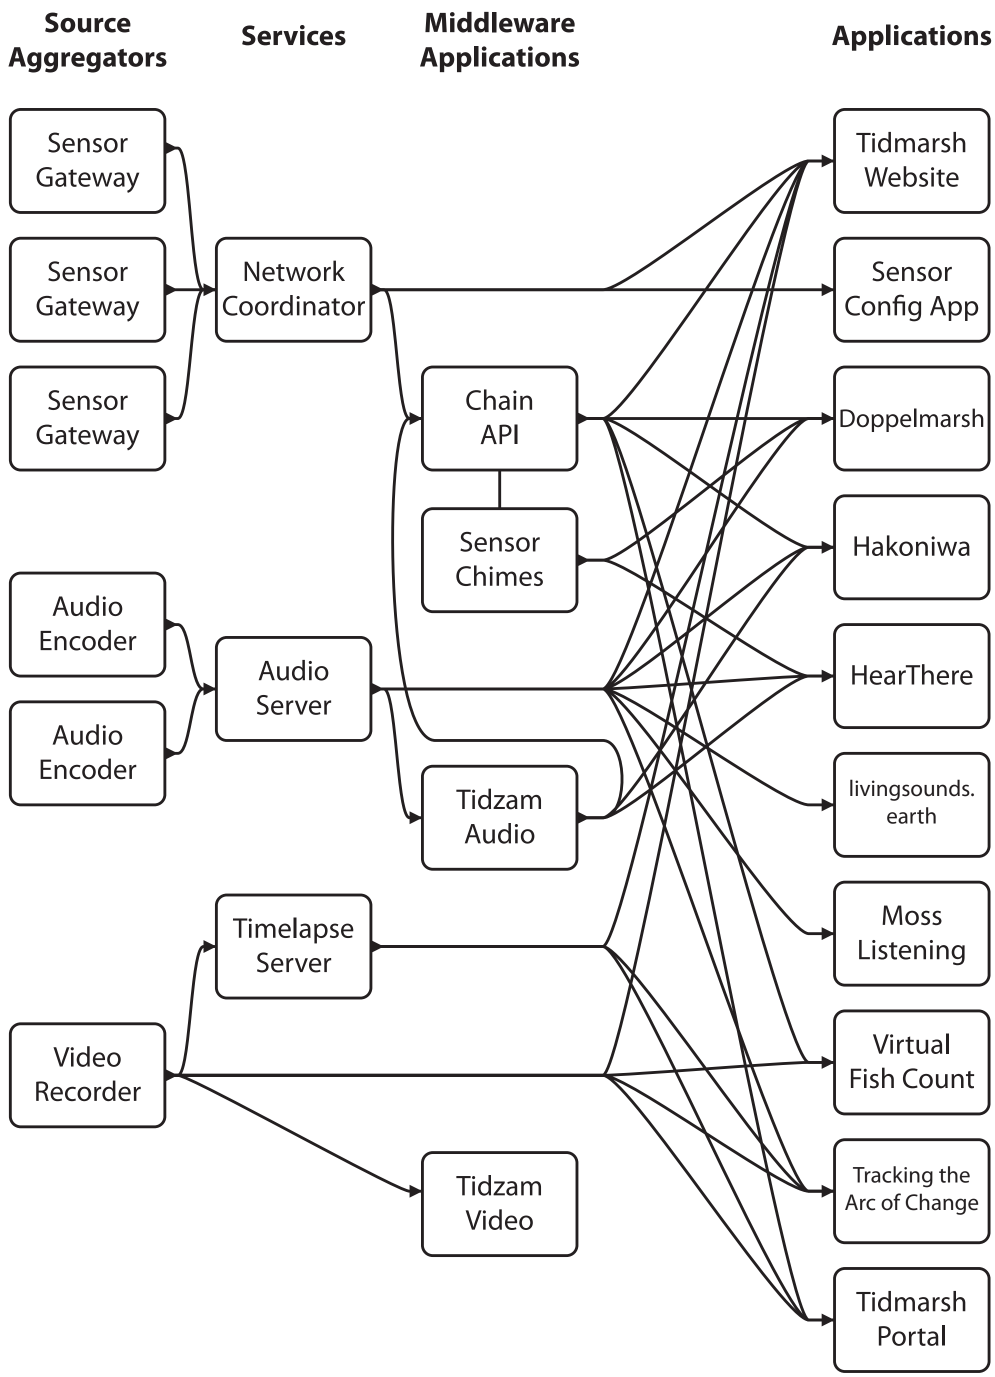

Figure 3.2: Data flow from sensor network inputs, through infrastructure services, to a selection of applications.

generation" or "v2" node design improved on the first, incorporating many of the findings from earlier deployments to produce a more capable and extensible node.

3.2 First-Generation Sensor Nodes

3.2.1 Sensors

The sensors that were integrated into the first-generation sensor node platform were chosen to be readily available from electronics suppliers, reasonable cost, low power, PCB[1]-mountable, and easily ported to the environment where required without excessively modifying the enclosure. These sensors include temperature and humidity, atmospheric pressure, ambient light, and acceleration.

These sensors are all small, relatively inexpensive, and low power. The biggest "cost" of including a chip-size sensor in a sensor node are the extra manufacturing steps to expose it to the environment so that it can sense what it needs to (Section 3.2.3) without compromising the waterproofing of the node. The selection of sensors on the v1 node is somewhat arbitrary; if a sensor was available, compatible with the power budget, and didn't excessively complicate the manufacturing, and looked interesting, I included it. At the time the v1 node was designed, I was primarily designing for a Doppellab-like cross-reality experience; any type of environmental data suited that purpose.

The temperature and humidity sensors are precise enough to measure variations between sensor nodes when deployed a few meters apart, which might provide insight into microclimates due to the topography or proximity to water.

Atmospheric pressure does not vary across even the entirety of a site as large as Tidmarsh (unless the weather conditions are *extreme*), but it is economical

[1] Printed Circuit Board

to include it in all of the sensor nodes. In theory, the pressure sensor is precise enough to measure differences in altitude on the order of tens of centimeters and could be used to help localize the sensor nodes in 3D space, though we have not yet used it for that purpose.

The light sensor is also useful for measuring variation between sensors, which again might point to microclimates due to one sensor node being shaded more than another. The light sensors can also track the progression of shadows across the landscape as clouds move overhead or the sun sets behind a ridge, gradually casting a shadow across the network. The sensor is a part designed for ambient light detection (such as for automatic brightness adjustment of a mobile device) and thus is tuned to a wavelength response curve that mimics the human eye. The light sensor also has an infrared channel, but the acrylic window in the sensor node blocks most IR.

I included the accelerometer as we were initially intrigued by the idea of hanging sensor nodes in trees, where the accelerometer might be a proxy for sensing wind. However, most of the sensor nodes have been deployed on fairly rigid stakes in the ground in areas without trees, and the accelerometer was never activated in a deployment of first-generation sensor nodes.

3.2.2 Audio DSP

The first-generation sensor node also included an optional feature that would permit audio streaming. This would have to be very limited, as the batteries could not supply enough power for very long and the 250 kbps theoretical maximum bandwidth of the 802.15.4 network (shared between all nodes) could not support very many simultaneous audio streams. Nevertheless, I designed the sensor node with this as an option. I included a footprint on the PCB for a VS1063A audio codec chip, which includes a microphone preamplifier, analog to digital converter, and digital signal processor capable of encoding to MP3 or Ogg Vorbis. The intent was to populate this on some of the boards, and selectively enable the audio stream on sensor nodes (augmented with solar panels to provide the extra power) only when lower-powered sensing identified interesting features, similar to the work done in [32].

While this feature was tested and did in fact work, network conditions had to be ideal to support more than one simultaneous audio stream (neither MP3 nor Ogg Vorbis deal particularly well with dropped frames, and memory on the sensor node was too limited to buffer enough data to make retransmission practical). Given the design goal of streaming as much as possible in real-time to facilitate rapid prototyping, it made more sense to use dedicated audio streaming hardware that could remain on continuously (described in detail in Chapter 4) with more appropriate power and network resources rather than compromising to make this a feature of the low-power sensor network.

3.2.3 Physical Design and Manufacturing

Enclosure and Porting

The first-generation sensor node was designed around the Hammond RP1065 enclosure, which was selected for its availability, reasonable cost, small size, and IP65 rating[2].

The case, which is approximately 85 mm square and 55 mm deep, accommodates a 72 by 76 mm square PCB with notched corners. The front side of the PCB, which is installed face-down into the case, provides more than enough room for all of the sensors and circuitry. The backside of the PCB contains a battery holder for three AA cells. The remaining volume in the case was intended for desiccant packets to prevent condensation from forming, as well as providing some space for add-on electronics modules.

In order to sense the environment, most of the sensors require ports that expose

[2]The IP, or ingress protection code, specifies the protection that an enclosure provides from contaminants in the environment. The first digit (6) indicates that the enclosure is dust-tight; the second digit (5) indicates that the enclosure can resist low-pressure jets of water, e.g. splashing or rain.

only the sensor to the properties that they sense, while keeping out everything else. This is perhaps the most challenging aspect to the mechanical design of a sensor node. It requires modifying the enclosure to add the required openings, and adding additional hardware to maintain seals so as to not compromise the waterproofing of the enclosure.

Three ports were required for the first generation sensor for humidity, light, and atmospheric pressure. The cases were modified to add three small circular holes. The modification work was performed in the Media Lab shop using a CNC-driven vertical mill. This additional manufacturing step was reasonably fast for each sensor node since only one face of the enclosure required milling and the enclosure could be quickly aligned and fixtured in the vise.

The Sensirion SHT21 combined temperature/humidity sensor needs to be exposed to moisture in the air in order to sense the humidity. Sensirion sells a specialized filter cap that snaps over the SHT21 sensor. The cap seals against a hole in the surface of the enclosure with an O-ring. The front surface is a filter membrane that is water vapor-permeable but keeps out liquid water, dust, and other contaminants. The cap snaps tightly onto the PCB and completely covers the sensor, encapsulating it in a small chamber so that water vapor can reach the sensor but not the rest of the enclosure. Unfortunately, the screw standoffs inside the Hammond enclosure are taller than the SF2 cap. To close the gap, an acrylic washer was affixed to the enclosure around the inside of the opening with cyanoacrylate (CA) glue.

Porting for the light sensor is somewhat simpler, since this only requires a transparent surface. A small acrylic disc was press-fit into the opening and sealed with 2-part epoxy on the inside of the case. The front surface of the acrylic disc was sandblasted to diffuse the light and make the sensor response less directional.

The port for the atmospheric pressure sensor allows the pressure to equalize between the inside and outside of the case by permitting air to pass through slowly but blocking water vapor and other contaminants. It is constructed

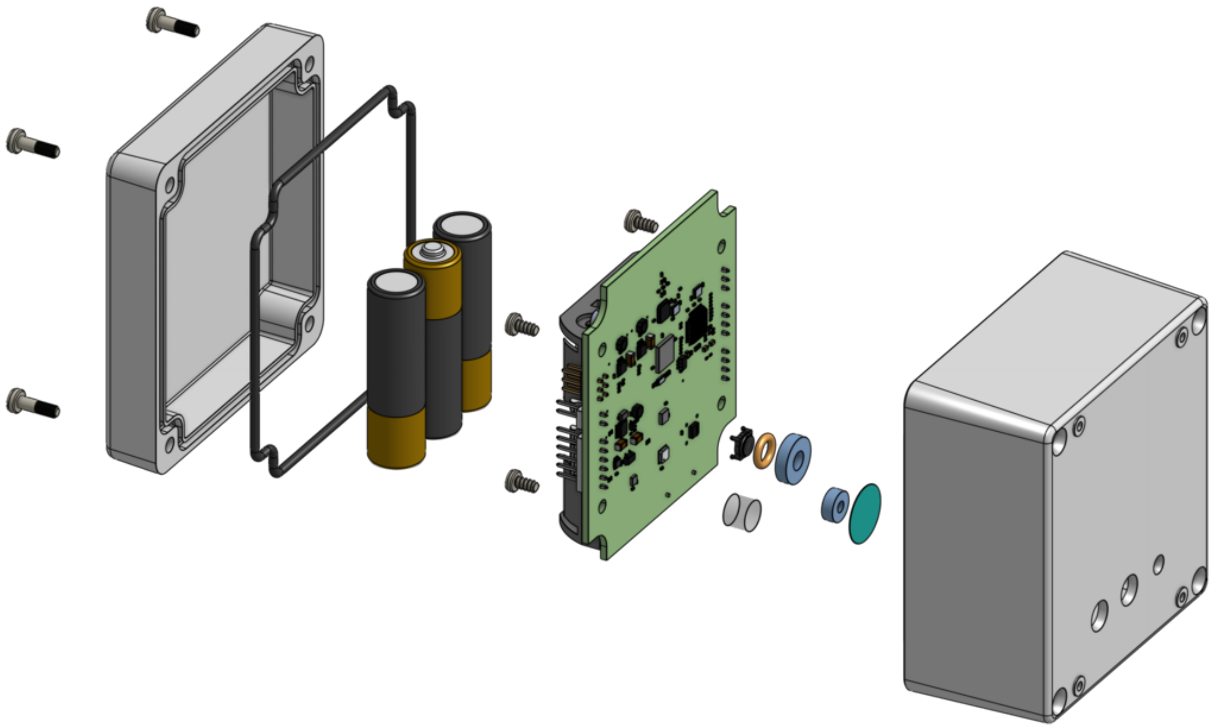

Figure 3.3: First-generation sensor node assembly, exploded view

from a composite expanded polytetrafluoroethylene (ePTFE) fabric, commonly known by the brand name Gore-Tex™. We purchased the fabric in sheet form, laser cut it into small discs, and press-fit them into the opening in the enclosure with acrylic washers before sealing around the interface with 2-part epoxy.

In addition to enabling the pressure sensor to function, the pressure port also plays a key role in the waterproofing of the sensor node. Without pressure equalization, rising atmospheric pressure effectively creates a slight partial vacuum inside the case, which often results in moist air and liquid water (e.g. raindrops) being drawn in through any imperfections in the seals.

PCBs and Assembly

The printed circuit board (PCB) design was sent to a contract manufacturer for full turnkey manufacturing. We received fully assembled boards ready to be integrated into the final node assembly.

To complete the assembly of the sensor node, the PCB is screwed into the plastic standoffs in the bottom of the prepared enclosure with four stainless steel screws. This holds the PCB in place and compresses the O-ring seals against the front of the case.

An exploded view of the complete assembly is shown in Figure 3.3.

3.2.4 Firmware

The sensor nodes run a custom real-time operating system (RTOS). The core of the RTOS is a task scheduler, which uses the microcontroller's 32 kHz internal real-time clock[3] (RTC) to execute tasks at regular intervals. At each wakeup, the scheduler evaluates the tasks that are due to execute and then calculates the time of the next scheduled task, which is used to program the next wakeup interrupt.

The main sensing task executes in two phases. At the first wakeup, the sensors are powered up and instructed to begin making a measurement. As many of the sensors require some time to complete an analog to digital conversion, the MCU goes back to sleep. After about a second has passed, the MCU wakes up again and assembles a packet using the now completed readings from each of the sensors. Only once the full packet is ready to transmit does the MCU wake up the radio. The packet is then transmitted, sensors and the radio are powered down, and the system goes back to sleep until the next execution of the sensing task.

The sensor node also understands a set of commands that can be used to configure and manage the node. Since the node is sleeping most of the time with the radio powered down, it is normally not able to receive these commands asynchronously. Instead, each time the node transmits sensor data, it listens for an acknowledgment from the base station. If the base station has queued

[3] This internal clock is an RC oscillator, suitable for scheduling tasks but not for accurate timekeeping.

commands for the sensor node, it indicates this through a byte in the acknowledgment frame. Upon receipt of this signal, the sensor node sets a timer for several seconds, and remains awake with the radio powered on to listen for further commands.

The command set allows basic management of the sensor node, such as adjusting the sensing interval, enabling and disabling certain sensors, and performing firmware updates through an over-the-air (OTA) update mechanism. This allows the code running on the sensor to be reprogrammed remotely in the field. The flash memory on the node is divided into two halves; the lower half contains the currently running program image, and the upper half is used to store a new program image. To update the firmware, it is pushed out through the network 64 bytes at a time through the command protocol. As the node receives these pieces of the new firmware image, they are written to the upper half of memory. Once the full image has been received, it is checksummed to verify the integrity of the received data. If the checksum matches, a flag is set and the node reboots into a bootloader program. The bootloader program checks for this flag, re-verifies the checksum, and then copies the new image into the active section of the flash memory. Finally, the bootloader reboots the system using the new firmware. This scheme allows the application code to do the heavy work of receiving the new firmware over the network and the bootloader can be relatively simple as it only has to copy the new image.

3.2.5 Processing and Communication

The microcontroller (MCU) is the Atmel ATxmega128A4U, an 8-bit processor using the AVR instruction set. This MCU was chosen for its low power consumption, extensive peripherals, and a well-supported open source toolchain. While more capable 32-bit ARM processors are now readily available, the simplicity and familiarity of the 8-bit AVR made it an attractive choice. Since the network is designed to transmit all data in real-time for maximum flexibility in prototyping, additional processing power on the node is generally not required. The MCU's function is simply to orchestrate the process of collecting measurements from the sensors and transmitting them.

The communication is handled by the Atmel AT86RF233 radio. It is a 2.4 GHz radio that implements the IEEE 802.15.4 PHY and MAC layers. 802.15.4 is best known as the protocol underlying ZigBee, which is most frequently used in home automation. The full ZigBee protocol is somewhat complex and software implementations are generally proprietary. To avoid this complexity, I based the Tidmarsh network protocol on a simpler stack called Atmel Lightweight Mesh (LWM) [4], which includes many of the desirable features of similar protocols like ZigBee, such as basic mesh networking, without a lot of the overhead and complexity.

I extended the base LWM protocol with a compact binary representation for encoding sensor data, as well as the command set described above.

The antenna is a PCB trace inverted-F design, based on a Texas Instruments application note [1]. A U.FL coaxial connector also allows the connection of an external antenna by changing the orientation of the final capacitor in the matching network to divert the signal to the connector.

3.2.6 Power

The main power source for the sensor node is a set of three AA-size batteries. The intent was that a set of non-rechargeable batteries would power the node for several years before requiring replacement.

The 4.5 V nominal battery voltage is stepped down to 2.1 V with a micropower switching regulator (LTC3103) with a nominal 1.8 μA quiescent current. The MCU and radio are most efficient at 1.8 V, but the sensors have higher minimum operating voltages. The SHT21 temperature and humidity sensor has the highest minimum voltage, which defined the 2.1 V system voltage. A separate 1.8 V regulator supplied power to the optional audio codec, which could be turned off independently of the MCU to save power.

To enable higher power nodes (such as repeaters that stay powered up in order to forward packets in a mesh network, or nodes with external sensors with higher power requirements) I included a solar charger on the PCB. This was implemented with the Linear Technology LT3652, which implements maximum power point tracking (MPPT) to match the impedance of the charger to the solar panel. To use this optional feature, the batteries would be replaced with rechargeable nickel metal hydride (NiMH). This feature turned out to not work very well: extra holes had to be drilled to connect the solar panel, the charge current was set too high, and the charger turned out to not support NiMH[4].

3.2.7 Expansion

The sensor nodes were designed with the possibility for expansion in mind. The on-board sensors were intended to provide a solid baseline for sensing the environment, but are certainly not exhaustive. It made sense to be able to customize the sensors attached to some sensor nodes, both in order to use sensors that were more expensive (either in cost or in power consumption) at lower density, and to add sensors that only made sense in particular locations (such as water sensors for nodes near water).

On the first-generation sensor nodes, the expansion capability was provided through a set of connectors on the back side of the PCB on either side of the battery holder. These connectors broke out a handful of extra GPIO pins from the microcontroller, including a few with analog-to-digital (A/D) capability, and I²C and asynchronous serial buses.

These pins were not practically usable as-is for connecting external probes directly. The limited input voltage range and low input impedance of the internal A/D, and the available supply voltages (2.1 V system voltage and 4.5 V

[4]An early version of the LT3652 documentation implied compatibility with NiMH, which requires different fast-charging techniques that most other battery chemistries. This was actually only possible with additional external circuitry, and Linear removed any mention of NiMH from later revisions of the datasheet.

unregulated battery voltage) made it poorly matched to most off-the-shelf external probes. Furthermore, neither supply could be switched independently, so a directly connected sensor could not be easily powered down unless it included its own sleep mode.

The intent was that add-on sensors would connect via a specialized PCB that would contain whatever buffer circuitry and power management were required for each individual sensor that might be connected. However, this required a full PCB design/manufacturing cycle for each sensor, creating too much friction for practical expansion. In an attempt to solve this problem and make a more generally expandable sensor node, I developed a flexible add-on box that would connect to the main board via the I^2C bus, and provided four analog channels via a 12-bit A/D with a 3.3V range. Two variable power supplies could be digitally programmed to any voltage between 0.8 V and the battery voltage (this is a feature inspired by programmable dataloggers) or turned off completely when not in use.

The expansion box also included an interface for the 1-Wire protocol, which allows multiple low-cost digital temperature probes (DS18B20) to be connected to the same bus.

3.2.8 Testing, Evaluation, and Lessons Learned

We manufactured a total of 150 v1 sensor nodes. At the peak of the deployments a total of 64 were concurrently installed in the field at Tidmarsh. These nodes were used in several deployments (described in Section 3.8) spanning continuously from early tests in late 2013 through the summer of 2017, at which point we began replacing them with second-generation nodes.

While this design worked well for close to four years as our primary sensor platform, there were plenty of bumps in the road and lessons learned along the way, many significantly informing the design of the second-generation platform that would follow. These findings are detailed in the following sections.

Sensor Performance

The most significant fault with the design was that the temperature sensor was very susceptible to solar heating in direct sunlight. The sensor was thermally coupled to the main PCB, and the exposed black filter membrane at the front of the case would heat up especially if the front of the node was oriented towards the sun, sometimes reading as much as 10°C higher than the actual ambient air temperature. The light sensor on the front of the case was similarly sensitive to the orientation of the sensor node.

Power

In the first winter, I observed a power issue related to the network discovery protocol. The base station, which was powered by a 100 W solar panel and 100 Ah lead-acid battery, failed to keep the gateway and Wi-Fi uplink powered through week-long stretches of overcast skies. Even when there was a sunny day, the short winter days and low angle of the sun meant the system struggled to refill the battery, inevitably leading to network downtime.

The parameters of the base station discovery protocol on the sensor nodes were tuned too aggressively: after 32 failures in reaching the gateway when transmitting sensor data, the node would begin searching for a new gateway. To do this, it had to transmit on all available channels and dwell on each listening for a response, consuming significantly more power than normal. The nodes would repeat this every few minutes until re-establishing communication with a gateway. With the gateway offline for significant lengths of time, this depleted a significant amount of energy from the batteries that were intended to last for years. The network discovery algorithm could ultimately be re-tuned in the firmware and I pushed out new code via the OTA update mechanism, but this demonstrated why it might be useful to have rechargeable batteries in an experimental sensor node platform. Even if the node consumes very little power under normal operation, unforeseen conditions (which arise more frequently in an experimental network) can burn through power that is impossible to recover without opening individual nodes and changing batteries.

Extreme temperatures and debugging adventures

The second winter with sensors deployed brought another unexpected challenge. Shortly after dramatically expanding the sensor deployment in cells 3 and 4 (Section 3.8.3), nodes began reporting obviously incorrect temperatures before crashing and restarting[5]. After a long debugging session (which involved a sensor node in a small refrigerator at the lab with a thick bundle of cables leading back to various testgear), I concluded that the falling winter temperatures were slightly depressing the setpoint on the main voltage regulator due to component drift[6]. The temperature/humidity sensor normally operated at the minimum end of its voltage range, and the slightly lower system voltage (especially as the system came out of sleep mode and the voltage would droop slightly before the regulator could react to the increased load) was causing the sensor to behave erratically, reporting incorrect temperatures before latching up the I²C bus and preventing further communication between the MCU and any sensor on the board.

Fortunately, it was possible to deploy a software fix—testing in the refrigerated setup confirmed that delaying the temperature measurement to later in the cycle after waking the processor allowed the voltage to stabilize and the sensor would operate reliably. Deploying this fix during the winter proved to be a more significant challenge. After rebooting, the sensor nodes would operate normally until the scheduler ran the sensing task, which would attempt to read from the temperature sensor and cause the node to crash. This meant the sensor node would only run for about 30 seconds between reboots. The OTA update mechanism was not fast enough to transfer the complete firmware image within this time—it would make it almost to the end before the sensing

[5]The MCU has a watchdog timer that is reset each time through the main program loop; if the program crashes and stops resetting the timer, the MCU is automatically rebooted when the timer expires.

[6]Every electronic component is a temperature sensor, whether you want it to be or not.

task would execute and the node would reboot, losing all progress. (In this version of the firmware it was not possible to resume a partial transfer).

To work around this challenge, I ended up creating a special version of the firmware that had everything that wasn't essential to the OTA process stripped out. This resulted in a much smaller image that could be fully transferred within 30 seconds, which could be pushed out before the node crashed. The full firmware image with the fixed temperature sensing could then be fully transferred without crashing.

Death by "natural" causes

Deploying sensors in nature requires that they withstand the environment for long periods of time. While not a truly extreme environment like the antarctic or outer space, wetlands and electronics do not readily coexist.

At least one sensor node failed several months into its operational life due to the omission of the O-ring sealing the temperature/humidity sensor against the case. While the location of the gap was reasonably well-shielded from the rain, ants found their way inside the enclosure (Figure 3.4) and decided that it was a fantastic spot to build a nest. The resulting moisture from a thriving ant colony eventually caused the sensor node to fail.

Several more nodes failed due to being fully submerged as water levels rose; this was not something the nodes were designed to withstand for any length of time.

Wear and tear

The majority of the sensor nodes survived their four-year deployment without issue, and were still operational when we replaced them with second-generation hardware. They were, however, certainly starting to show their age. The ABS plastic enclosures were not intended for long-term exposure

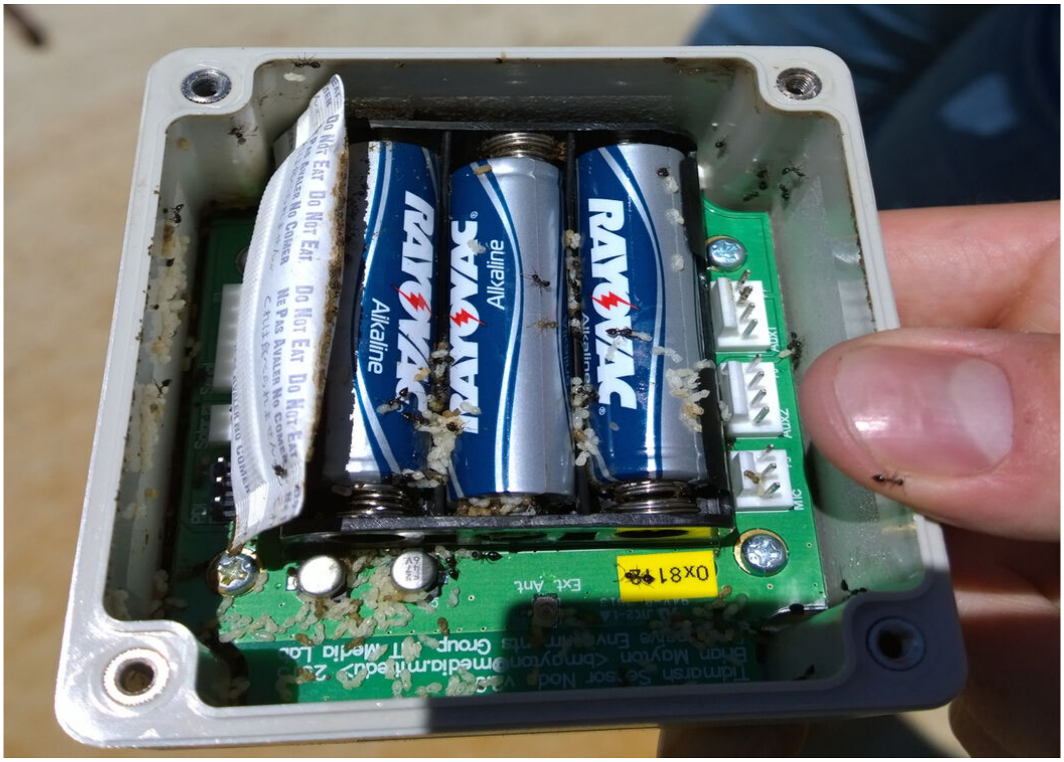

Figure 3.4: Ants nesting inside a sensor node.

to ultraviolet from the sun and were beginning to yellow and become brittle. None had actually failed due to this, but clearly this would be a weak point in a longer-term installation. Similarly, the exposed ePTFE fabric that formed the pressure equalization valve on the front of the exposure had been bleached by years of sun exposure, and some were beginning to fray around the edges.

Expansion issues

The expansion capability of the first-generation sensor node was a significant shortcoming in the design. While the add-on expansion box (Section 3.2.7) effectively added general purpose capability for using external probes, practical considerations made it very difficult to use. As an external box, it meant expansion was more of a special exception than a first-class feature of every sensor node. Wiring sensors into the expansion box was a lot of work, requiring a hole to be drilled and waterproofed to allow the cable for each probe to enter. Minimal space for wiring inside the box made connecting a full set of four probes very tedious, especially when working in the field, and was prone to mistakes and spotty connections. Finally, the expansion boxes were programmed by manually writing a configuration file and physically uploading it to the EEPROM on the board by plugging in a laptop. Again, this was very tedious and error-prone, and required going back out into the field and opening the box if any configuration changes needed to be made.

Once installed and configured, the expansion boxes worked well; the features they provided were well suited for connecting a variety of probes, including the soil moisture and temperature sensors we were most interested in.

Secondarily, the expansion interface reinforced the need for a rechargeable power source. The increased power consumption of the additional probes shortened the life of the sensor nodes that had them.

3.3 Second-Generation Sensor Nodes

As described in the previous section, I learned a significant amount from the deployments and testing of the first-generation of sensor nodes. Using these findings, in 2016 I began designing the second-generation platform. The following sections describe the changes that were made to improve upon the design.

3.3.1 Expansion

The most significant engineering finding from the first-generation deployments was that extensibility was a very important feature. The basic sensors internal to the node were sufficient for experimenting with Doppellab-like sensor-driven cross-reality environments, but the network would be significantly more effective at answering real ecological questions if we could readily customize

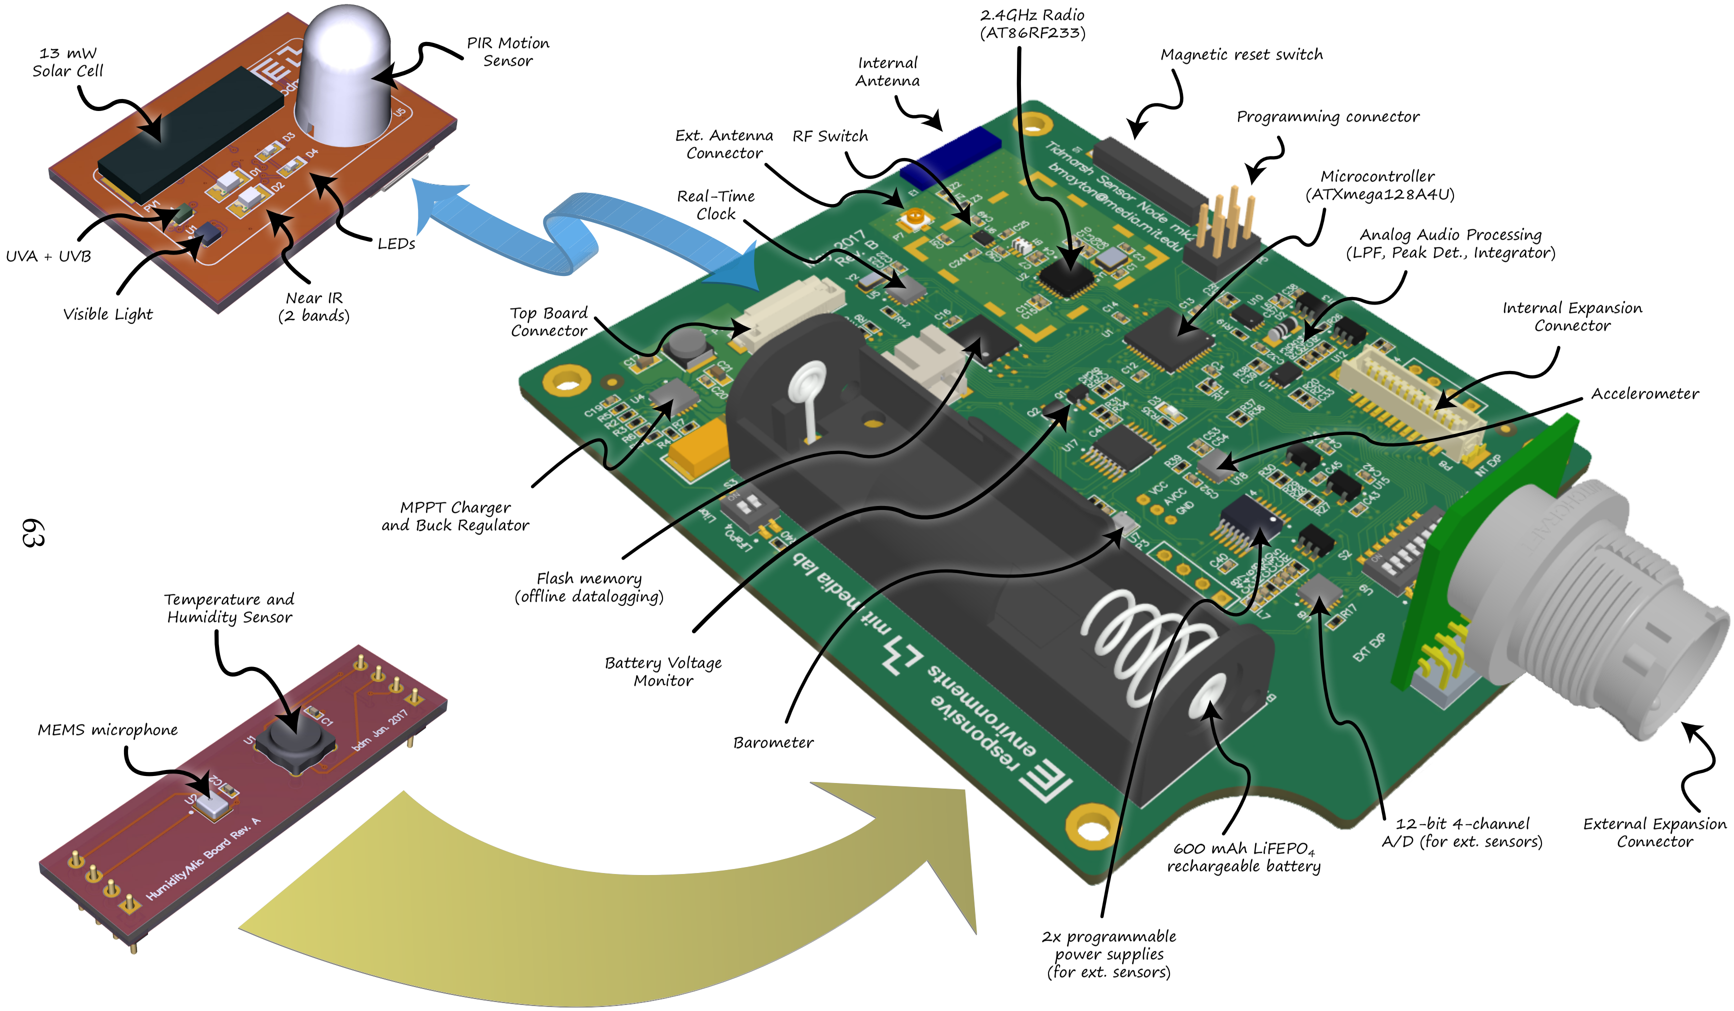

Figure 3.5: Second-generation sensor node electronics.

it and add external probes. Many questions about the effectiveness of wetland restoration naturally concern water, and while humidity is part of that story, it is important to observe what is happening in the soil and in the water. As burying the sensor nodes creates wireless communication issues (among others), external soil moisture and temperature probes are important add-ons.

At the time I was sketching the designs for the second-generation node, we had already begun planning large-scale soil monitoring deployments, potentially with hundreds of probes. Connecting and configuring them as we had done with the first-generation nodes would have required far too much effort to be practical.

These observations led to several important design decisions for the second-generation node. First, the same features provided by the expansion box in the old system would be included as standard features of the new node. This includes four analog inputs connected to a 12-bit A/D with an 0-3 V range, two programmable and switchable power supplies, I²C, asynchronous serial, and a 1-Wire bus.

Second, to streamline the connection of external sensors, these expansion functions are exposed through a 9-pin[7] circular connector with an IP68 rating[8]. The connector mounts directly to the main PCB (installed by the contract manufacturer) and secures to the enclosure with a nut, which also holds a captive protective cover in place.

To connect external probes, we ordered a set of custom pigtail cables that have the overmolded circular connector on one end of a 30 cm cable, and the other end fans out into unterminated color-coded wires. To connect probes to these wires, 3M Scotchlok™ connectors may be used. These connectors allow two or three unstripped wires to be inserted and the connection is made by compressing a plastic button which pushes insulation-displacement contacts into

[7]Some functions, such as the 1-Wire bus and asynchronous serial port, are multiplexed with the analog inputs.

[8]IP68 means that it will withstand full immersion in water at 1 meter or more.

the wires. The inside of the connector is filled with petroleum jelly, so the resulting connection is completely waterproof. This arrangement allows complex probe assemblies to be put together in just a few minutes with basic tools. I enclosed the resulting splices in a small plastic box to provide extra protection and for aesthetic reasons, but this enclosure need not be waterproof as the splices themselves are. The complete instructions (including photos) for the probe assemblies are included in Appendix B.

Third, when installing the sensor node and probes in the field, the sensor node needs to be configured to know which external probes are installed so that it knows what to measure and how to report it. Instead of manual configuration files and specialized programming equipment, I extended the task scheduler in the real-time operating system to also include a "sensor scheduler". The sensor scheduler is configured by defining *sensors* and *sampling tasks*.

Sensor definitions are primarily used by the network coordinator, and include information about how to decode the received value into meaningful data: what the metric is called, what the units will be, and what needs to be don to transform the raw value to those units (such as applying a scale factor or an offset).

Sampling tasks are sent to the sensor node itself, and specify a sampling interval, voltages for the programmable power supplies, a "setup time" (delay between enabling the power supplies and reading the A/D value), and a mapping between *sensors* and physical analog channels or 1-Wire slots. Multiple sampling tasks may be defined so that not all sensors need to have the same sampling interval.

The same configuration mechanism also allows *derived sensors* to be defined. These are recomputed (by the network coordinator) when the primary metric that they reference is updated. They can compute an arbitrary function (in JavaScript syntax) using any combination of metrics from the sensor node (including the internal sensors), named constants, and *calibration parameters* (which are stored per-node). For example, a calibrated pH output could be defined as a derived sensor that takes a raw voltage from a pH electrode, applies temperature compensation using the reading from a temperature probe, and finally applies a scale factor and offset for an individually calibrated electrode.

This system allows essentially any collection of analog or 1-Wire sensors to be connected to the node and configured to output meaningful values to the database without modifying the firmware on the node. The configuration can be written in YAML or JSON syntax, but is generally specified and edited through the network coordinator's graphical interface (Section 3.6).

3.3.2 Power

To avoid the first-generation node's issues with non-rechargeable batteries, the second-generation node uses a rechargeable battery instead, with a small solar panel (like the one you'd find in a calculator) providing the energy to recharge it. The solar panel provides about 13 mW of power in direct sunlight—this would take weeks to recharge the 2 watt-hour battery from fully discharged while still powering the sensor node, but is well in excess of the approximately 50 μW average power consumption of the sensor node[9].

Battery charging in an outdoor setting requires some care. Lithium ion (Li-ion) cells will form crystals that can cause shorts when charged quickly below 0°C. I selected lithium iron phosphate ($LiFePO_4$) cells; these have an energy density somewhat lower than Li-ion and lower nominal voltage (3.2 V rather than 3.7 V), but are generally more tolerant of abuse. Charging below 0°C still requires care, but the very low current from the solar panel shouldn't cause damage. They can be charged in the same manner as Li-ion, but with a lower

[9]This number is an estimate; while I intended to make careful measurements of power consumed by various operations with the final sensor node firmware and hardware, COVID-related lockdowns have prevented me from easily accessing the equipment required to do so. These numbers are based on a quiescent power draw of 12 μW (measured) and about 30-40 ms of awake time per minute at about 30-90 mW depending on how long the radio remains on to receive the acknowledgment frame. External probes that require power will of course increase the average power consumption.

charge termination voltage. They are readily available in AA (or 14550) size as they are commonly used in older solar-powered garden lights. The sensor node uses a single AA cell with a 600 mAh capacity (just under 2 watt-hours). Even with no sunlight, the sensor node should run for over a year (the node's power consumption is low enough that self-discharge is significant).

Charging is controlled by the Texas Instruments bq25570, which is a single-chip charge controller and low-power switching regulator intended for energy-harvesting applications. This chip also provides the main 2.8 V system voltage for the node. It includes brownout protection that shuts the sensor node down if the voltage falls below a programmed threshold, allowing the battery to recharge partially before re-enabling it.

With the rechargeable battery and solar cell, the only limit to the second-generation node's lifespan is from components and the battery physically wearing out. And importantly, it is able to recover from unexpected circumstances (such as firmware bugs, network issues, or updates) that temporarily draw more power.

3.3.3 Sensors and Performance

The solar heating issue with the temperature sensor was improved by moving the temperature/humidity sensor to a small daughterboard offset from the main PCB. This both moves the sensor forward to attain the correct spacing to seal with the case and thermally isolates the sensor from heat conducted through the main PCB. This does not completely eliminate solar heating, but the effect is significantly less pronounced than it was on the first-generation node. (Further improving accuracy in the sun would likely require an external temperature sensor with a large radiation shield or forced aspiration.) We further mitigate the effects of solar heating by installing all sensor nodes with the front face (where the temperature sensor is ported) pointing north, so that the sun never strikes it directly.

The new design also adds several new sensors. The new porting arrangement for the solar cell, described in the following section, permits more optical sensors to be easily added. In addition to the visible light sensor (which has been upgraded to sense up to 256 klx) I added an ultraviolet sensor that provides separate UVA and UVB channels, which can also be used to compute the UV Index. I also added two infrared photodiodes: one with a clear lens and broadband sensitivity, and the other with a daylight-blocking lens sensitive only to the 800-1050 nm range.

A MEMS microphone, ported through the front of the case, can sense basic audio features. It is wired to an analog low-power integrator and peak detector, and can be used to measure both average and peak audio levels over a long sampling interval while the MCU is asleep. The raw output is also wired to the MCU's A/D, which can be used to sample a short window of audio and compute a fast-fourier transform. The microphone and filter circuitry have a bandwidth of about 4 kHz.

At the top of the case, a passive infrared (PIR) motion sensor was added with the goal of potentially detecting large wildlife as it passes by.

3.3.4 Physical Design and Manufacturing

The physical design of the sensor node was also altered for the second-generation node (Figure 3.6). I switched to a slightly larger enclosure made from UV-stabilized polycarbonate for better longevity in the sun and an IP67 rating for better water resistance. To streamline the assembly process, the modifications to the enclosure (milling and drilling holes for sensors) were contracted out to the case manufacturer.

The new design added the solar panel and several light sensors, all of which need to be exposed to the environment in an optically transparent but waterproof manner. This is accomplished by milling out the top of the case to create a rectangular opening, affixing an additional PCB in place, and casting over

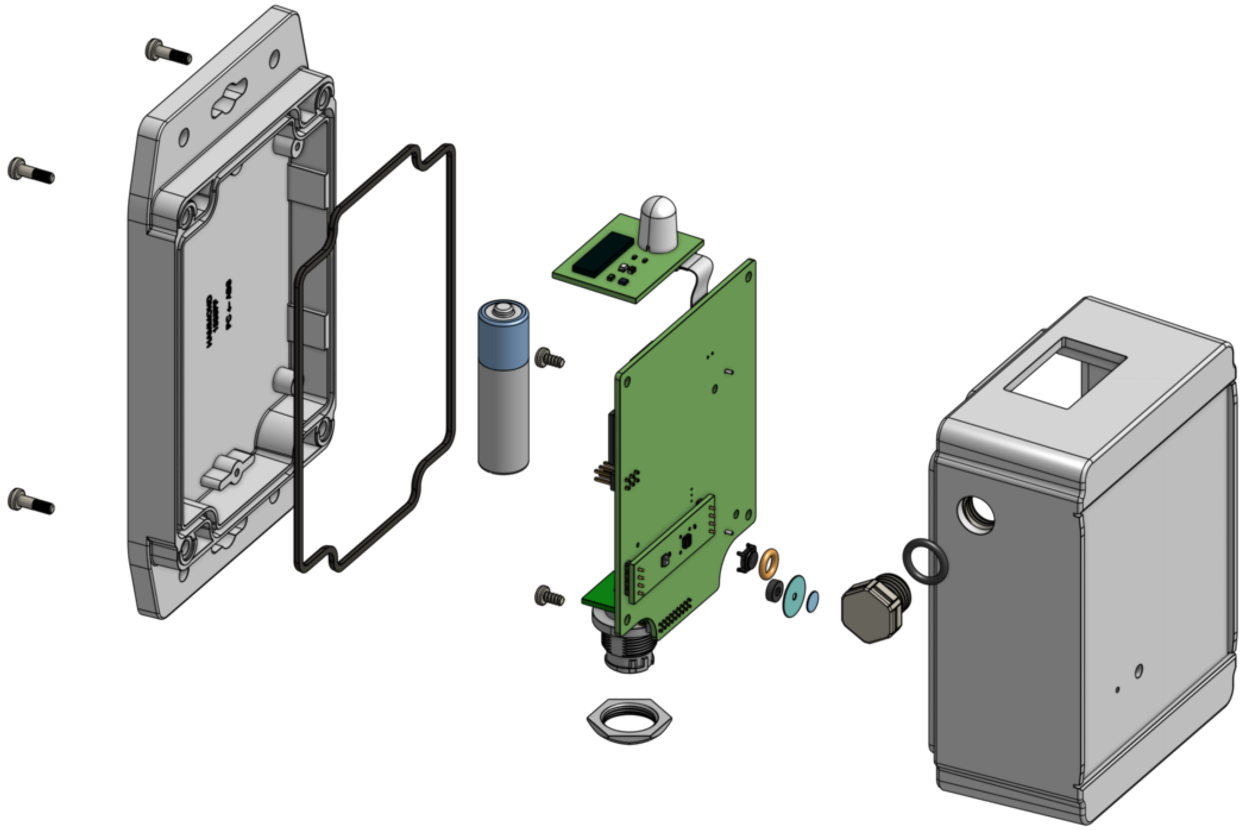

Figure 3.6: Second-generation sensor node assembly, exploded view

the top to form a watertight seal. This potting material needs to be optically clear (including UV and IR), not yellow or degrade in the sun, and not crack or come loose under thermal cycling. I evaluated several specialized epoxy resin compounds before settling on Smooth-On Solaris, which is a clear silicone rubber compound intended for potting solar panels.

This top PCB also accommodates the PIR motion sensor, which protrudes from the top of the casting.

The pressure vent assembly was changed from the homemade ePTFE fabric arrangement to an off-the-shelf threaded plug. The plug contains the same ePTFE membrane, but it is shielded from UV radiation inside the plug, making it easier to assemble and longer lasting.

Table 3.1: Sensor Node v2 Unit Cost

Components	$117.79
PCB assembly	$48.71
Enclosure and machining	$16.55
Battery	$4.00
Sensor node unit cost	$187.05

Table 3.2: Soil Moisture External Probe Assembly Cost

METER Environment EC-5 soil moisture probe	(2x) $90.00
DS18B20 waterproof temperature probe	(2x) $2.80
Switchcraft EN-3 pigtail cable (custom part)	$21.95
Splice box	$4.40
Probe add-on cost	$211.95

3.3.5 Cost

The unit cost of the v2 sensor node is summarized in Table 3.1. In addition to the per-unit cost, the non-recurring setup costs for the manufacturing process totaled about $2000. These prices were for a manufacturing run of 300 units, and do not include the labor for the final assembly steps of installing the boards in the enclosure, adding seals and vents, programming firmware, and testing, which were performed by volunteers from the Responsive Environments group and myself.

Sensor nodes used in the soil moisture experiments (Chapter 5) were augmented with an external probe assembly that plugs into the waterproof connector on the bottom of the node. The cost of each probe assembly is summarized in Table 3.2.

3.4 Base Stations

The sensor nodes report data to base stations, which form the bridge between the low-power 802.15.4 network and the internet. There are currently five base stations in the network, each serving as the hub for a sensor site. The base sta-

Figure 3.7: The off-grid base station at the herring site.

tions vary in their power and network configurations. Two are completely off-grid and employ solar panels and large batteries for power and directional Wi-Fi antennas for the internet uplink. The others have wired power and internet.

The component of the base station that actually bridges between the sensor network and the internet is the gateway. It is based around a small single-board computer (SBC) and a custom PCB that contains the 802.15.4 radio and communicates with the SBC via USB. The radio in the gateway uses the same transceiver as the sensor node, but with an external low-noise amplifier (LNA) used while receiving, power amplifier (PA) used while transmitting, and antenna diversity switch, which dynamically selects whichever of the two antennas provides the best link for each packet.

The software on the gateway connects to the network coordinator (described in the following section). Most network logic is controlled by the coordinator with the gateway acting as a transparent bridge, but some time sensitive operations (such as time synchronization and acknowledgments) are handled directly by the gateway.

Figure 3.7 shows the base station at the herring site, which is one of the two off-grid stations. The dish antenna at the left of the pole is the Wi-Fi uplink and points back toward the network head end. The gateway is the box to the right; the 2.4GHz antennas are on top. The solar panel hides the storage battery, charge controller, and network switch underneath. All of the components of the base station are connected to the network switch via Ethernet, which also provides power to each device.

3.5 Head End and Backhaul

The base stations connect to the internet via our "head end". This is the location on-site where our internet connection to the outside world is located. The original head end was located at the Agway barn on the northeast corner of the property (refer to Figure A.1 in Appendix A) and later moved to the Morton barn.

The off-grid base stations (the west side and the herring installation) connect to the head end via antennas on top of the barn, which form one end of the directional Wi-Fi links.

To reach the base stations in the impoundment, wireless links would not work due to a forested region in between. Instead, we ran fiber optic cable buried in the ground to these stations.

The head end also houses servers and routers (Figure 3.8) that manage the on-site network. It also contains a high-performance GPU server that is used for on-site data processing, especially on the video streams where we do not have enough upstream bandwidth to transfer all of the data to the Media Lab in real-time.

Figures A.5 and A.6 in Appendix A show the head end and base station loca-

Figure 3.8: The indoor portion of the Tidmarsh network head end.

tions as well as the fiber and wireless backhaul links at two different points in time.

3.6 Network Coordinator and Management Interfaces

The network coordinator software is responsible for managing the sensor network, configuring sensor nodes, and collecting all of the data. The latest version of the software is implemented in the Node JavaScript framework, which is well-suited to the asynchronous nature of low-power sensor network communication.

Each gateway connects to the network coordinator over the internet. The coordinator dynamically keeps track of which sensor nodes are accessible via which gateway and automatically routes commands appropriately.

The coordinator receives all of the sensor data from the network and decodes the binary data into individual measurements. Node-specific sensor configurations and calibration data are taken into account during this process (see Section 3.3.1 for details). The decoded data are available through the coordinator's own API in real-time, and can also be pushed to several other databases and services (the latest version of the coordinator sends data to Chain-API [64], which is Responsive Environments' in-house sensor database backed by InfluxDB, as well as MQTT [51] and Graphite [25], and can be extended with plugins to support other databases and services.)

The coordinator software also includes a user interface for configuring sensor nodes in the field. The interface is implemented as a web application and can be accessed from any browser. I designed the user interface to be easily usable on a mobile phone so that it can be accessed in the field while installing and maintaining sensor nodes.

Figure 3.9 shows several screenshots of the configuration tool, outlining the

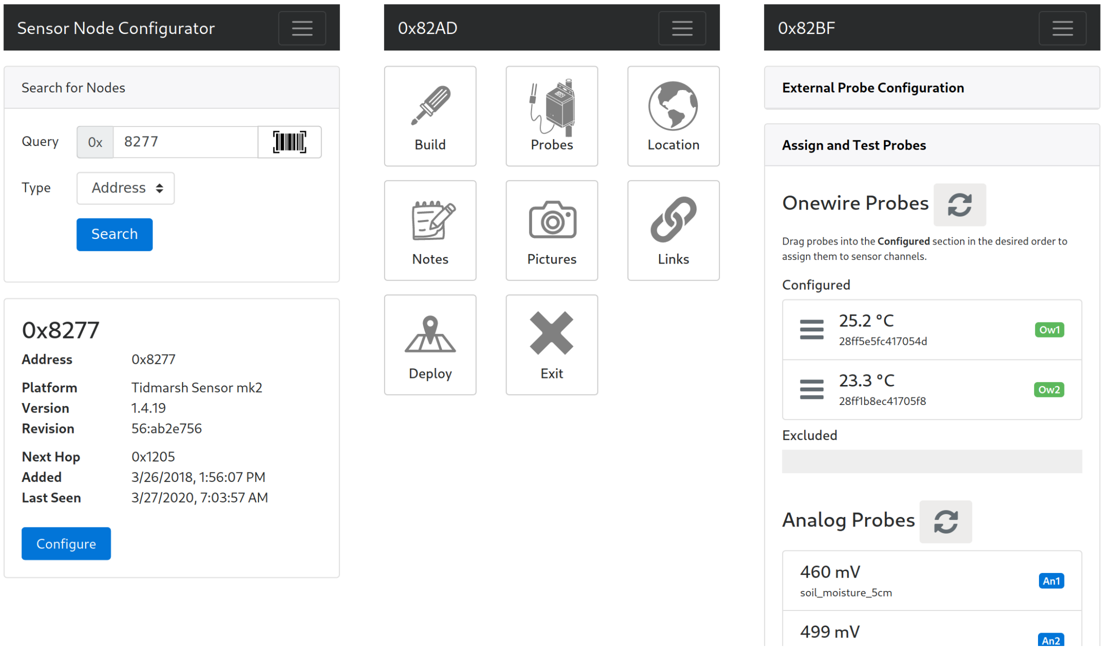

(a) selecting a node to configure **(b)** main menu **(c)** probe configuration

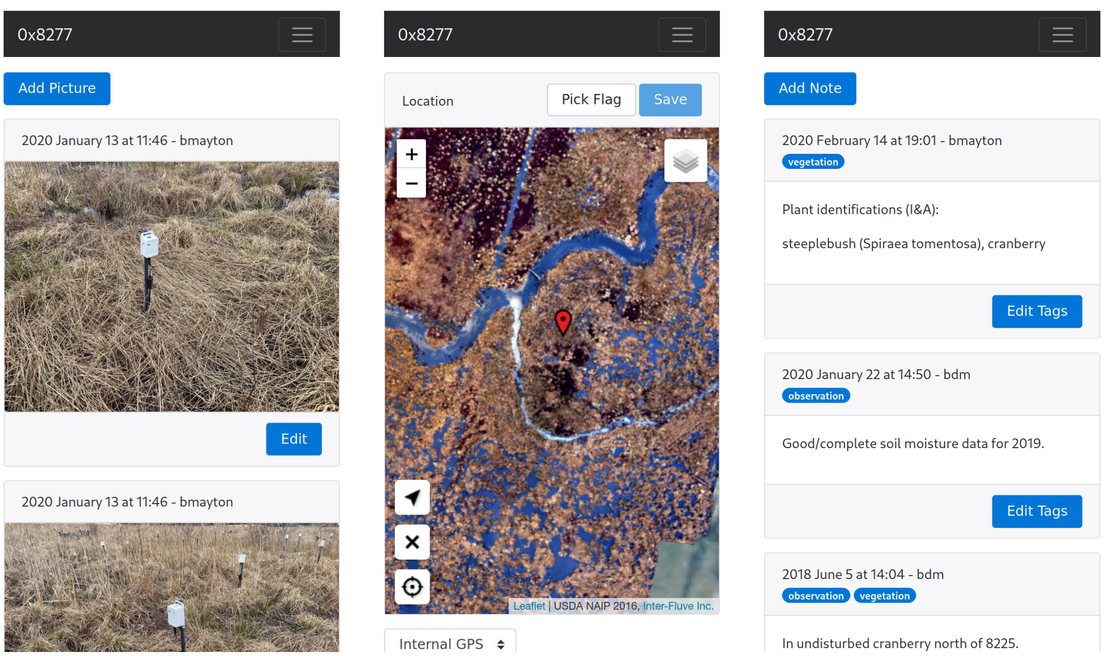

(d) node photographs **(e)** setting sensor location **(f)** notes

Figure 3.9: The sensor field configuration interface.

process of configuring and documenting a sensor node as it is installed in the field. In (a), a sensor node is selected by typing in its address (or scanning the barcode on the front of the node with the phone's camera). (b) shows the main menu that appears after selecting a node, which allows the user to pick a submodule. (c) shows the configuration of external probes; selecting this screen from the menu sends a wakeup request to the sensor node. The configuration process interactively communicates with the node, permitting the user to see live data from the probes, and storing the configuration in the node's memory as settings are updated.

Data from sensor nodes are practically worthless without metadata. It is vital to know where each sensor node is, when it was installed, and so on. Keeping track of this information can be a challenge, especially when working in the field with many people installing sensor nodes.

To address this challenge, the configuration tool has several modules for recording metadata. (d) shows a database of photos associated with each individual sensor node. Photos can be added by tapping the "Add Picture" button, which brings up the phone's camera to take quick snapshots of the sensor node, its surroundings, probes in the soil, etc.

(e) shows the interface for setting the location of a sensor node. The location can be entered by dragging the marker on the map, using the phone's GPS, or by referencing an existing marker in the database (such as a list of positions surveyed with a high-end GPS).

(f) shows the notebook, which is a generic scratchpad for comments about the sensor node.

Comments and photos can easily be added by any collaborator at any point during the lifecycle of the sensor node. These become part of the record of that particular node, and can be accessed through the API alongside the actual data. Notes and photos also support tagging with key words.

At the end of the installation process, the interface presents a deployment checklist (not shown), allowing the user to verify that each of the steps have been completed and that all of the important metadata has been entered.

The configuration tool is also integrated with the initial assembly and commissioning process, the "Build" module (not shown) is used to record information specific to the node (such as PCB serial numbers, date of assembly, etc.) and interfaces with a hardware programmer to physically load the initial code and configuration onto a brand new sensor node.

3.7 Data Storage and Access

As sensor data are streamed back to the Media Lab, they is sent to a system called Chain API [64], which is the primary means of storing and programmatically accessing sensor data from the Tidmarsh network. This software, developed by my collaborator Spencer Russell, is a database that stores data from sensor networks and metadata about those sensors. It is designed to allow applications to automatically discover the available data through following links to related resources. It allows access to archived data through a JSON+HAL [38] application programming interface (API) and to real-time sensor data streams through WebSocket connections.

3.8 Sensor Deployments

3.8.1 Prototypes

Three prototypes of the first-generation sensor node were constructed and installed at the Arm alongside the audio installation in the spring of 2013. These differed slightly from the hardware described earlier in this chapter; they used a different accelerometer and a different antenna design, and ran different firmware with a simpler network stack that was not capable of meshing.

These prototype nodes validated the hardware design, reporting data for the duration of the experiment at the Arm. The design was altered into the production hardware described in Section 3.2.

3.8.2 Cell 3 Initial Deployment

During the summer of 2013, we began manufacturing of the v1 sensor node, and started to have complete units ready in the fall. At the same time, the first of the long-range Wi-Fi links at Tidmarsh was installed, which provided connectivity to a location on the west side of the property near the berm dividing cells 3 and 4 (refer to Figure A.1 in Appendix A).

We chose this location because a nearby hill was a good vantage point for line-of-sight back to our head end antenna and thus was a fitting place for our base station. This would allow us to install sensors in both cells 3 and 4. Cell 4 was one of the wettest areas of the property prior to the construction, and seemed like somewhere we'd have the best chance of measuring any variation significant enough to be visible in a cross-reality sensor browser. To the north, cell 3 was drier and seemed likely to undergo a measurable change during the restoration.

We installed nine sensor nodes in cell 3 and on the hill with the base station. These would be our first real test of the v1 sensor node hardware in the field, and our first experience of operating through the winter (refer back to Section 3.2.8 for some of the lessons learned).

In the spring of 2014, as the ground thawed and it became warm enough to work outside, we continued expanding this installation into cell 4, adding about another 15 sensor nodes.

3.8.3 Cell 3 and 4 Large Deployment

In the fall of 2014, we expanded the sensor installation significantly, scaling up to a total of 64 sensor nodes installed in a grid pattern that covered a region of

Figure 3.10: First-generation sensor node redeployed after the construction.

cells 3 and 4 about 150 meters on a side (see Figure A.4 in Appendix A).

3.8.4 Refresh and Reinstallation

The end of 2014 and beginning of 2015 brought heavy snows, which frequently covered our sensors and base station solar panels. By early March 2015, the batteries in the nodes were starting to run low. (This was also the winter in which the issues described in Section 3.2.8 occurred.)

We went to Tidmarsh and retrieved all of the sensor nodes (save for a few that we could not find buried beneath the snow) and brought them back to the Media Lab. I replaced all of the batteries, and took the opportunity to make a few modifications to the hardware. I corrected a fault in the radio circuitry that limited the transmission range, reduced the power consumption by removing the unused solar charge controller chip and power supply circuitry for the unpopulated audio DSP, and added a waterproof connector to the bottom of each node that exposed the I^2C bus for expansion.

The refreshed nodes were reinstalled in the spring of 2015.

3.8.5 Construction

In the fall of 2015, construction to remove the berms and reshape the stream channel through cells 3 and 4 was ready to begin, which meant we needed to remove the sensor nodes to avoid having them bulldozed. We left a few in place and flagged them so that the construction would avoid them, but removed the rest.

To maintain some sensing throughout the construction, in addition to the few nodes we left, we installed about 15 nodes in the impoundment, where we were also beginning to install microphones (Chapter 4).

When we returned to reinstall the cell 3 and 4 nodes in the spring of 2016, we came back to a transformed landscape. The dry surface of the retired farm had been comparatively easy to navigate prior to the restoration, and constructing the grid pattern had been straightforward. The post-construction landscape was much more varied. We installed as many of the nodes as we could in their original locations, but found that many were now located within a body of water or in locations that were too muddy to safely reach.

3.8.6 Microtopography Cluster and Transect

In the summer of 2017, the v2 sensor hardware had been completed, and started to replace the v1 nodes as well as expand the network. I began installing a dense cluster of sensor nodes in a new site at the center of the property. We called this the "herring" site due to its proximity to a stream where the river herring run in the spring. This area of the property made heavy use of a tech-

Figure 3.11: Second-generation nodes in the microtopography cluster.

nique called microtopography, which is described in Chapter 5 along with the experiments that these nodes were installed to test.

Concurrent to installing the sensors around the herring station, we also began swapping out the v1 sensor nodes that remained in cells 3 and 4 and in the impoundment.

One of the variables in a sensor network is the arrangement of the sensor nodes. Dense clusters, like the installations described above, provide high density data in two dimensions, but require a large number of sensors to cover a significant area. Reducing the density allows a larger area to be covered, but with less spatial resolution. Transects maintain density along one axis while having none along the other, which allows a slice of the property to be monitored.

As a point of comparison to the clusters of sensors and to cover more areas of the site without adding more base stations, I drew a line between the base sta-

Figure 3.12: Sensor node monitoring saplings in the nursery.

tion at cell 4 to the east side of the property just north of the barn. To construct the transect, I planned nodes spaced about 6 meters apart. We did not have enough soil probes to include them on every node in the transect, so I planned to install these on every other node.

To date, between the herring cluster, the transect, and replacing the existing nodes in cells 3 and 4 and in the impoundment, just over 100 v2 sensor nodes have been installed.

The locations of all currently deployed sensors may be seen in Figure A.7 in Appendix A. A detailed view of the microtopography cluster is shown in Figure A.8.

3.8.7 Greenhouse and Nursery

In addition to sensors on the bog surface, we have also used them to instrument an on-site greenhouse and nursery used to start tree saplings and other wetland plants that are used in the restoration efforts. These nodes have been extended with soil probes. Measurements of moisture in plant pots (Figure 3.12) and in the nursery soil, as well as temperature and humidity in the greenhouse, are an example of how the sensor network can be used for management decisions.

Chapter 4

Audio and Video: Rich Media Streams

Audio captured from microphones in the building had been a significant part of our earlier work on Doppellab. Even though the sounds of a lab and office building might seem mundane at first, listening for a few minutes reveals a lot about the environment. There are the background sounds of the building itself: the drone of the HVAC system. There are sounds of activity: doors opening and closing and the chimes of the elevators. And the voices of the building's inhabitants: in Doppellab, we developed an algorithm for scrambling these so that speech could not be understood, but a listener still got a concrete sense of people coming and going and the tone of their interactions. This is a lot of information that can come from an inexpensive microphone.

For the purposes of creating cross-reality environments, audio has a number of useful properties. While a camera has a fundamentally limited field of view through a lens (barring complicated VR setups that require moving cameras or extensive optics on the capturing end, and head-tracking displays on the viewing end), the experience of listening to a remote audio feed is quite like the experience of actually being there. This is true even with a single microphone, but the effect is greatly enhanced in stereo and can continue to improve with the addition of more channels (though this requires more elaborate playback setups). And even with common inexpensive microphones and headphones or speakers, the quality can be high enough to sound convincingly real and not obviously like a reproduction. (The same cannot be said today about visual VR viewing devices, where limited resolution is quite apparent and the bulk of the devices themselves makes it hard to forget that one is wearing such a unit.)

Audio is also quite malleable when synthesizing cross-reality environments. Video is particularly hard to include naturally; it inevitably ends up looking like a TV screen in the virtual world or is boiled down to a few features that influence the virtual world but no longer resemble an image. Audio in its simplest form is a convincing backdrop that sets the mood—even a static stereo stream played back while the virtual world is explored can provide a similar sense as being there. Most modern game engines also support spatialization, which alters the levels and panning of multiple virtual sound sources as the player's "head" moves around the environment. (More advanced spatialization engines even simulate the Doppler effect, reverberations, and other atmospheric effects for a more realistic experience). Feeds from multiple microphones placed on virtual sources in corresponding locations in the virtual world can simulate how different sounds would be heard at different locations. While this technique can create unnatural artifacts (especially when the same sound from the real world is loud enough to be picked up by multiple microphones with different delays), the effect is not obviously jarring to most listeners. (These artifacts, and techniques for minimizing them, are the subject of my colleague Spencer Russell's dissertation [65]).

Given this and what we had learned from using audio in Doppellab, it naturally follows that audio would be a significant component of the work at Tidmarsh. In the same way that the background sounds of a building provide insight into the activity of its inhabitants, we anticipated that audio would tell us a lot about the activity of a restored wetland, perhaps even how healthy it was. Many animal species make sounds that are loud enough to be captured by microphones at some distance, so an expert could potentially assess biodiversity by listening to the audio. And nearly anyone can probably tell the difference between the sounds of a recently operational farm and a thriving wetland.

Indeed, the audio has proven to be one of the most impactful components of this work.

4.1 Initial Experiments: The Arm

Shortly after my involvement in the Tidmarsh project began, an opportunity presented itself that would end up driving our first experiments with capturing and streaming audio from Tidmarsh. (For details of the project, see *Moss Listening* in Section 6.5.1).

The first challenge was selecting a site and setting up a network connection. Our first candidate site was in the red maple swamp at the north end of the property. There was a location where water pumps had only recently been removed, and the utility poles that supplied power to the pumps and a phone line for controlling them still remained. This location would have provided us with fantastic nature sounds, as it was relatively untouched by the farming operation and was forested with a large flowing stream channel. We investigated the possibility of a DSL connection on the phone line, but ran into roadblocks with the phone company, whose systems would not allow us to proceed without a building or structure with a service address. This probably would have been a surmountable problem with enough effort (and perhaps the right connections), but we ended up moving on to other options.

Another potential location was a barn at the northeast edge of the property ("Agway" barn). This was the one structure on the property as of early 2013 that was within 200 feet of a road where Comcast had service, and we were able to sign up for a normal business internet package. Unfortunately, this was one of the least interesting spots on the property for capturing nature sounds. Its proximity to the road meant that there was a lot of traffic noise. The nearby bog surface still closely resembled a farm, with stagnant water in ditches and not a lot of animal activity. But there was a forested area to the north that held

Figure 4.1: Gershon running cable to a microphone across the Arm.

some promise.

The third and final site came to us rather serendipitously. An abutter to "the Arm," one of the two reservoirs where water was impounded at the south end of the property, took an interest in the project and graciously allowed us to put equipment in his garage and use his internet connection. The Arm was a fantastic location for capturing nature sounds, being a large pond with a floating marshy island.

Having identified two viable (and very different) sites with internet connectivity to proceed with our experiments, the next step was to design and install the technology to actually capture and stream audio. We installed two multichannel setups. At the Arm, we placed four microphones roughly in a square, with two in trees along the eastern shore and two on the island. Four hydrophones were placed underwater, below the microphones. This gave us a quadrophonic setup both in the air and underwater.

At the barn, we ran cable both into the forested area and onto the bog surface. Four microphones went into trees, two were on the bog surface, and two hydrophones went into the standing water in the ditches.

We streamed the audio back to the lab and created the first version of the tidmarsh.media.mit.edu could listen to a live stream from the Arm. We also used the live audio in a physical installation (see *Moss Listening* in Section 6.5.1).

From each of these two sites (which are indicated in Figure A.3 in Appendix A), we recorded dramatically different audio. The sounds from the Arm made it clear that it was *alive* at all times of the day, especially early in the morning as nesting geese called to each other and in the evening as the spring peepers came to life. Near the barn, we recorded almost entirely silence, with a bit of traffic noise from the nearby road. There wasn't much life on the unrestored bog, at least not that was making sound.

This installation also taught us about some of the challenges we'd face in keeping cabled audio installations up and running. After the early hydrophone failures at the Arm (Section 4.4.2), the installation ran smoothly for a while. Water slowly found its way into our microphones and they would need to be swapped (this led to the improved design in Section 4.4.1).

At the barn, the audio streams did not stay running for very long. When we returned to investigate, the microphone we had hung from a tree at the edge of the forest had been pulled down, and the connector had been ripped off of the end of the cable. The vandalism was likely the work of dirt bikers, who were upset at being told not to ride on the property and likely thought that the microphones were for surveillance. The audio interface was also severely electrically damaged. We don't know whether that was also the work of the vandals, who could have connected power to the severed cable to deliberately damage our equipment, or a coincidentally timed lightning strike that found its way down our cables. Either way, these bumps in the road foretold some of the challenges that we would continue to face with our audio installations.

Figure 4.2: The wireless stereo audio streaming setup.

4.2 Cell 3 Streaming Box

As we began installing sensors on the west side of the property in cells 3 and 4, we wanted to stream audio from there as well. However, in this location we were entirely off the grid: our only power came from solar. Running a desktop computer and multichannel audio interface, as we had at the Arm and the Agway barn, would take more power than we could practically draw from the sun and keep running 24/7.

Instead, I designed a battery-powered setup around a single-board computer (SBC) and stereo USB audio interface. The Raspberry Pi SBC, which had recently become available at a significantly lower price than anything like it, did not have quite enough horsepower in its first version to reliably encode two channels of audio in real time. We found that the BeagleBone Black, which was Texas Instruments' take on a small, low-cost SBC did, and consumed less power.

The resulting "streaming box" (shown in Figure 4.2) contained the BeagleBone, a 2-channel USB audio interface (we used interfaces from both Focusrite and Behringer), an 18 amp-hour lead acid battery, and a maximum power point tracking (MPPT) capable solar charge controller. Outside the box, we connected a 60 watt solar panel and a pair of microphones. The audio was transmitted back to the base station through Wi-Fi, and a v1 sensor node monitored the humidity inside the box and the charge level on the battery. It could also send a command to the charge controller to power cycle the setup if there was a problem.

The first iteration of the stereo audio box was installed in Cell 3 in 2014, with one microphone near the box and another about 30 meters away (Figure 4.3). It went through a few revisions, including a change to the Intel Edison SBC, which consumed less power to perform the same function. In 2015 as we relocated the sensor nodes to allow the construction to proceed, the audio box moved up on the hill near the base station where it would be out of the way.

4.3 Multichannel Streaming: The Impoundment

As the construction work started and we transitioned our focus to the impoundment area (which would remain undisturbed by the construction), the setup there provided an opportunity to construct a larger network of microphones. The impoundment was separated from our head end by a forested area, which meant that a wireless link would be impractical. Instead, we connected the base station by running fiber optic cable through the forest back to the head end. Since we were digging trenches and pulling cable anyway, it made sense to include AC power as well. This meant that the base station at the impoundment had both wired internet and power, and could support a larger setup.

Figure 4.3: A microphone in cell 3.

To act as an audio interface, we used the rackmount version of the Behringer X32 digital mixer (Figure 4.4). This gave us 16 inputs with microphone preamplifiers, and we could apply digital signal processing (DSP) in the mixer before sending the audio to the computer to be encoded and streamed.

From the base station, we ran microphones into the forest on one side and into the marsh on the other.

We expanded this setup with a satellite box another 100 meters down the path, which connected via CAT6 cable to the mixer. The satellite box provided another 16 inputs.

The full network of microphones covers an area about 350 meters in diameter. The number of microphones has fluctuated over the years. At times we have had over 20 channels operating at once.

Figure 4.4: Audio installation at the Impoundment: microphones ready for deployment and in the field, main base station (right), satellite box (lower left).

Our decision to implement a cabled network had two main advantages. One is that no separate power is required at the microphone locations; this is provided by the same cable that carries the signal. The other is that each channel is synchronized with the others since they are all captured by the same audio interface.

The downside to a cabled installation is that it requires a lot of cable, which is easily damaged and requires frequent maintenance. Rodents tend to chew on everything, and microphone cable is no exception. Where possible, we buried the cable or ran it through the water, but it is impractical to completely cover the entire length. Later in the installation, we switched from actual microphone cable with soft insulation to gel-filled CAT5e ethernet cable intended for direct burial, which seems to be less attractive to the rodents.

This cycle of maintenance has in a way led us to the optimal density for the microphones. When we started the installation, we weren't yet sure of the right density or how long we could make our cables without running into signal integrity issues. The first microphones we installed were on 30 meter cables clustered around the base station. As the installation grew, we tried longer cables and found that we could go above 150 meters without any problems, and the satellite box also added another 100 meters from the main base station. Over time as cables have failed, we've focused on fixing microphones that are further apart and in acoustically distinct environments and have removed failed microphones where the density is unnecessary.

The impoundment installation now has about 15 operational microphones, and is shown in Figure A.9 in Appendix A.

4.4 Audio Hardware

4.4.1 Microphones

While microphones are arguably one of the most commonly used sensors, microphones suitable for continuous outdoor use present a number of challenges. In order to conduct our audio experiments, we sought a microphone design with the following characteristics:

1. **High fidelity.** In addition to processing the audio data, we wanted it to sound natural and pleasing to listen to. Very high quality microphones are numerous, but are generally designed for studio use or live sound reproduction and can be rather fragile. Microphones designed for field recording (or specialized housings for them) can tolerate use in windy and rainy conditions, but generally for short periods of time.

2. **Low noise and wide dynamic range.** Nature can often be very quiet, and we didn't want the hiss from a high noise floor to become the dominant sound. At other times, nature can be quite loud.

3. **Low cost.** In order to experiment with high channel counts and put microphones into environments where they might incur damage, we wanted a pricetag lower than a typical studio microphone.

Figure 4.5: Weather-resistant omnidirectional microphone (v1), cut-away and exploded views

4. **Weather resistant.** To be suitable for permanent field installation, microphones need to withstand weather conditions such as rain and freezing temperatures, and be resilient to damage from dirt and wildlife.

5. **Low maintenance.** With large installations and microphones in locations that are difficult to access, it's undesirable for the microphones or housings to need regular service.

6. **Standard interfaces.** We wanted to use commodity low-cost equipment for capturing sound from the microphones. Standard microphone preamps for studio and live sound applications use 48 V phantom power and balanced cabling, which is frequently used with long cable runs in high-noise environments. Designing microphones that work with this standard allows us to work with readily available equipment.

Weather Resistant Omnidirectional Microphone v1

The first version of the microphone design, shown in Figure 4.5, uses a standard XLR connector as the enclosure. The microphone capsule is pressed into the end of the rubber boot with the capsule orifice facing out through the hole where the cable would normally pass. Two-part epoxy resin poured over the back of the capsule affixes it in place and forms a seal between the capsule and the electronics.

The printed circuit board contains electronics to provide the bias voltage to the capsule from the 48 V phantom power and to drive a differential signal into the cable. It is based on a circuit popularly attributed to an early transformerless Schoepps design and appearing in various forms in the DIY microphone community [21].

I built the first batch of this design with the Panasonic WM-61A capsule, which is a small electret capsule that at the time was popular for its flat frequency response and low cost. (It has since been discontinued).

The next batch, and all subsequent microphones, were built with the Primo EM-172 capsule. This was a capsule popular in the DIY wildlife recording community due to its high sensitivity and very low noise.

For the third batch of microphones, which were in place for the latter half of the installation at the Arm, I switched from standard XLR connectors to Neutrik NC3MX-HD XLR connectors. These have a stainless steel body and extra O-rings, giving them an IP65 rating. This improved the longevity of microphones in the field.

Weatherproof Omnidirectional Microphone v2

For the second revision of the microphone hardware, the main goals were improving ease of manufacture and longevity. The compact design of the v1 microphones inside the XLR connector was not strictly necessary, and added complexity to the assembly procedure. The IP65-rated NC3MX-HD connectors were also not truly well-suited to long-term outdoor use; without a pressure equalization valve, water on the seals and humid air tends to get drawn in as the atmospheric pressure changes and creates a partial vacuum inside.

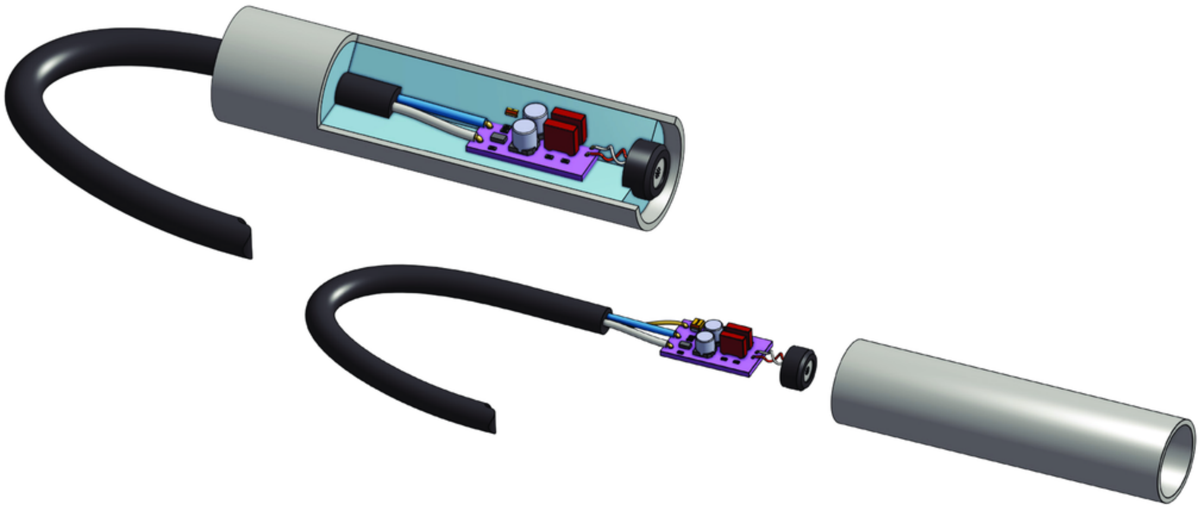

Figure 4.6: Weather-proof omnidirectional microphone (v2), cut-away and exploded views

The v2 microphone, shown in Figure 4.6, is based around a 100 mm length of aluminum tubing with an approximately 25 mm outer diameter. The electronics are attached to a short pigtil cable and placed inside, and the entire interior of the tubing is potted with a nickel-cure silicone rubber compound (Smooth-On OOMOO). This leaves only the front of the capsule exposed at the front of the tube and the pigtail cable exiting at the back. The lack of airspace inside the tube prevents condensation from forming, and the silicone completely seals the electronics away from water ingress. The flexible silicone will not pull away from the cable and capsule as it thermally expands and contracts.

The schematic of the circuitry is generally the same as the v1 microphone, but several improvements were made by substituting components. The v1 microphone used ceramic capacitors in the audio signal path, which can themselves be microphonic and are generally undesirable in audio circuits. The v2 design, which has more space to accommodate the electronics due to the larger aluminum tube, uses polyester film capacitors in these locations. Similarly, the v1 design had limited space for bulk capacitance to filter noise from the zener diode used as a regulator for the capsule bias voltage; the v2 design has space for larger electrolytic capacitors. The v2 design also switches to a lower-noise zener part. Together, these modifications improve the sound qual-

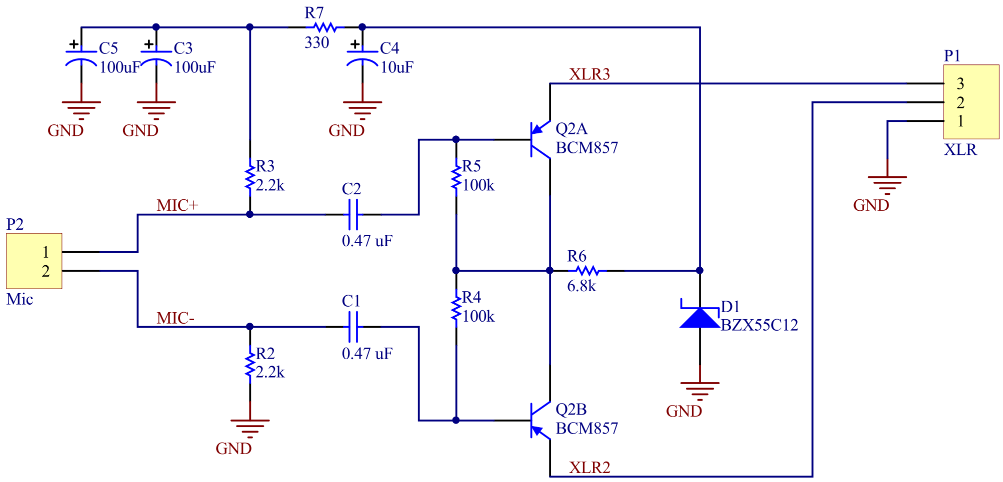

Figure 4.7: Weather-proof omnidirectional microphone (v2) schematic

ity and reduce the noise floor substantially. The two discrete transistors that form the balanced line driver were also replaced by a single-package matched pair, speeding up assembly as transistors no longer need to be hand-matched for best performance. The full schematic for the v2 microphone is shown in Figure 4.7.

We have assembled and deployed over 30 microphones of this design, beginning with the audio streaming boxes in Cell 3 and continuing through our large installation in the former impoundment. (A batch of assembled microphones, with foam windscreens, is shown at the top left of Figure 4.4). There have been a few failures due to rodents chewing off the pigtail cables and a few were damaged by a lightning strike, but the design has overall been extremely reliable. Most of the v2 microphones installed in 2015 are still problem-free almost 5 years later in 2020.

For the most recent batch of v2 microphones, the capsule was switched to the Primo EM-272 (the EM-172 has been discontinued).

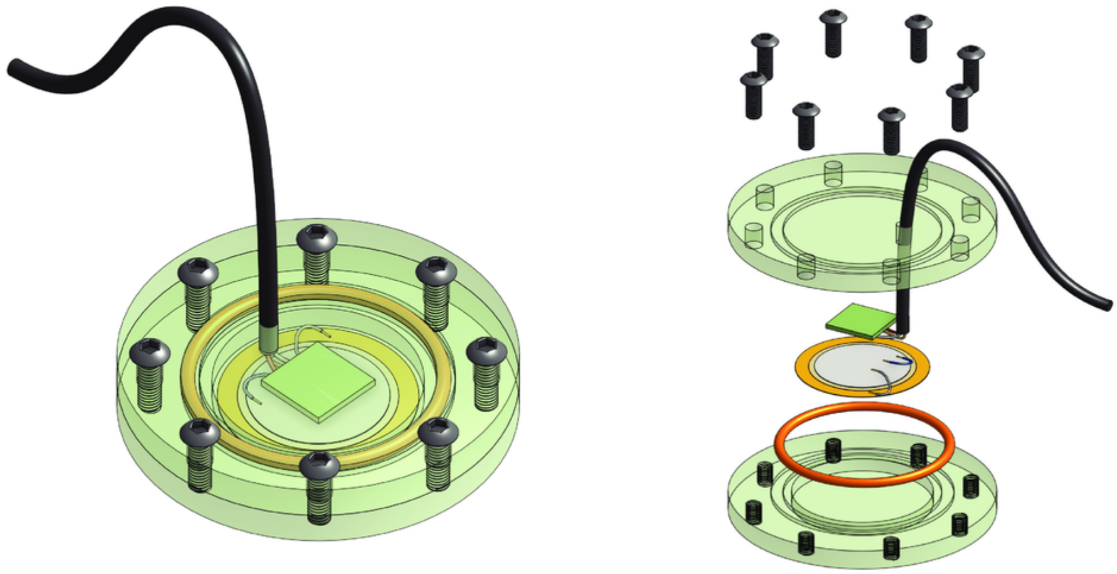

Figure 4.8: Hydrophone prototype (v1), assembled and exploded views

4.4.2 Hydrophones

When we started designing audio installations at Tidmarsh, we wanted underwater sound to be a part of the experience. True research-grade hydrophones are very expensive and thus are not conducive to experimentation with high channel counts. In proper Media Lab fashion, we began designing and fabricating hydrophones from parts on hand.

Our work here was generally aimed more at being able to create an experience of being underwater, capturing sounds that would be interesting for a human to listen to that could be used in artistic expressions. This work happened concurrently to designing and building out the sensor network and didn't receive particular focus. These *explorations* should be considered as such and are included here mostly for completeness, not as well-researched examples of proper hydrophone construction.

This is an area that remains ripe for future work. Our efforts to build and deploy a good hydrophone have been plagued with issues, especially 60 Hz noise pickup. We suspect that the proximity of the Pilgrim Station nuclear plant and the associated high-voltage distribution lines, some of which cross directly over the stream channel at the north end of the property, may be a contributing factor.

Prototype Hydrophone v1

The first of these designs, shown in Figure 4.8, uses a piezoelectric disc inside of an acrylic case consisting of two round discs sandwiched together. The acrylic was laser cut from a 6 mm cast sheet. A circular pocket was milled into one half, reducing the material thickness between the piezo disc and the water to 1.5 mm. The piezo disc was affixed into the bottom of this well with epoxy.

The air space above the piezo disc accommodates the electronics, which consist of a 48 V phantom-powered JFET-based high-impedance differential amplifier, based on the design from [62, 56] and adapted to a small surface-mount board. The PCB was glued to the half of the acrylic case opposite the piezo, and the cable was routed through a hole in the face of the acrylic and sealed with epoxy. Most units were built with a 2 meter section of flexible lavalier cable terminating in a Switchcraft EN-3 connector to attach to the long cable runs back to the recording interface. These connectors are IP68-rated when mated together and can be submerged.

The entire assembly was held together with eight stainless steel M4 screws fed through clearance holes in the top side of the case into tapped holes on the bottom. The two halves seal against an O-ring placed in a groove milled into both sides.

Benchtop testing at the lab produced results in line with what one would expect for a hydrophone cheaply assembled from parts on hand and not carefully impedance-matched to the water. While the resonance of the rigid acrylic resulted in a peaky frequency response, it did pick up the sounds of objects dropped into the bucket and the surface of the water being swished around.

Results from the installation at the Arm were less promising. The first apparent issue was significant hum pickup when the preamplifier gain was set high enough to pick up sounds. While the metal piezo disc acts as a ground plane and the amplifier and cabling are differential, this was not sufficient to avoid significant noise pickup.

The second major issue with this design was its ability to withstand water. While the O-ring seal was quite effective, inspection of dead hydrophones removed from the field revealed that water entered through the cable penetration where the only seal was a small amount of rigid epoxy resin around the flexible cable where it passed through the hole drilled in the case. In hindsight this was an obvious design flaw, but as we were rapidly building the installation at the Arm we'd hoped that it would last a bit longer than it did. Of the four hydrophones deployed, two failed within the first few hours; the other two worked for a little over a week.

Low-Cost Commercial Hydrophones

We also attempted to use low-cost commercial hydrophones. We identified the Aquarian Audio H2a-XLR as a possible reasonably priced option, at about $200 USD per unit. With a phantom-powered buffer circuit and an XLR connector, we expected it to work well with our system. However, despite being sold as *compatible* with balanced XLR systems, the units we received only used pins 1 and 2 in the XLR connector, with pin 3 left floating—the output on these units is in fact single-ended. Due to the lack of a differential output, the noise picked up by 100 meters of cable made these unusable for our purposes.

We attempted to make these units work by adding another phantom-powered buffer with a differential output to convert the single-ended signal from the hydrophone to a balanced signal before sending it down the length of the cable. While this helped, the short length of cable between the hydrophone (which is completely potted and impossible to disassemble) and the buffer circuitry still resulted in unsatisfactory hum.

Modified Microphones as Hydrophones

Our best underwater recordings were obtained accidentally when a microphone installed close to the water became partially submerged due to rising water levels. The sound was completely clear and free of hum.

This led to experimentation resulting in our most successful self-built hydrophones, in which we simply extended the silicone casting on our v2 omnidirectional microphone design to cover the entire electret capsule with a thin layer of silicone. The full shielding provided by the aluminum tubing and differential output result in much lower hum.

4.5 Video

As the network has grown, we have also added cameras as another type of rich media stream. I began adding network cameras to the base stations just prior to the start of construction with the goal of capturing the transformation through timelapse. The cameras are commercial products: we started with the Unifi UVC Pro model and have since upgraded to the UVC G3. These cameras are normally marketed towards surveillance applications and are designed for outdoor use. The power requirements are low enough that we can reasonably run one from a solar-powered base station, and the cost is very reasonable at about $150 per camera.

The cameras send video to the manufacturer's network video recorder (NVR) software, which we host on a server in our head end. I developed additional software that queries the NVR every 10 seconds and records a frame to disk, which we can later use to construct timelapse video. These still frames are also served via the website.

While we don't have enough storage infrastructure to record full-motion video constantly, the NVR software is configured to record video clips whenever motion is detected. This has captured some spectacular wildlife footage, particu-

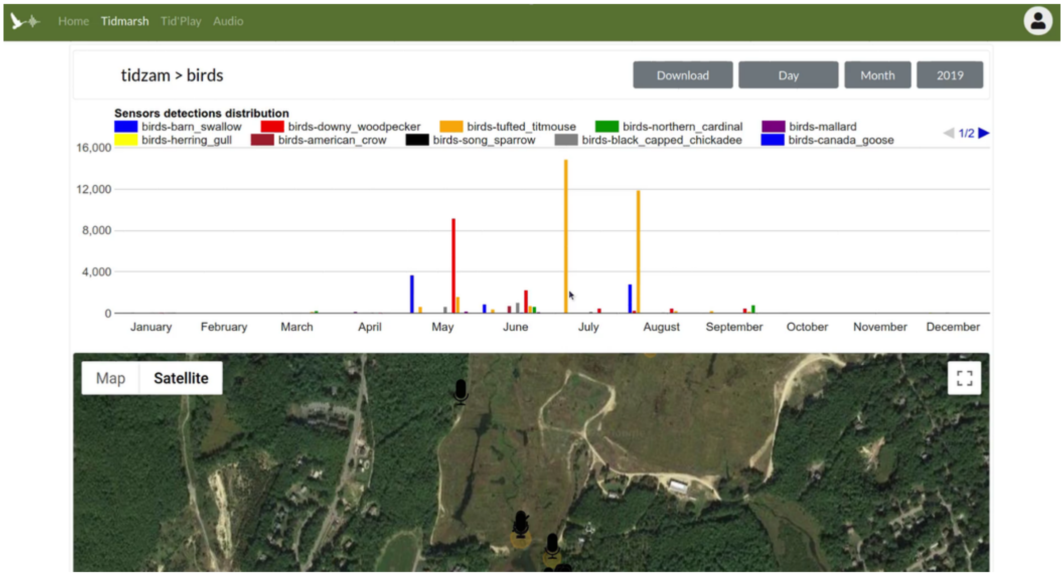

Figure 4.9: Tidzam, showing bird species classifications over a 1-year period.

larly from our camera at the herring site that overlooks a flowing stream.

4.6 Tidzam

When we started installing microphones, it was mainly because we were interested in listening to nature and hearing how the restoration would change the sounds of Tidmarsh. But these audio and video streams can also be processed to produce quantitative data, creating "virtual sensors" that look for features in the audio and video.

Tidzam [19, 18] (Figure 4.9) is an application created by my collaborator Clément Duhart that leverages machine learning techniques to process the audio and video. The audio version of Tidzam has been trained on libraries of bird and insect sounds, and can detect these in real-time in the audio streams from each microphone. It records these detections to the Chain API database, and can plot the activity history of different species over time.

Another version of Tidzam operates on the camera streams and can detect and highlight wildlife in video, pulling out video clips that not only have motion (which can easily be triggered by the wind) but actually have identifiable wildlife.

These tools demonstrate how relatively simple and inexpensive sensors like microphones can be used not only to listen to a place but to automatically monitor how it evolves over time.

Chapter 5

Soil Hydrology Experiments

As a multi-purpose sensor network, one of the goals is that the sensor data should be useful in evaluating and improving the state of the art of wetland restoration practice. As the processes repaired by restoration can span many years, long-term monitoring is an important tool for learning about restoration outcomes.

To explore the use of the sensor network for long-term monitoring, I conducted a study monitoring soil moisture across portions of the site over a period greater than one year. To design this study, I collaborated with Dr. Christine Hatch, a hydrogeologist at University of Massachusetts Amherst, to identify five research questions that the sensor network might be able to answer. These questions were as follows:

1. **How much does soil moisture vary across the site over time?**
2. **How much groundwater is present across the site?**
3. **How much does soil moisture vary across the pit and mound topography over time?** (*Pit and mound topography*, or *microtopography*, refers to a restoration technique that roughens the surface of the soil, creating hollows and hummocks (pits and mounds) on the scale of a few meters or less. The intent is to create a more heterogeneous landscape, hypothetically increasing the diversity of plant and animal life by creating varied

microclimates.

4. **How much does soil moisture vary from the pit and mound topography over time relative to the areas that are essentially undisturbed from the former farmed surface?**

5. **Is water flowing through the sand beneath the site?**

To answer these questions, I planned three experiments and sensor node installations, two of which have been completed to date. The data gathered so far address questions 1, 3, and 4. Questions 1 and 5 concern groundwater, requiring probes at greater depths, and thus remain as future work.

The first installation consists of a dense cluster of sensor nodes in an area of the site (around the "Herring" base station) where microtopography was heavily employed. Sensors were placed on various aspects of the microtopography: at the tops of mounds, the bottoms of pits, and on the north, east, west, and south faces of mounds, with multiple replicates of each position. This installation was designed to primarily address question 3, while also providing data for question 1.

The second installation took the form of a linear transect, spanning from the western boundary of the property to the eastern boundary with sensor nodes spaced 6 meters apart. A transect provides high spatial resolution in one dimension while still spanning a complete slice of the site with a reasonable number of sensor nodes. The transect as designed would contain 88 sensor nodes. By spanning across the entire site and including areas with and without microtopography, the transect was intended to address questions 4 and 1.

For these first two installations, nodes would be augmented with four external probes. Two soil moisture probes (METER Environment EC-5), which determine water content of an approximately 200 mL volume by measuring the dielectric constant of the soil, were installed at 5 and 15 centimeter depths. Temperature probes (Maxim DS18B20) were placed at the same depths—near the corresponding moisture probes but far enough away to be outside of the measurement volume.

The third installation, which has been planned but has not yet been completed due to time and resource constraints, is to consist of sensor nodes sparsely distributed across large portions of the site. The locations of the sensors were chosen to coincide with gravimetric soil moisture measurements taken prior to the construction phase of the restoration. The sensor nodes at these locations would be outfitted with the same set of probes, but installed at greater depth (25 and 45 centimeters) in order to answer questions 2 and 5 about groundwater.

5.1 Sensor Installation and Data Collection

The sensor installation began with the microtopography cluster around the "Herring" site. The first seven nodes were installed early in the summer of 2018, and were the first of the production v2 hardware to be installed. Additional sensor nodes continued to be installed through the summer of 2019, eventually reaching a total of 27 operational nodes (Figure A.8 in Appendix A).

At each sensor node location, soil was removed with a post-hole digger and an aluminum jig was used to measure depth and aid in insertion of the probes horizontally into undisturbed soil on the side of the hole (Figure 5.1). Each installation was documented with several photographs directly within the sensor configuration software. These photos show the probes installed in the soil so that the soil types can be seen and later associated with soil-specific calibration curves. The photos also document the immediate environment around the sensor node so that observations can be contextualized. After installing, configuring, and documenting the probe installation, the hole was then backfilled with the original soil.

The installation of the transect began in the autumn of 2018, starting from both the western and eastern sides of the property and moving towards the middle.

Figure 5.1: Soil moisture probes inserted into the side of a hole, with installation jig

Another small group of sensors was installed at the middle of the transect. The installation of the transect would be mostly completed during the early summer of 2019 with the assistance of a large group of volunteers from Responsive Environments. As of this writing, a total of 55 nodes have been installed along the transect. (Some locations were skipped during the installation due to the difficulty of installing sensors in deep standing water, and a short span remains between the middle and east sides of the property where sensor locations have been surveyed and marked but sensors have not yet been installed).

While it is always ideal to have an uninterrupted dataset from all sensors for the duration of an experiment, this study, which was concurrent to the development and testing of the hardware itself, was subject to the realities of prototype hardware and software. Both the staggered installation stretching over a year and sensor failures resulted in partial coverage of the study period from some of the sensor nodes. However, a good number of the sensors did produce a good record of the 2019 field season. While we did not obtain data from as many replicates of the microtopographic features as initially planned, the resulting data did provide insight into the hydrology of each sensor location. These results are discussed in the following section. Furthermore, this study showed some of the ways that a permanently installed sensor network may be useful in learning about a wetland site.

5.2 Results

I began analyzing the data in early 2020, and selected a total of 22 sensor nodes that recorded clean data for a significant fraction of the year, both from the microtopography cluster and the transect.

To facilitate discussions around the dataset, I developed a web-based tool that combines all of the documentation from each sensor node with plots of data from the 2019 year. The documentation shown includes the photos of the probes installed in the soil, photos of the sensor node and its immediate environment, and any notes stored in the sensor database, which include manufacturing information, installation comments, and research observations. A map view also shows the precise location of the sensor nodes and others nearby, and is overlaid with high-resolution drone imagery and peat depth contours.

This tool was shared with the Living Observatory community, along with specific links to the 22 selected sensors. Alexey Zinovjev and Irina Kadis provided plant species identifications in the photos around each sensor node, and I worked closely with Dr. Hatch to interpret the soil moisture data from each sensor. From our discussions, we identified a number of sensors that were representative of different soil types and microtopography and provided evidence to help answer the questions outlined earlier.

The EC-5 probes produce an analog voltage as their output. To transform this to volumetric water content, I first used the equation in the EC-5 documentation [20] that relates probe output voltage to apparent permittivity (ϵ_a). I then used the equations and coefficients in [73] to calculate volumetric water content from $\sqrt{\epsilon_a}$. Separate coefficients are provided in the documentation for

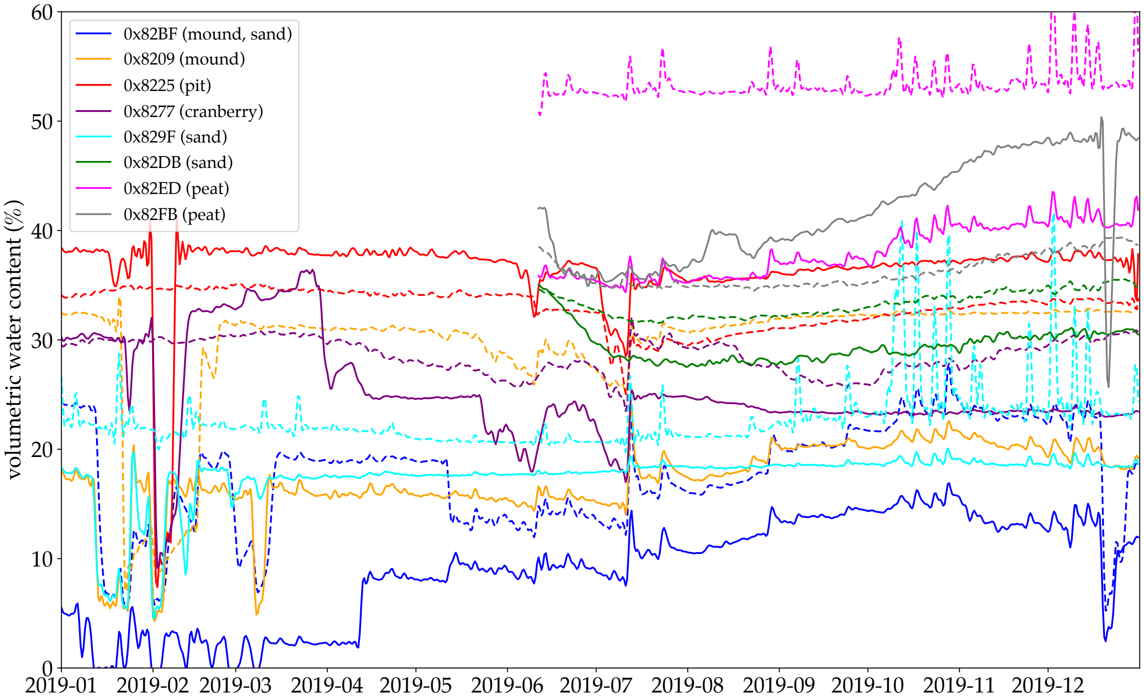

Figure 5.2: Soil moisture from selected sensors across Tidmarsh, 2019. Solid lines are probes at 5 cm depth; dashed lines are at 15 cm depth.

organic soil (with a high organic matter content) and mineral soil (containing sand, silt, and clay). I chose which coefficients to use for each probe based on the photographs of the soil taken at the time of installation.

Figure 5.2 shows a plot of many soil moisture sensors at both 5 and 15 cm depths from different locations, aspects of the microtopography, and types of soil. This shows that there is significant variation in soil moisture, addressing the first question about how soil moisture varies across the site and over time. This variation appears on a local scale rather than gradually changing from one end of the site to another. Sensors geographically near each other show substantial differences in both wetness and response to rainfall due to local variation in soil type and topography. Over time, there is a general drying trend for the first part of the year, and increasing wetness for the second half. In this 2019 dataset, this transition is particularly centered around a large storm that produced 11 cm[1] of rainfall over a 24-hour period.

The apparent sharp drops during the winter months are a result of freezing, which the moisture sensor measures similarly to dry soil. (This behavior is noted in the probe manufacturer's Frequently Asked Questions [23]).

The spikes on the plot are the soil moisture responding to rainfall events, and the extent of this response shows how water flows through the soil. High responsiveness, particularly when seen by both probe depths, suggests vertical water flow and connectivity. Lack of significant response to rainfall but long-term seasonal increase in moisture is indicative of lateral flow. The shape of these curves following the July storm event as the water drains and the soil dries down illustrates the degree to which water is retained by the soil. Figure 5.3 shows a closer view of the July storm, which followed a hot and dry period resulting in drier soil across the site. (This can be seen in many of the sensors at the center of Figure 5.2).

To understand how the soil type, morphology, and topography affect water flow, we can compare sensors in these different conditions and features.

Sensor 0x8277 (shown in purple) is installed in a patch where the surface was left undisturbed, and the vegetation remains as primarily cranberry. The cranberry substrate consists of numerous horizontal layers. Well-sorted glacial outwash sand was applied as part of the farming efforts, and creates layers interspersed with organic matter. This creates a highly anisotropic substrate that is significantly more permeable horizontally than in the vertical direction, with the water flowing readily within the horizontal sand layers. The sensors show little evidence of vertical connection between the 5 cm and 15 cm probes.

Sandy soil, on the other hand, shows significant vertical connection. Sensor 0x829F (shown in cyan) is installed in sand with little fine-grained or organic matter. This pure mineral soil is isotropic and water flows readily in all directions. Both the upper and lower probes respond readily to precipitation and

[1] As measured by the Foothills Preserve weather station about 1 km away.

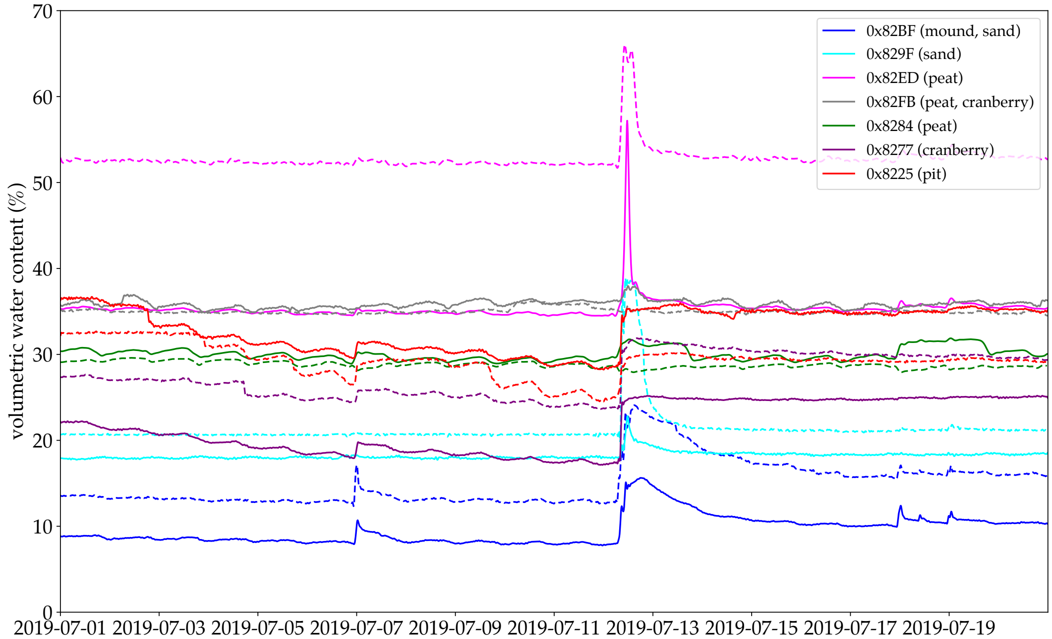

Figure 5.3: Soil moisture preceding and following a significant July storm. Solid lines are probes at 5 cm depth; dashed lines are at 15 cm depth.

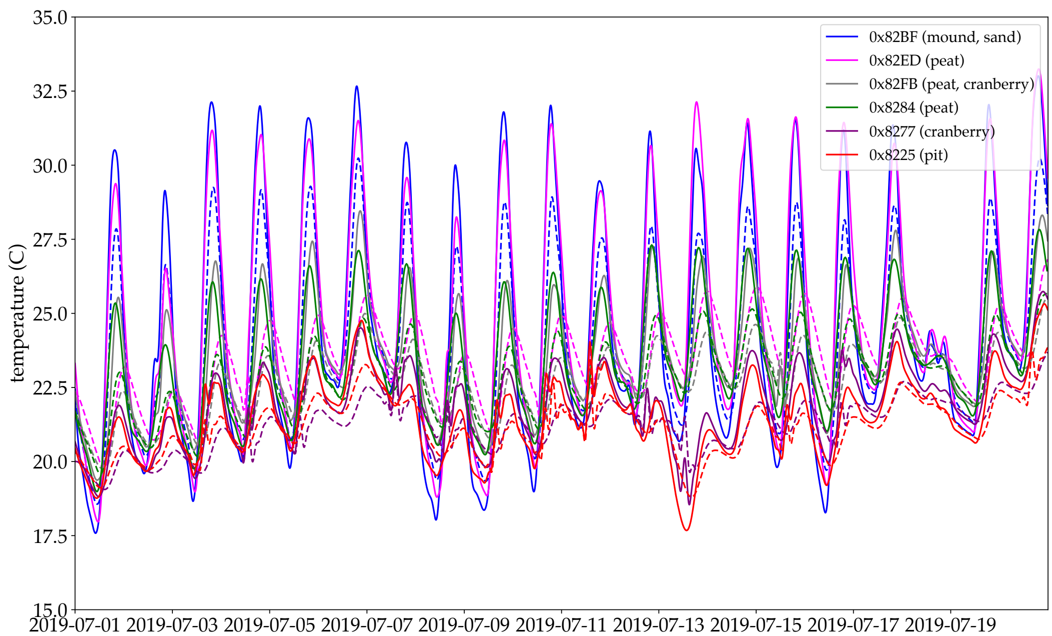

Figure 5.4: Soil temperature preceding and following a significant July storm. Solid lines are probes at 5 cm depth; dashed lines are at 15 cm depth.

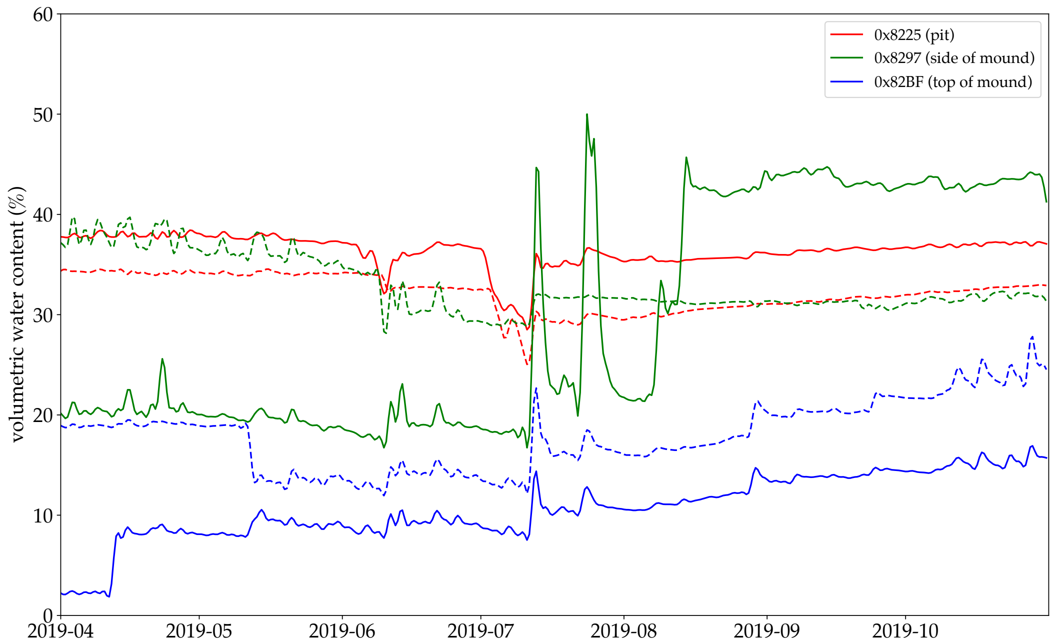

Figure 5.5: Soil moisture from different microtpoographical aspects. Solid lines are probes at 5 cm depth; dashed lines are at 15 cm depth.

dry out rapidly due to quick drainage. These heavily sandy areas are some of the driest on the site as they lack the means to retain moisture.

Sensors in peat show the slowest response and greatest capacity to retain moisture. Sensors 0x82FB (gray) and 0x82ED (magenta) are installed in peat-rich soil, and represent some of the wettest locations on the site. Changes in moisture occur very slowly; once the peat is wet it stays wet.

We can also compare moisture across features of the microtopography. Figure 5.5 shows data from three nearby sensor nodes. 0x82BF (blue) is installed at the top of a mound; 0x8297 (green) on the side of a different mound, and 0x8225 (red) is in a low spot/pit.

The top of the mound, which is composed primarily of sand scooped up from the layers placed by farming, shows a response similar to other sensors located in sand. Both probe depths represent some of the driest conditions measured, and are strongly responsive to rainfall, showing significant vertical flow.

The pit, in addition to being much wetter than the top of the mound, shows wetter conditions at the surface than at depth. This indicates a less permeable layer between the surface (5 cm) and the deeper probe (15 cm) that prevents water from percolating deeper, resulting in less than saturated soil even when there is ponding on the surface. (This finding is consistent with prior observations made with hand probes by Dr. Hatch's team). The lack of response to individual storm events also supports a lack of vertical flow. The side of the mound, intuitively, is in between. The deeper probe is very similar to the deeper probe in the pit, and shows a gradual drying down for the first half of the year and a gradual wetting over the second half. The 5 cm probe begins the year very similar to the sandy top of the mound, relatively dry and responsive to rainfall. Midway through the year it appears to switch to another mode, becoming saturated and staying wet.

These results begin to suggest some answers about the pit and mound topography. Constructing the pits and mounds does appear to break up the highly anisotropic cranberry substrate with primarily lateral flow left at the end of farming, allowing local vertical flow and creating areas both wetter (pits) and drier (mounds) than the undisturbed cranberry surface.

5.3 Conclusions

These experiments are a concrete example of how features of the sensor network can be utilized to design and execute experiments that test hypotheses. While these results are preliminary and would benefit from more replicates to account for the many variables across the pit and mound topology (which can be continued in future work), this experiment shows how the sensor network can detect real-world changes as a result of restoration.

In addition to the specific set of hypotheses included in the designed experiment, the long-term data from the sensor network can also lead to unexpected observations. As an example, Figure 5.3 shows significant diurnal cycling in the soil moisture, with decreases in moisture during the hottest hours of the day suggesting that plant transpiration is cyclically drawing down the moisture in the soil. After the storm, this is greatly reduced—suggesting enough moisture is present such that plant use no longer significantly affects the local moisture around the sensors. (Temperature also somewhat affects the response of the sensor and can account for some diurnal cycling in the data, but the adjacent temperature measurements shown in Figure 5.4 are not substantially different before and after the storm).

A final important result from this work is the demonstration of how the sensor network has facilitated collaboration between researchers with different backgrounds. The interface that I built for accessing the data and metadata from the experiment was instrumental in enabling me, a computer engineer with only a novice background in hydrology, to share the data (and metadata necessary for contextualizing the data) with expert environmental scientists and interpret the results.

Chapter 6

Experiences, Creative Expressions, and Public Outreach

The first key research question of this thesis concerns the extent to which a sensor network can serve the dual roles of being a research tool for furthering the practice of wetland restoration and a canvas for creative expressions, experiences, and tools that connect people to wetlands and wetland restoration. To answer the latter half of that question, this chapter catalogs the projects that make use of the sensor network and its data in one form or another.

While I played a part in the creation of some of these works, many are independent works by others, who are credited below the description of each project. The number of works and diversity in medium and background of the creators serve as an evaluation of the sensor network as a creative tool. As this chapter hopefully demonstrates, the sensor network has been successfully used in many different expressions that have been experienced by a wide audience through various installations, events, and online content and software.

The projects presented here are summarized at the end of the chapter in Table 6.1, which lists the specific inputs from the sensor network employed by each. This list is meant to be representative rather than comprehensive; not ev-

Figure 6.1: Doppelmarsh.

ery project that uses the sensor network to date is listed here, and we hope that more will follow as the network lives on in the years following the publication of this document.

6.1 Doppelmarsh: Tidmarsh in Cross-Reality

Building on Doppellab and our other previous experiments in cross-reality, Doppelmarsh brings the landscape of Tidmarsh into 3D game engine software and allows the user to freely explore it virtually as it is affected by real-time sensor data.

I began building Doppelmarsh (Figure 6.1) in 2013, in the early days of my involvement in the Tidmarsh project. Like Doppellab, I started with CAD data to construct the virtual world. Where Doppellab used 3D architectural drawings of the Media Lab buildings, Doppelmarsh used topographic maps, primarily based on high-resolution LIDAR scans from USGS. I converted the digital elevation map (DEM) data into heightmap textures that could be imported into the Unity game engine and used to create a virtual terrain. Once in Unity, the terrain required some post-processing to smooth out the quantization inherent in the DEM format. I then used Unity's terrain modeling features to paint on textures, grass, and trees to match the vegetation from the real Tidmarsh.

As sensor data was received from the network and decoded into measurements, it was piped into S. Russell's Chain-API, which became the real-time data source for Doppelmarsh and many other applications. We developed a Unity script called ChainSync, which retrieves the site summary from Chain-API and automatically populates the virtual world with representations of the sensor nodes that visually mimic their real-world counterparts, with text floating in the space above them showing the sensor readings. After the environment is initialized, ChainSync then connects to Chain-API's WebSocket feed and updates the virtual environment as new measurements arrive.

Doppelmarsh has continued to grow and evolve, with many sub-projects building on the basic framework of the cross-reality environment. S. Russell added a sonification of the sensor data, creating a textural soundtrack driven by temperature and humidity using samples of Tibetan singing bowls, and transient ukulele plucks (whose pitch vary with temperature) when a new measurement is received from each sensor. The soundtrack is spatialized so that the background textures are driven by the sensors closest to the user's position in the virtual world, and the ukulele plucks sound like they are coming from the sensors that triggered them. This work would later be extended into a full framework for building data-driven compositions (See *SensorChimes*, section 6.3.1).

Real-time audio also found its way into the virtual environment. I built *StreamChannel*, a native Unity plugin (with builds targeting Mac, Windows, Linux, and iOS) that receives SHOUTcast audio streams in Ogg Opus or Ogg Vorbis formats and outputs PCM audio data via Unity AudioSource objects that can either be mixed into the background soundtrack or attached to game objects and spatialized. Each channel from the input stream is mapped onto a separate AudioSource object and thus can be individually spatialized. We used this plugin to bring spatialized audio from the microphones at Tidmarsh into the virtual environment.

Figure 6.2: Sensor Creatures in Synthetic Menagerie.

D.D. Haddad contributed significant effort to improving the visuals in Doppelmarsh, and further integrated the effects of the sensor data by using measured wind speed to control the motion of the grass and trees and weather station data to render rain, snow, fog, and other weather phenomena [27].

From the end of 2015 into 2016, Tidmarsh underwent the construction phase of its restoration, which transformed the physical landscape by removing berms, digging new channels, and reshaping the topography. Since we did not have new LIDAR scans reflecting these changes, we made the same changes to the virtual landscape that the earth-moving machines were making to the physical terrain, using the construction engineering plans as a reference. We updated the vegetation in the virtual environment to match the rapid transformation after the construction equipment left and the natural seed bank in the soil leapt into action.

Project credits: Brian Mayton, Gershon Dublon, Don Derek Haddad, Spencer Russell.

6.1.1 Synthetic Menagerie

Many of the visualizations in Doppelmarsh are concrete: either literal numeric displays of sensor data or virtual renderings of physical phenomena. But as Doppelmarsh is a virtual environment, this doesn't have to be the case—representations of data can take on more abstract forms. *Synthetic Menagerie* [26], by D.D. Haddad, populates the landscape with virtual creatures (Figure 6.2). These are purely virtual: a deer in the virtual world does not represent a deer at Tidmarsh. But these virtual creatures *do* respond to sensor data from the physical world by changing their behavior and appearance.

Early versions of *Synthetic Menagerie* experimented with realistic renderings of animals. Some, like deer, would not be out of place at Tidmarsh. As such, they fit in naturally to the virtual environment. However, in a rendering where many physical phenomena are represented literally, a realistic rendering of an animal may cause users to believe that a deer has actually been sensed at Tidmarsh and is present at that location. D.D.H. also experimented with creatures that would be obviously out of place at Tidmarsh (e.g. a giraffe) or anywhere in the natural world (such as Pokémon-esque sprites).

In its later iterations, *Synthetic Menagerie* settled on real and locale-appropriate animals (deer), but changed the rendering to abstract and ethereal representations. The deer are recognizable but are drawn as clouds of light. Their virtual "fur" can change length and become "frizzy" to indicate temperature and humidity conditions.

D.D.H. also envisioned (but as of this writing has not yet implemented) a way that a user might form lasting relationships with these creatures—"collecting" them as virtual pets that "feed" on sensor data. A user could curate their collection of creatures, being able to interact with them from anywhere through Doppelmarsh, and through augmented reality (*à la* Pokémon Go) when physically visiting Tidmarsh.

Project credits: Don Derek Haddad.

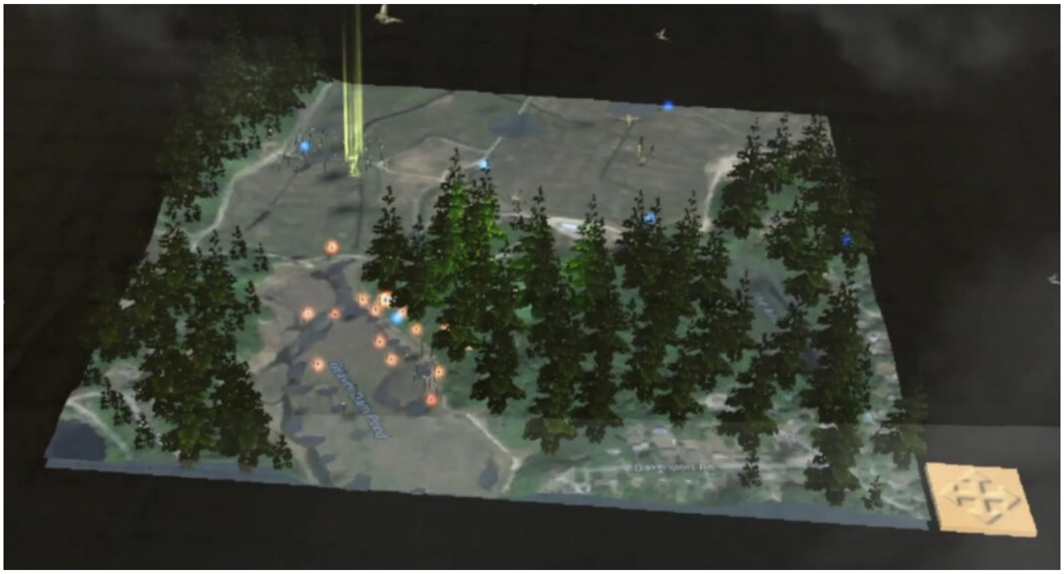

Figure 6.3: Hakoniwa as seen on a tabletop through the Hololens headset.

6.1.2 Hakoniwa

Doppelmarsh was originally developed for a desktop computer, with the user experiencing the virtual world in first-person perspective through a keyboard and monitor, and later also through low-cost virtual reality (VR) headsets. Visual augmented reality (AR) displays, like the Microsoft Hololens, allow virtual objects to be placed in the wearer's immediate environment. *Hakoniwa*, from the Japanese 箱庭 (lit. *box garden*), places a portion of Tidmarsh on a tabletop as if it were a miniature boxed garden (Figure 6.3). This top-down view allows the user to physically walk around the table and view the miniature from any angle. The user may examine the model more closely simply by moving their head closer.

Hakoniwa is also an auditory experience. The live microphones are depicted in the miniature, and the user hears the sound coming from them. At a distance, the sound is a mix of all of the microphones. As the user moves their head in closer, they begin to hear only the microphones near to their focal point. In addition to the live streams, static recordings of interviews and stories [6] are attached to points of interest. The user navigates these in the same manner as the live sounds, by moving in closer and focusing on them.

Project credits: Gershon Dublon, Spencer Russell, Halsey Burgund, Brian Mayton.

6.1.3 Baguamarsh

In *Baguamarsh* [34], Fred Jiang combines environmental data from the Tidmarsh sensor network with 360° panoramic photography, ambisonic recording, and personal biometric data (ECG and EEG) to capture a visitor's subjective experience. The user records a short clip with a GoPro 360 camera, which also captures ambisonic audio. Concurrently, an Apple Watch and and Muse BioHarness record EEG and ECG data. The sensor network is queried via Chain API to capture data from nearby sensor nodes, and weather conditions are fetched from AccuWeather.

After these data are captured, the user can later re-visit the experience through an immersive VR environment based on the Eight Trigrams, or *Bagua*, of the *I Ching* (the Chinese *Book of Changes*). Each trigram corresponds to a physical phenomenon in nature, and in *Baguamarsh* is mapped to an appropriate type of sensor. The user can toggle between the real environment (reconstructed from the panoramic photos) and an abstract environment that represents the sensor data. In the abstract world, the user navigates the data space through the Bagua menu (Figure 6.4a). Different visualizations are used for each sensor, such as lines of varying width on the ground to represent soil moisture (Figure 6.4b) or spheres that vary in density to represent humidity (Figure 6.4c).

Project Credits: Fred Jiang

(a) Bagua menu

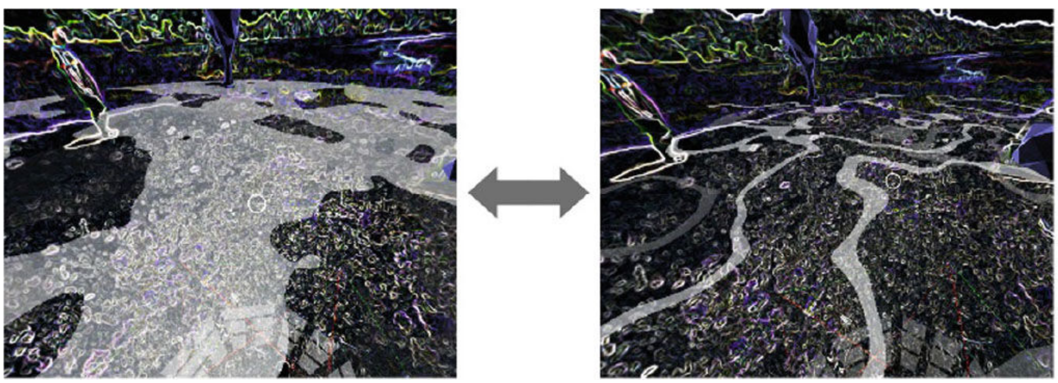

(b) Soil moisture is represented by lines on the ground that vary in thickness

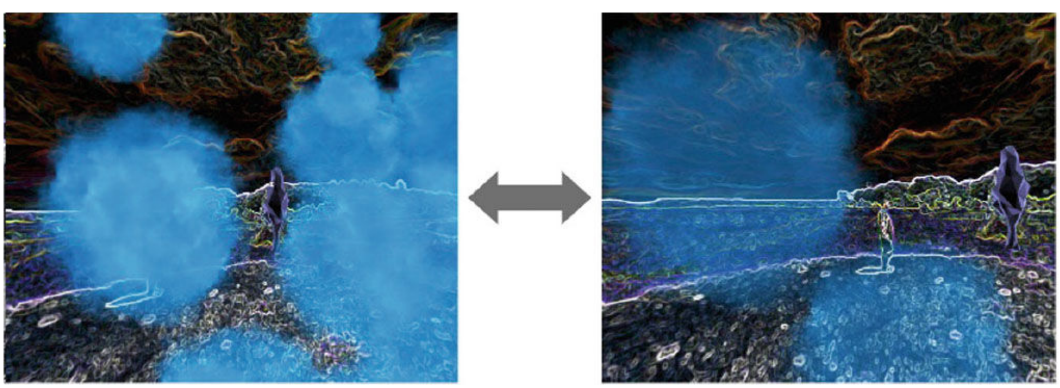

(c) Humidity is represented by the density of blue spheres

Figure 6.4: Elements of the Baguamarsh VR interface.

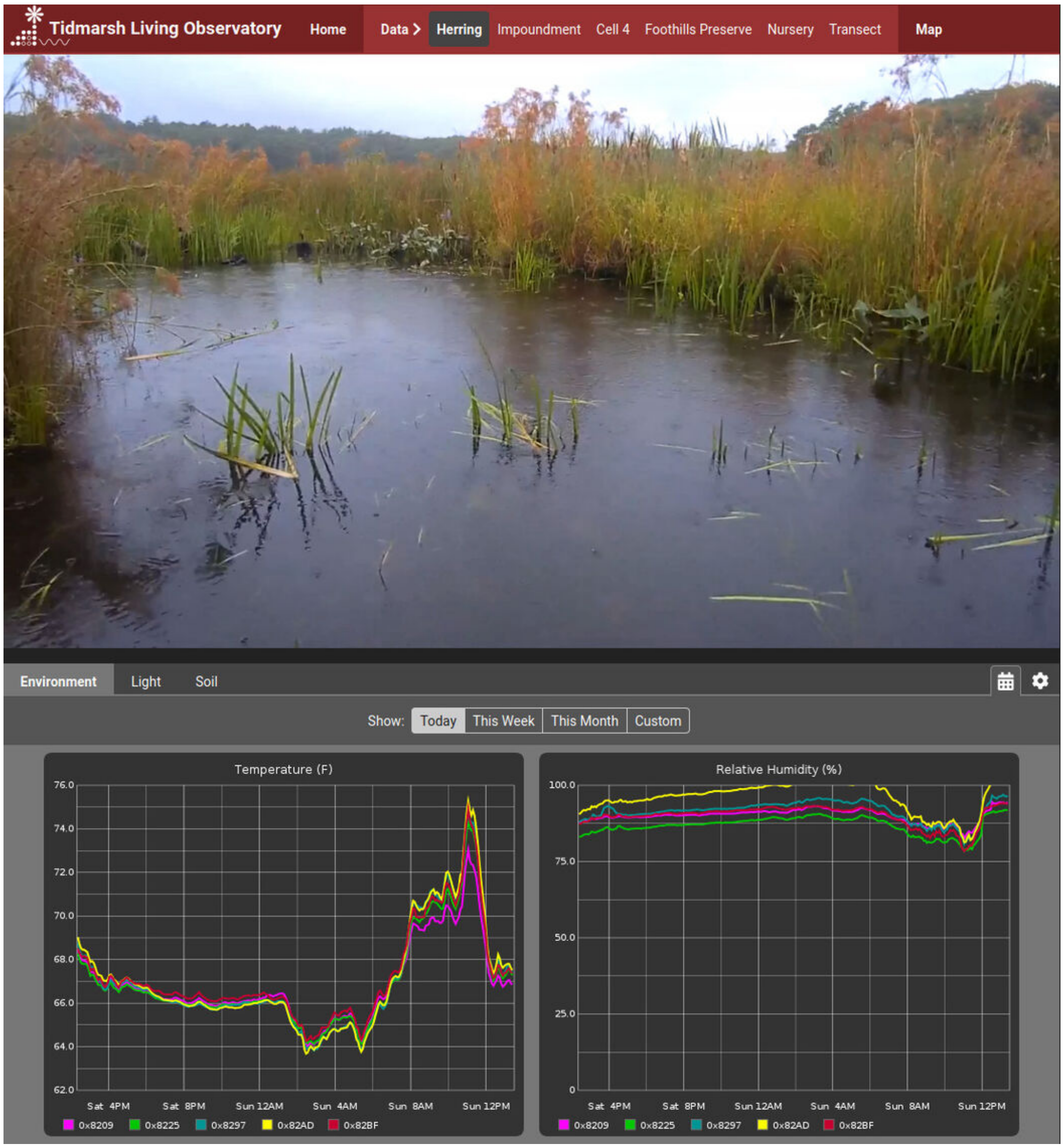

Figure 6.5: The tidmarsh.media.mit.edu website.

6.2 Web Interfaces

6.2.1 tidmarsh.media.mit.edu

The project website at https://tidmarsh.media.mit.edu (Figure 6.5) is perhaps the most concrete and straightforward representation of the data collected by (and metadata about) the sensor network. In addition to a short introduction and links to publications and project videos, the website provides visitors with access to both live and archived sensor data.

The website is organized into "sensor sites," grouping data around clusters of sensor nodes (and in some cases, particular experiments) at Tidmarsh. Each site shows plots of environmental data from the low-power sensor nodes, live video and images from nearby cameras, and live audio feeds. The sensor data plots by default show the most recent day's worth of data, but the user may also view the last week or month, or choose an arbitrary range of dates to plot from the database.

Each site shows plots for different sensors, depending on the capabilities of the hardware at that location and the nature of any experiments being conducted. Some sites, for example, include soil moisture and temperature data. A few sites also show data about the sensor network itself, such as the state-of-charge of batteries and available energy from solar panels.

Since much of the data and context for the project has a strong geospatial component, the web interface also provides a map view. This is implemented using "slippy" tiles (familiar to users of Google Maps, though here implemented with Leaflet.js), allowing visitors to pan and zoom by dragging with a mouse or touchscreen.

The base imagery is selectable from a number of public domain aerial imagery datasets from multiple different years. For some datasets, both visible light and infrared views are available. The primary sources of this data are the USDA NAIP program [52] and USGS High Resolution Orthoimagery [74] but

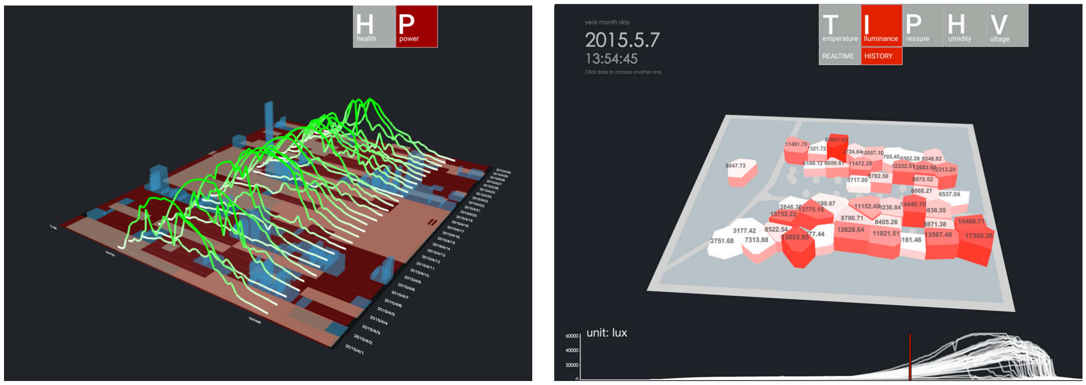

Figure 6.6: MarshVis displaying base station power (left) and illuminance across the sensor network (right).

the system is set up to allow other data sources to be easily added as well. Public domain datasets are used so that related projects and works can freely use map data hosted on the Tidmarsh server without negotiating licenses or paying fees.

Additional imagery collected by Living Observatory collaborators may also be overlaid on portions of the site. This includes very high resolution drone imagery captured by Inter-Fluve Inc. during the winter following the restoration, and infrared thermal imaging of portions of the stream channel from USGS [29].

Vector GIS data layers can also be shown on the map, such as subsurface data collected with ground-penetrating radar [28] and live markers for sensor nodes that allow visitors to select a sensor metric to be plotted on the map at the locations of the sensors.

Project credits: Brian Mayton.

6.2.2 MarshVis

Another take on visualizing the sensor data in a web browser, titled *MarshVis* [40], was created by animator and graphic artist Nick Lee. It presents sensor data in a number of different views.

Several visualizations represent the area densely covered by sensors during the time that Lee was working on the project in 2015. Since the sensors were not installed on a uniform grid, he chose to represent the landscape as a Voronoi tesselation with a sensor node at each centroid. The cells in the tesselation change color and vary in height to represent various sensor metrics.

Another view focuses on solar charging for one of the base stations. 2D plots of battery charging current are shown at an angle in 3D space, allowing multiple days of data to be shown along the third axis. Underneath each plot, ambient light levels from across the network are rendered as histograms, clearly showing the correlation between bright sunny days and rapid battery charging. The traces on the plots occasionally disappear when the base station went offline for lack of power.

Project credits: Nick Lee.

6.2.3 livingsounds.earth

At *livingsounds.earth* [15], Gershon Dublon encourages visitors to the website to *listen*. Following a short introduction, visitors hear a stereo mix of the live audio from the network of microphones at Tidmarsh. Demonstrating how the audio can really stand on its own, the only visuals are abstract and simple: Nan Zhao's pixel-art birds flitting from tree-to-tree and Orcun Gogus's rippling wavy lines.

While the live sounds of Tidmarsh are normally the sole performer on this 24/7 radio program, Dublon has created an interface that enables invited guests to connect and perform alongside the marsh. This could take the form of a spoken interview, with researchers discussing their work and telling stories about their time at Tidmarsh, or a musician improvising in real time along with the sounds of the bog.

Figure 6.7: Part of a composition using the SensorChimes interface inside the Max/MSP graphical programming environment.

Project credits: Gershon Dublon.

6.3 Auditory Experiences

6.3.1 SensorChimes

SensorChimes [45] by Evan Lynch is a project that further enables others to build creative works using data from the sensor network. Inspired by Russell's original sonification of sensor data in Doppelmarsh, Lynch sought to enable non-programmers to build their own musical compositions. *SensorChimes* is a set of plugins for the Max/MSP environment, which is a graphical programming language often used by electronic musicians.

These plugins (or *externals* in Max parlance) provide new objects that can query the sensor network for both live and historical sensor data (Figure 6.7), as well as metadata (such as the locations of sensors). By removing the technical friction of querying the database, Lynch enables composers to focus on the music and build data-driven compositions in a familiar environment.

SensorChimes can communicate with a running instance of Doppelmarsh using the Open Sound Control (OSC) protocol, which enables the resulting compositions to be spatialized based on the avatar's current position in the virtual environment.

S. Russell later adapted *SensorChimes* to Pure Data, which is an open source environment very similar to Max/MSP. Compositions translated to the Pure Data environment can be executed within `libpd`, which encapsulates the DSP core of Pure Data within an embeddable library. Russell developed a Unity extension that wraps `libpd` and enables compositions to be loaded directly into Doppelmarsh without the need for running a separate copy of Max/MSP.

Lynch contributed his own data-driven composition, and commissioned pieces from composers Evan Ziporyn and Ricky Graham. All of these pieces, including Russell's original sonification, have been integrated into the Doppelmarsh experience.

Project credits: Evan Lynch, Spencer Russell.

6.3.2 HearThere

Many of the projects described in this chapter are experiences for *remote* visitors—presented through computers and wearable displays or physical installations distant from Tidmarsh. *HearThere* [16, 63] is a device for *on-site* Tidmarsh visitors, enhancing the experience of being there. It is a wearable device (Figure 6.8) that makes use of bone-conduction transducers on the wearer's temples to augment the wearer's sense of sound without blocking natural sounds from reaching their ears. *HearThere* tracks the wearer's location (via GPS) and head orientation (with a 9-DoF inertial measurement unit) and plays spatialized audio content based on where they are currently looking.

Most of the audio content comes from the network of microphones installed in the former impoundment and adjacent forest. This can be heard live (with a short buffer delay). While some sounds would be heard directly by an on-site

Figure 6.8: Glorianna Davenport wearing the HearThere device.

listener, *HearThere* offers a way to listen in on quiet or distant sounds picked up by the microphones far away from the paths, where it is not possible for a visitor to easily walk. By looking out into the marsh, the listener can hear sounds from the microphones in front of them. As *HearThere* detects the listener standing still and focused in a particular direction, their augmented hearing also narrows and focuses in. This effect can also be accentuated or briefly muted by touching the sides of the device, similar to how a listener can cup their hand in front of their ear or plug their ears to alter their hearing.

Since all of the Tidmarsh audio is recorded and stored, *HearThere* can also be used to listen to a different season or time of year. The spatialization still works the same way: the wearer hears the audio that they would be hearing from their present location and focal direction at the time the recording was made.

Listeners can also choose to hear one of the *SensorChimes* sonifications of the sensor data spatialized to their location, either live or from any time recorded in the database.

The interface of *HearThere* is intentionally minimal and the effects are subtle so as not to detract from the experience of being present in nature by putting an obvious piece of technology in front of the user's face. While there is a smartphone-based interface to set the configuration of the device, the user only interacts with this when they put on the device or if they want to change to a different time period or sonification. Most of the time, the phone remains in the listener's pocket and they interact with it by simply walking around and listening.

Project credits: Gershon Dublon, Spencer Russell.

6.3.3 ListenTree

Unlike the installation works described later in this chapter, *ListenTree* [57, 16] is completely invisible. Dublon and Portocarrero turn trees into audio-haptic playback devices by introducing vibrations through transducers buried under the soil. A person passing by such a tree might hear whispers of sound or feel slight vibrations coming from the ground. Upon leaning in and placing their head against the tree, the sound becomes clear as it is conducted from the transducer, through the tree, and into the bones in the listener's head, finally exciting the sensitive structures within the ear.

Two early installations of ListenTree (one outside the Media Lab and one at the MIT Museum) utilized live audio from Tidmarsh, connecting listeners in the city to the sounds experienced by trees in a wetland.

Project credits: Gershon Dublon, Edwina Portocarrero.

6.4 Tracking the Arc of Change

Tracking the Arc of Change [47] is a short film that tells the story of the Tidmarsh restoration through timelapse, video, and sound. It begins with a timelapse clip looking out over Tidmarsh in the fall of 2015 before construction began. The timelapse accelerates and shows the earth moving equipment arrive, remove the berms, add logs and other dead wood, and reshape the stream channel. Winter turns into spring and the muddy terrain begins to grow.

The view transitions to the herring camera, looking out across the water. Seasons pass, the vegetation grows and dies back, the water rises and fall as storms come and go, and the water freezes and melts. The speed of the timelapse changes to show changes and cycles happening at different timescales. Occasionally, the timelapse slows to normal video to show macrofauna coming to visit the camera: a deer jumping over the water, a kingfisher making a catch, and a blue heron pecking at the camera.

The timelapse is set to a soundtrack of recordings from the microphones nearby. The soundtrack plays at normal speed, but changes to reflect the seasons and time periods represented by the timelapse.

The view changes to look out across the marsh, and a late summer day passes by. The next morning, a storm rolls through, and in a dramatic flash of lightning, the view goes black. (This lightning strike damaged some of the recording equipment).

The timelapse continues from the herring camera, showing more seasons pass by. The sediment, visible at the bottom of the stream channel under the surface of the water, migrates in serpentine patterns.

The timelapse slows and proceeds into night; the view goes black but for the reflection of the moon transiting the sky, and the focus shifts to the sounds of the night.

Gradually, the sun rises, and winding tendrils of fog clear away from the ground as the day begins.

Late spring arrives and the view once again slows to normal speed, showing a school of river herring running the channel.

The sequence concludes as the sun sets, leaving the viewer to listen to the

Figure 6.9: Nan Zhao listening to the audio inside the *Moss Listening* installation.

sounds of the night before finally fading away.

Project credits: Brian Mayton.

6.5 Physical Installations

6.5.1 Moss Listening

Moss Listening was an installation piece that focused on the sounds of Tidmarsh. It took place in the early days of the project and was the first use of the live audio streams from the experimental setup at the Arm. The installation was an approximately 3 by 4 meter space constructed from black curtains in a corner of the Responsive Environments lab. We filled the enclosure with freshly collected organic matter from Tidmarsh, including peaty soil, grasses, stumps, and moss[1]. A humidifier pumped moist air into the space, creating wisps of fog along the perimeter.

Behind the plants and curtains, we hid an 8-channel speaker system, placing

[1] And, unintentionally, a stowaway toad.

the speakers at the vertices of the space. The speakers played audio from the microphones at Tidmarsh, spatialized using ambisonics. The audio from four microphones was panned around the sides of the space, and sounds from hydrophones were panned to the bottom.

As part of a spring 2013 event at the Media Lab titled *The Other Festival*, visitors were encouraged to enter the installation, where they could wander about or sit on a stump and experience the sounds of Tidmarsh, spatially rendered so as to convincingly reproduce the sonic experience of actually being there.

Project credits: Gershon Dublon, Brian Mayton.

6.5.2 Tidmarsh Living Observatory Portal

In contrast to *Moss Listening*, in which the structure was meant to be minimally visible, V. Sumini's *Tidmarsh Living Observatory Portal* [69] draws attention to its form. In this take on an enclosed telepresent Tidmarsh experience, a participant would find an egg-shaped structure located in an alcove of a building. The structure consists of a 3D-printed plastic frame and translucent fiberglass side panels. Projectors surrounding the pavilion display visuals from Tidmarsh onto the outside of the panels so that they appear rear-projected to a person sitting inside. Speakers around the egg play multi-channel audio and sonifications.

One can envision such pavilions as architectural features of a building, permitting the occupants to step out of their day-to-day activities and into Tidmarsh (and other environments) for a few moments.

Project credits: Valentina Sumini.

6.5.3 Virtual Fish Count

The spring of 2020 brought some unexpected challenges with the COVID-19 pandemic. As the world entered lockdown, Mass Audubon made the decision to close Tidmarsh to visitors and researchers. This disrupted the annual fish count, in which volunteers would normally take shift on site observing the stream at several locations to count river herring as they make their way back upstream to spawn.

Using the live video stream from the herring camera, we were able to successfully transition the 2020 herring count to an online volunteer activity. I developed an embeddable video stream player that shows the live camera view alongside live readouts of temperature sensors in the air near the stream channel and in the water. Robert Kearns authored introductory materials, created a page on Mass Audubon's Tidmarsh Wildlife Sanctuary website [37], and coordinated the virtual count. Volunteers could log in to the site, watch the stream for a 10-minute interval, and submit their count and observations through a form.

At the end of the fish counting season, volunteers had logged over 50 hours of counts and over 500 fish. The project concluded with an online celebration at which talks were given about river herring and the sensor network.

As I have presented this work over the summer of 2020, I've received many questions about why we didn't use polarizing filters to cut glare, or better position the camera to see through the surface of the water. While these are fantastic ideas for next year, and are certainly suggestions I would have implemented given the opportunity, I emphasize that we had no access to the property to make changes to the setup after the lockdown had started. While the arrangement we had was not what I ideally would have designed for conducting a fish count, it demonstrates how we were able to adapt part of the sensor network on short notice to design an experiment and experience for online volunteers. This underscores the value of a general-purpose and flexible network in permitting uses that were not envisioned when it was designed and installed.

Mass Audubon has expressed interest in continuing to use the live video streams, extending observations to other types of wildlife as well as running online fish counts over the coming years in parallel with the on-site count. We hope that this is also a step toward more significant collaborations with Mass Audubon and finding ways that we can use the sensor network in educational programs.

Project credits: Robert Kearns, Lauren Kras, Brian Mayton.

Table 6.1: Summary of creative works, experiences, and outreach projects

Project Title	**Media**	**Inputs**
Doppelmarsh	Virtual/cross-reality	environmental sensor data, live audio, camera images
Synthetic Menagerie	Virtual/cross-reality	environmental sensor data
Hakoniwa	Augmented reality, sound art	environmental sensor data, live audio
Tidmarsh website	Interactive web site	environmental sensor data, live audio, live video, GIS data
MarshVis	Interactive web visualization	environmental sensor data
livingsounds.earth	Web site, performance art	live audio
Sensor Chimes	Compositional framework, data-driven musical compositions	environmental sensor data
HearThere	Auditory augmented reality	live and historic audio, environmental sensor data
ListenTree	Installation art	live audio
Tracking the Arc of Change	Film	timelapse images, recorded video, recorded audio
Moss Listening	Installation art	live audio
Tidmarsh Living Observatory Portal	Installation art	environmental sensor data, live audio, camera images
Virtual Fish Count	Web site, volunteer activity	environmental sensor data, live video

Chapter 7

Conclusions and Future Work

When I began designing and building the Tidmarsh sensor network, experiences like Doppelmarsh and Moss Listening were the end goal, and set the initial requirements for the sensor network. *Real-time* data supported cross-reality telepresence, and *continuously recorded* data permitted time travel in the virtual world. We weren't yet aware of quite how far it would go, but we had an eye toward *long-term* installations to capture the restoration. *High bandwidth* audio streams provided an immersive experience, both directly on their own and within Doppelmarsh. Areas of *high density* sensing provided coverage at the resolution that would make a cross-reality browser interesting.

The hardware requirements for the first-generation sensor node were in part dictated by these early applications and in part by the environment. The real-time requirement influenced the network protocols and topology. Achieving density required a low-cost design that was practical to install and configure so that large numbers of sensors could be managed by a small team of researchers. The outdoor environment and potential for long-term monitoring required a durable node to withstand the elements and low power consumption to operate off the grid for years without maintenance.

As the network developed, its potential users expanded beyond ourselves. *Scientists* became interested in using the network to test hypotheses. *Artists* wanted to use the network to produce creative works. *Site managers* wanted to use the network to monitor conditions and make informed decisions. And *educators* wanted to use the network as a tool for teaching and learning.

Each of these groups of stakeholders has different requirements and uses the network in a different way. In some cases, my original design for the network was sufficient and existing systems found new purposes. For example, the network of microphones that we installed to capture and enjoy the sounds of nature has also become a tool for acoustic ecology and a testbed for new work in localization and soundscape resynthesis.

In other cases, I added features and requirements to the network to support new users. Using the network for scientific studies like our microtopography study required an *extensible* sensor node that could be customized around an experimental design, allowing the addition of sensors to monitor additional parameters. This led to a new design for the sensor node, and resulted in expanded capabilities for the network.

Each group of users may have a unique set of requirements, yet there isn't a single feature of the Tidmarsh network that has not been used by many. Every resource added to the network to support one application adds value for others: soil moisture sensors for studying microtopography can also make music or send alerts when it's time to water saplings in a greenhouse.

This work is not groundbreaking in being a wireless sensor network; my improvements there have been incremental. Commercial off-the-shelf nodes are now more readily available and can monitor many of the same parameters, though perhaps at higher cost and with less flexibility. *It is instead in bringing together all of these inputs and sharing them between a diverse set of applications and users that this work is unique.* To that end, the overarching goal of this work at Tidmarsh has been the development of a general-purpose, broadly capable sensor network that supports exploration and trying new prototype applications and uses.

The resulting network at Tidmarsh is extensive. The decision to ship as much data as possible to servers at the Media Lab in real-time requires significant infrastructure that provides the bandwidth and power. But making real-time data immediately available alongside the full recorded archive means entirely new applications can be prototyped without the need for changing or extending the network. All of the data are readily accessible.

In contrast, pushing details of the applications into the network itself permits more optimization, but places constraints on which data can be accessed and when. For example, tracking the return of bird species after a restoration could be done with an offline audio recorder and periodic collection and analysis of the recordings, but this would preclude an application like HearThere that makes heavy use real-time audio.

In the near future, many retired cranberry farms will undergo restoration projects through the Massachusetts Division of Ecological Restoration's Cranberry Bog Program [11]. These present an exciting opportunity to expand the work that began at Tidmarsh to a whole network of sites, creating a new *distributed* biological field station for cranberry bog restorations.

It is likely that not every site will have the resources to include the full range of capabilities provided by the network at Tidmarsh, and decisions will need to be made about which parts to include and how to manage the flow of data from sensors to applications. Here, the full network at Tidmarsh can serve as a testbed, giving the flexibility needed to prototype new ideas, which can then drive the decisions about what to include at other sites.

One of the big challenges in building a new application or experience is discovering what the network can do and what data are available. A few of our tools take steps in that direction: Chain API can associate related resources with links that can be programmatically followed. The tidmarsh.media.mit.edu website is a more human-friendly tool for browsing the network; it presents data using direct representations alongside geographical information and other metadata. To encourage more works to be created using data from the sensor network, a next step could be the development of a proper *toolkit* that makes it easier to explore the available resources and access the data through a single common interface.

This work has been a journey. Starting with a simple demo and sensor node, I've designed, built, and tested a large sensor network with hundreds of nodes, iterated and improved upon the many hardware designs and infrastructure systems that enabled and supported it, and kept it running for many years.

The sensor network, in addition to enabling the applications and experiences described in this document, has also played a central role in two Ph.D. theses [16, 65] in addition to this one and at least four masters' theses [45, 26, 61, 59].

The data collected by the network will be a lasting record from which we can continue to learn. It has played a part in the ecological understanding at Tidmarsh, and is helping to tell the story of complex ecological change. The lessons learned from Tidmarsh and its sensors will go on not only to improve future restorations, but can help us rethink the way we treat our entire planet.

Beyond Tidmarsh, this work demonstrates how heterogeneous sensor networks and rich streaming media can be a powerful tool for learning and communicating about ecological change across temporal and geographic scales. I close with the hope that this thesis inspires others to continue this work, and that sensor networks like this one are new way that we will experience, explore, and learn about our earth and beyond.

Appendix A

Atlas

All of the work done at Tidmarsh inherently has a geospatial component, and as a place in transition, Tidmarsh has changed significantly over the past several years. As not all readers will be intimately familiar with the locations described in this document (particularly the terminology we have used to refer to them) and the way that the geography has changed over time, this appendix is provided. It contains various maps depicting both pre- and post-restoration Tidmarsh, as well as various detail views and overlays showing the sensor and network installations that we have added and how these have changed over time.

The 2013 maps have been digitized primarily from the 2013 USGS High Resolution Orthoimagery [74]. The 2018 map was primarily digitized from the 2018 NAIP [52] dataset, with other references such as Inter-Fluve's engineering plans, particularly for the northern sections of the channel that are not clearly visible through the foliage in the NAIP images.

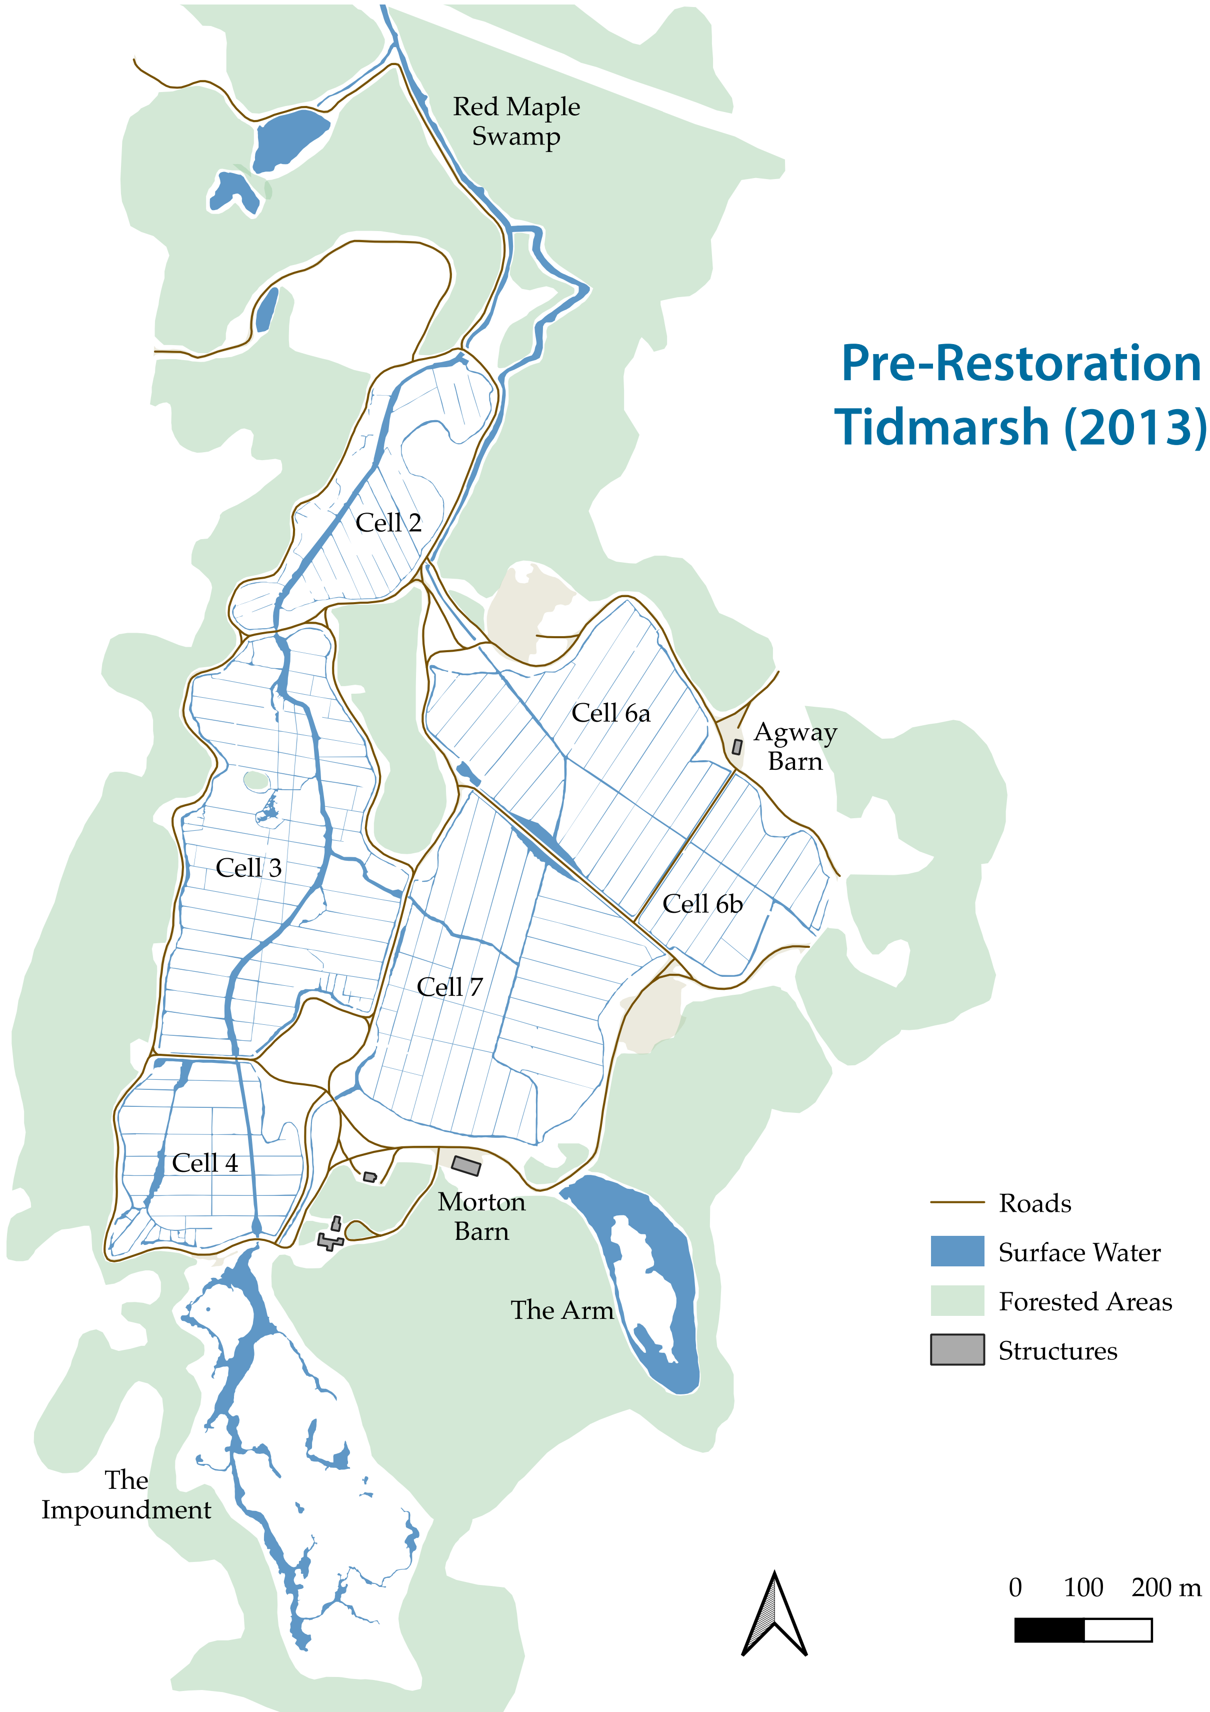

Figure A.1: Pre-Restoration Tidmarsh (2013)

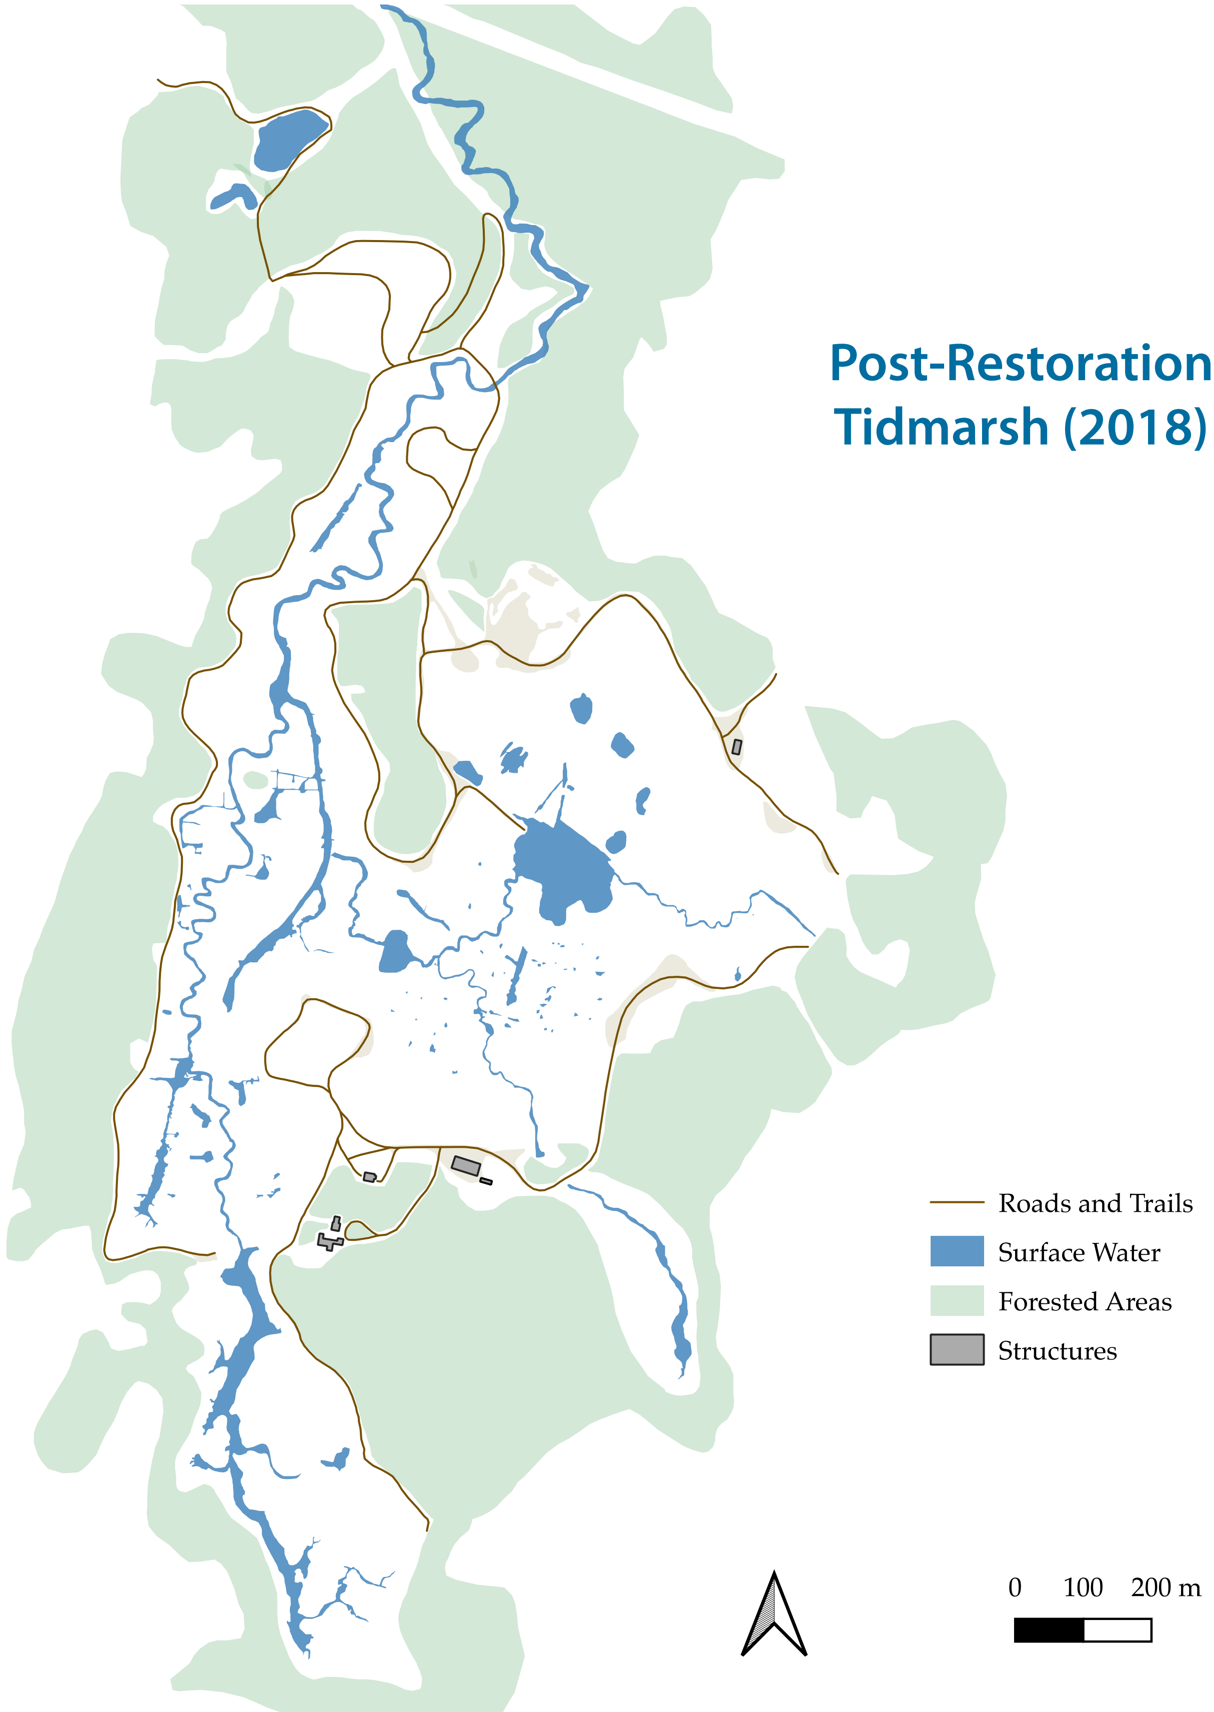

Figure A.2: Post-Restoration Tidmarsh (2018)

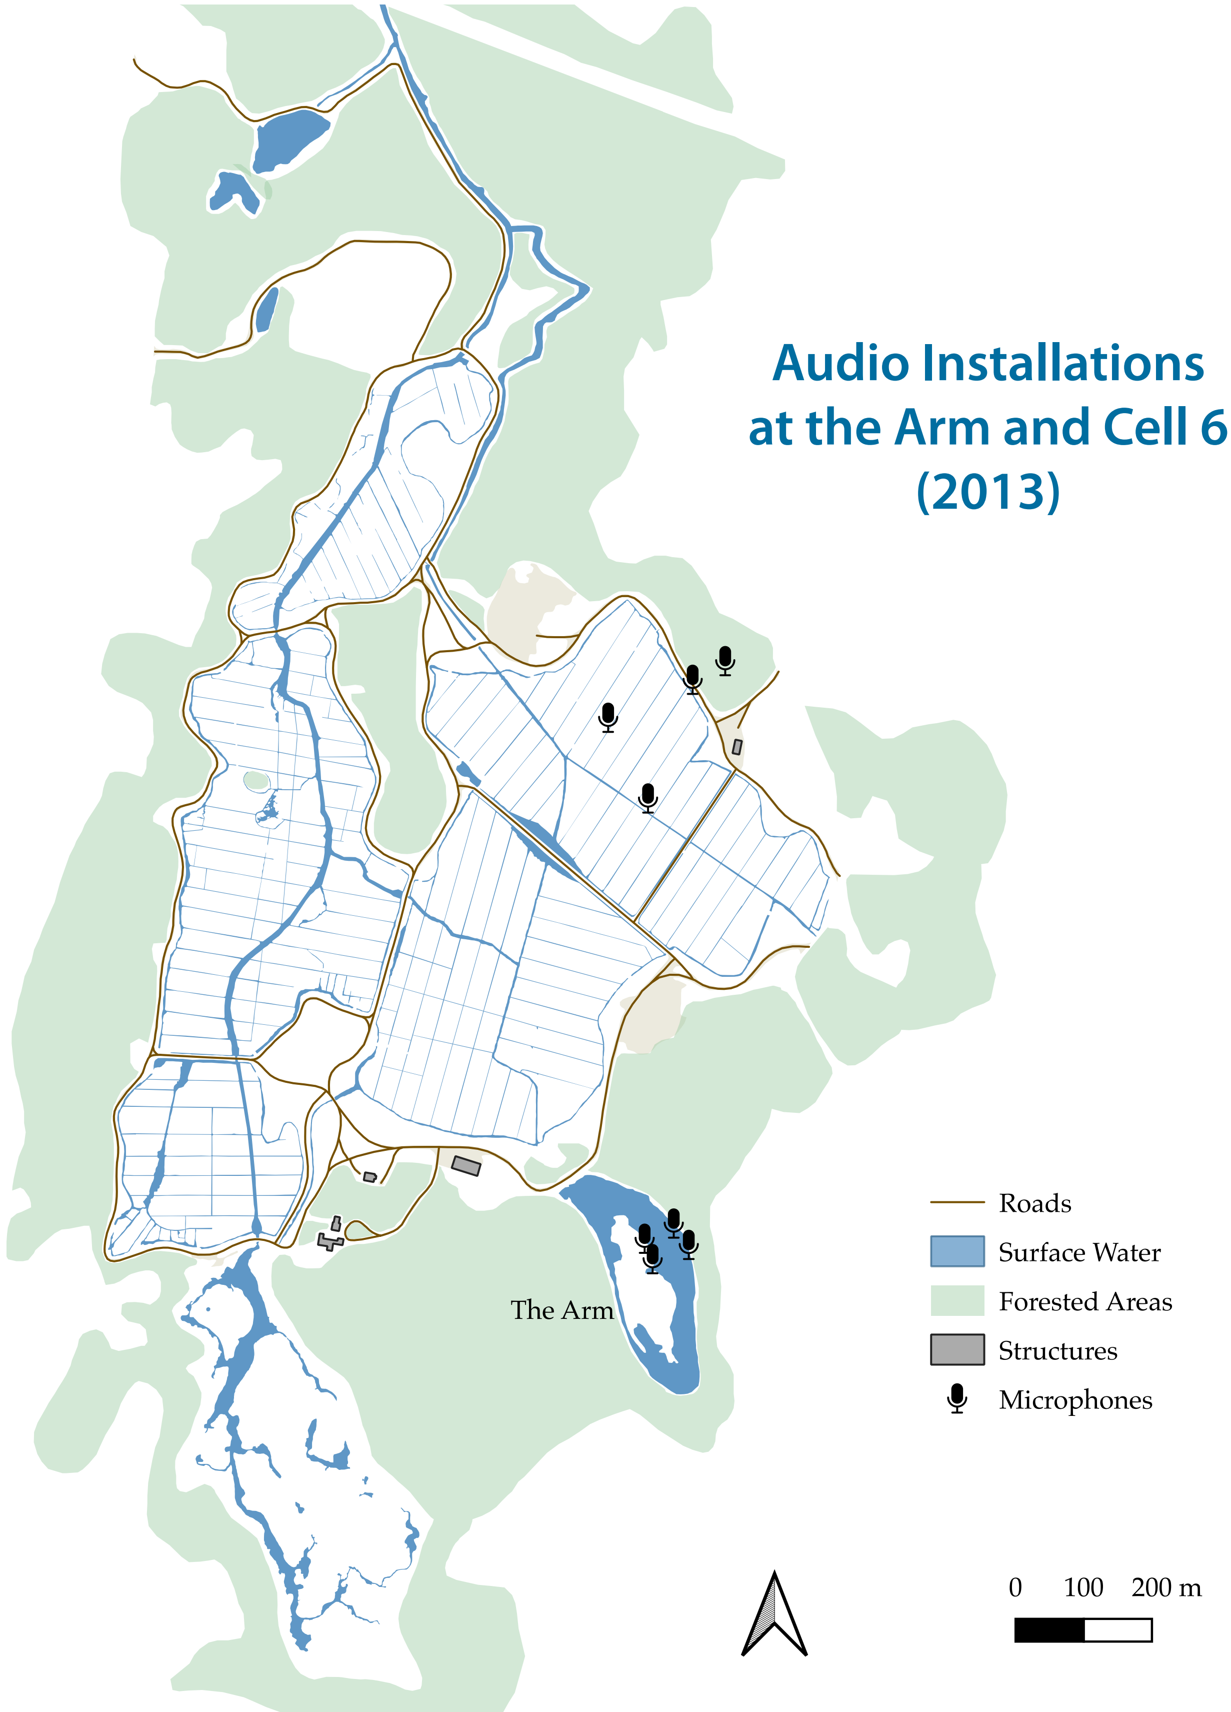

Figure A.3: Audio Installations at the Arm and Agway Barn (2013)

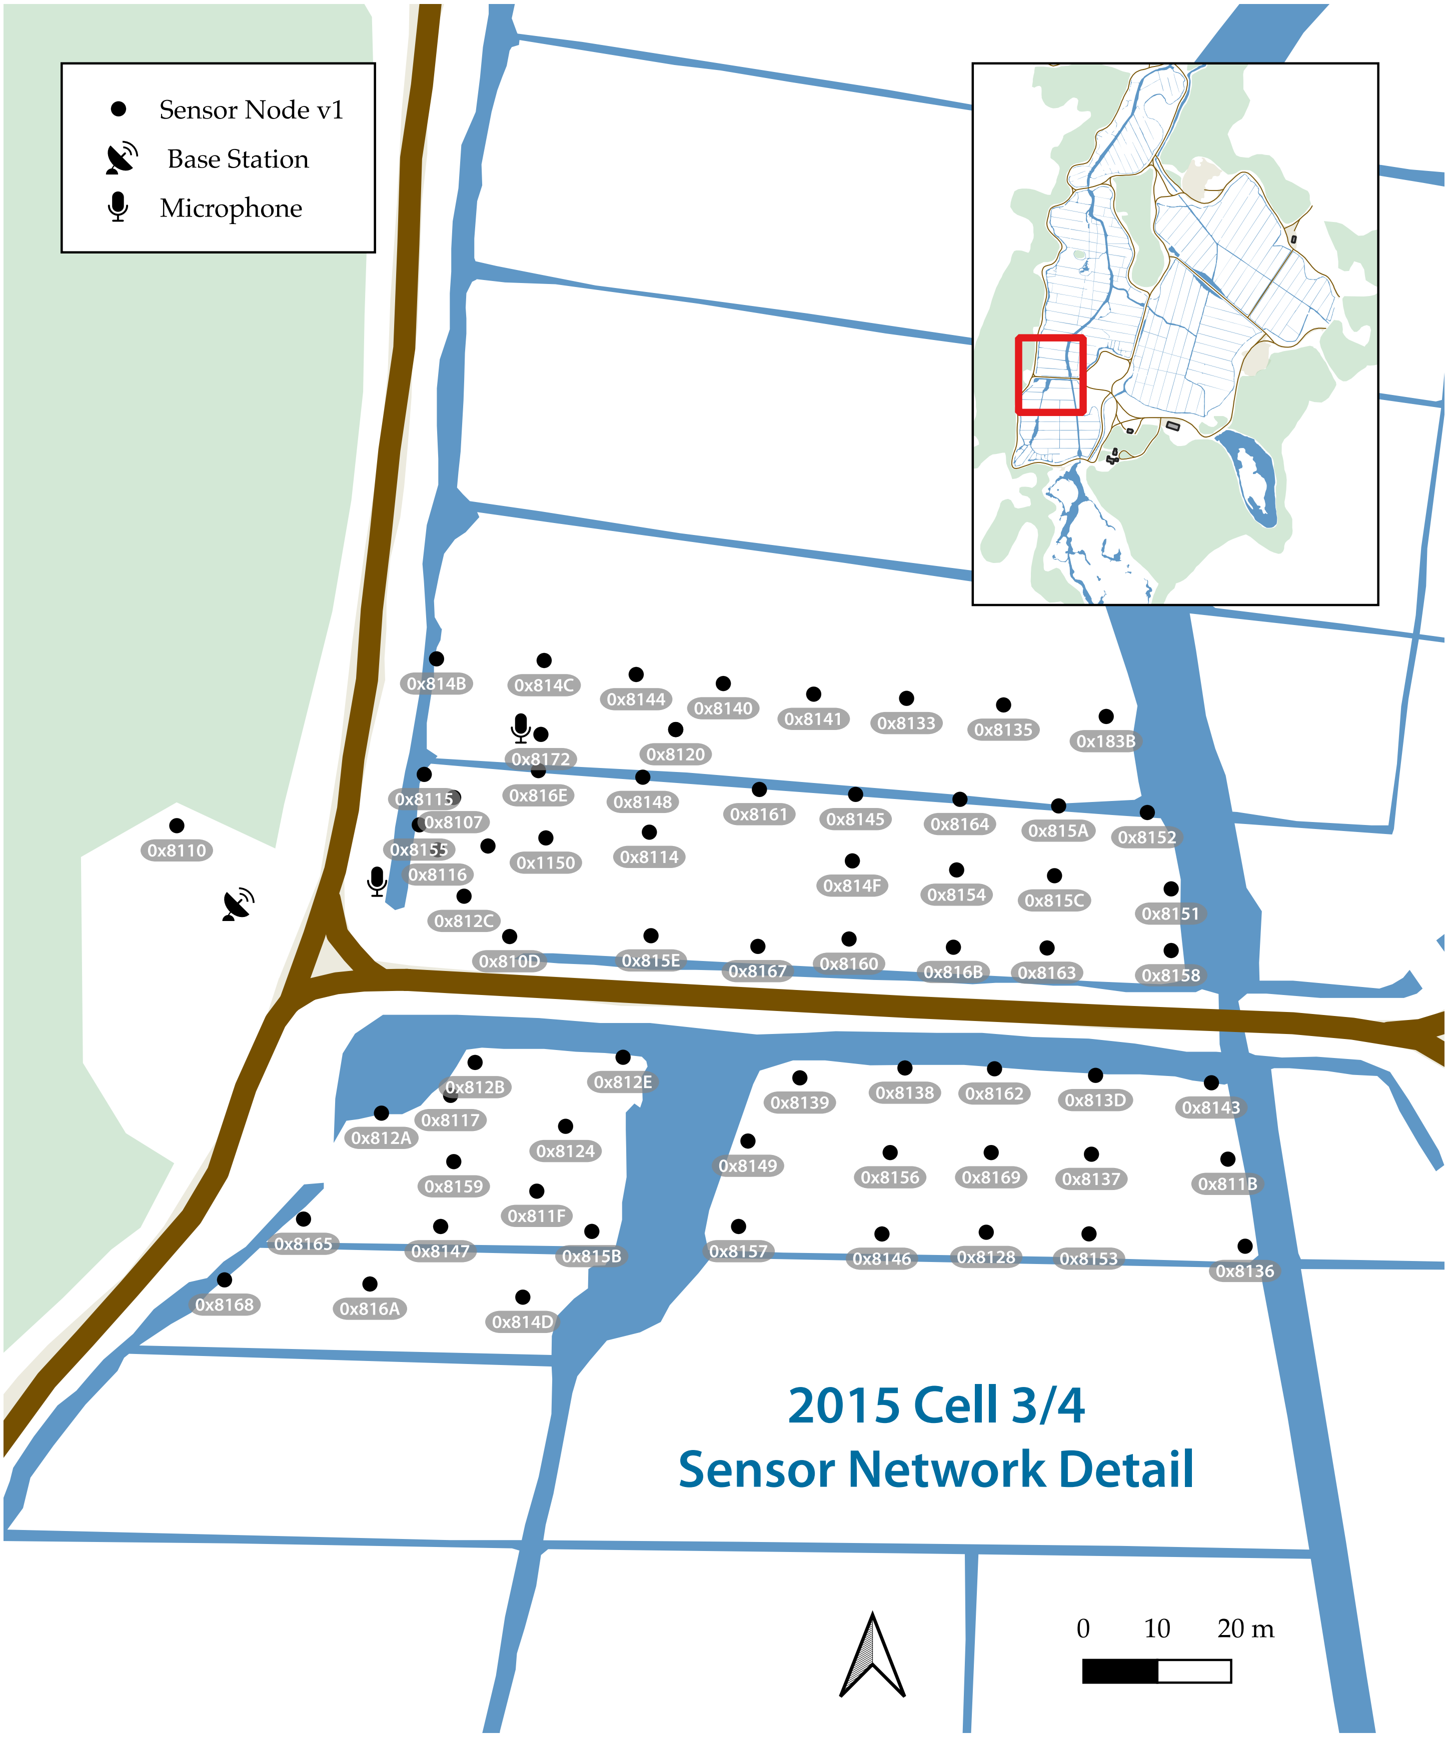

Figure A.4: Cells 3 and 4 Sensor Network Detail (2015)

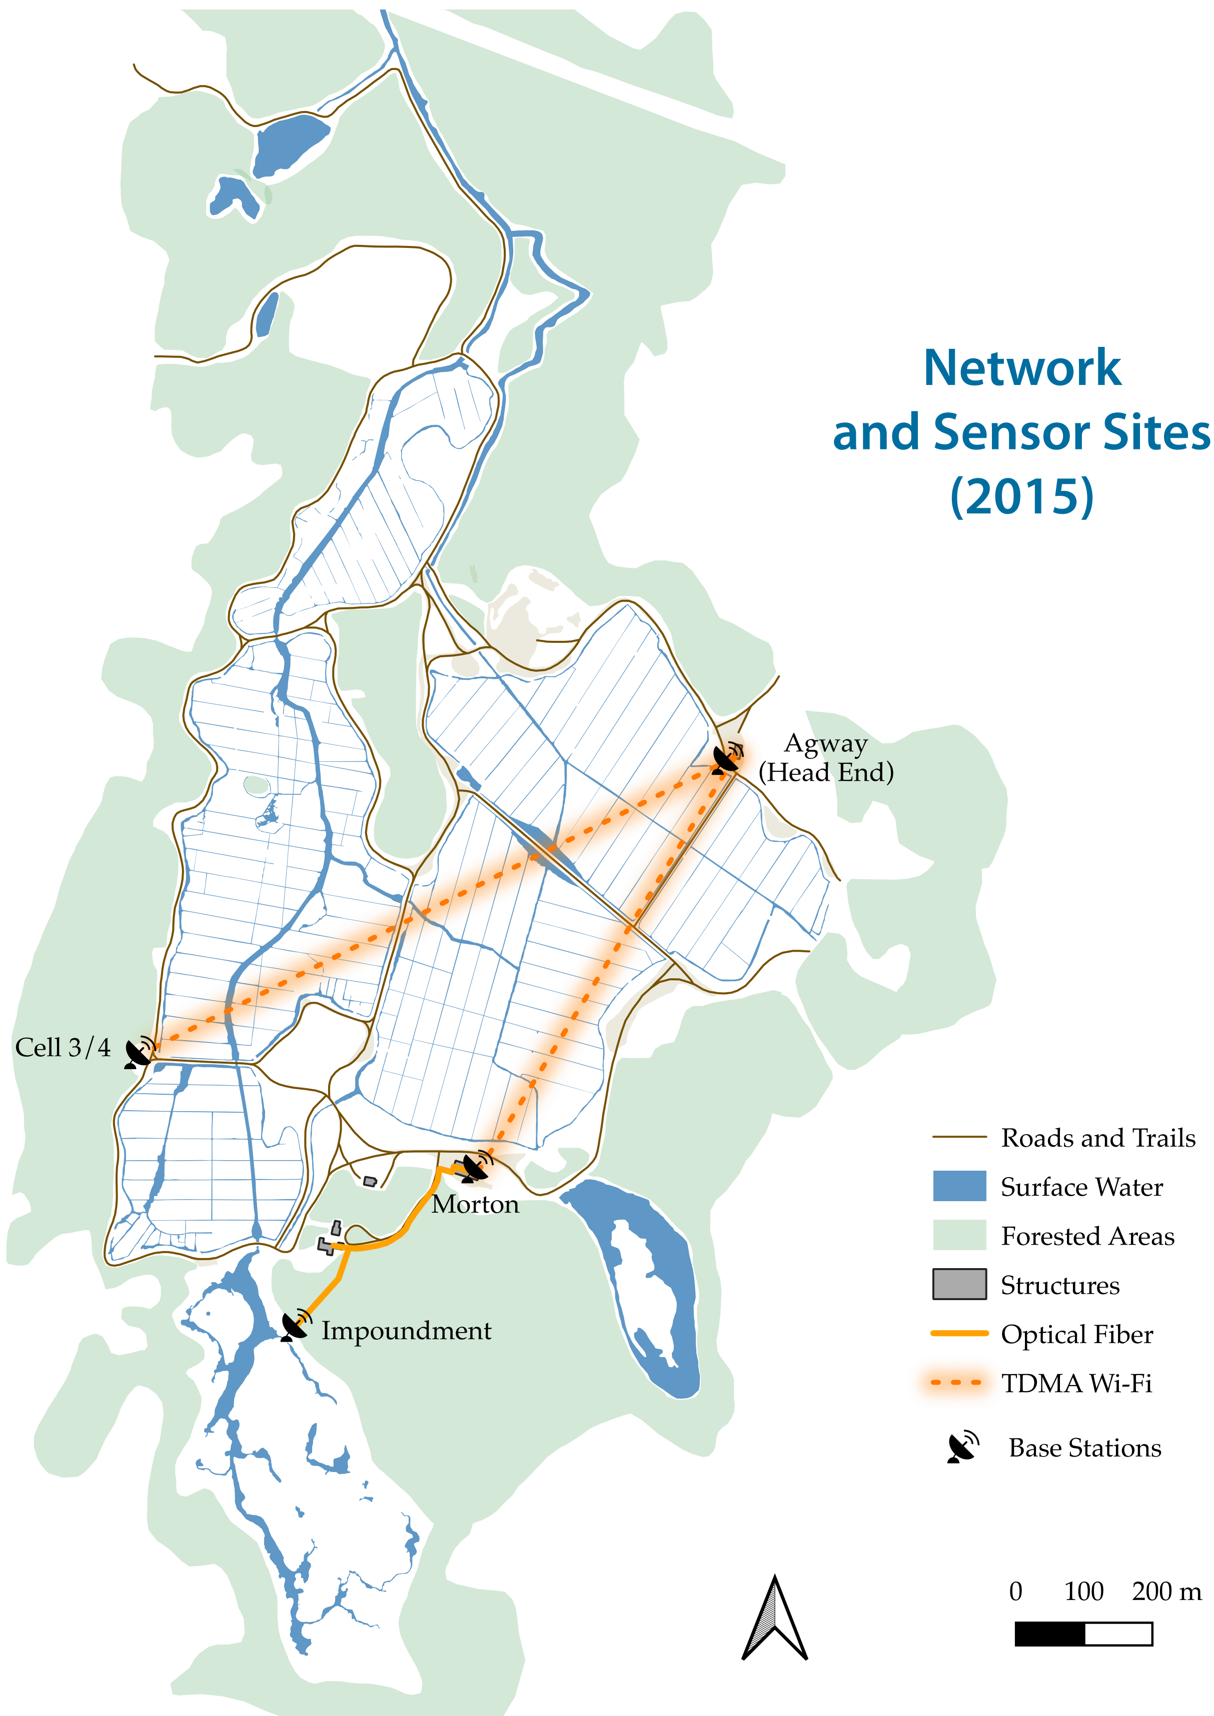

Figure A.5: Network infrastructure and sensor site names (2015).

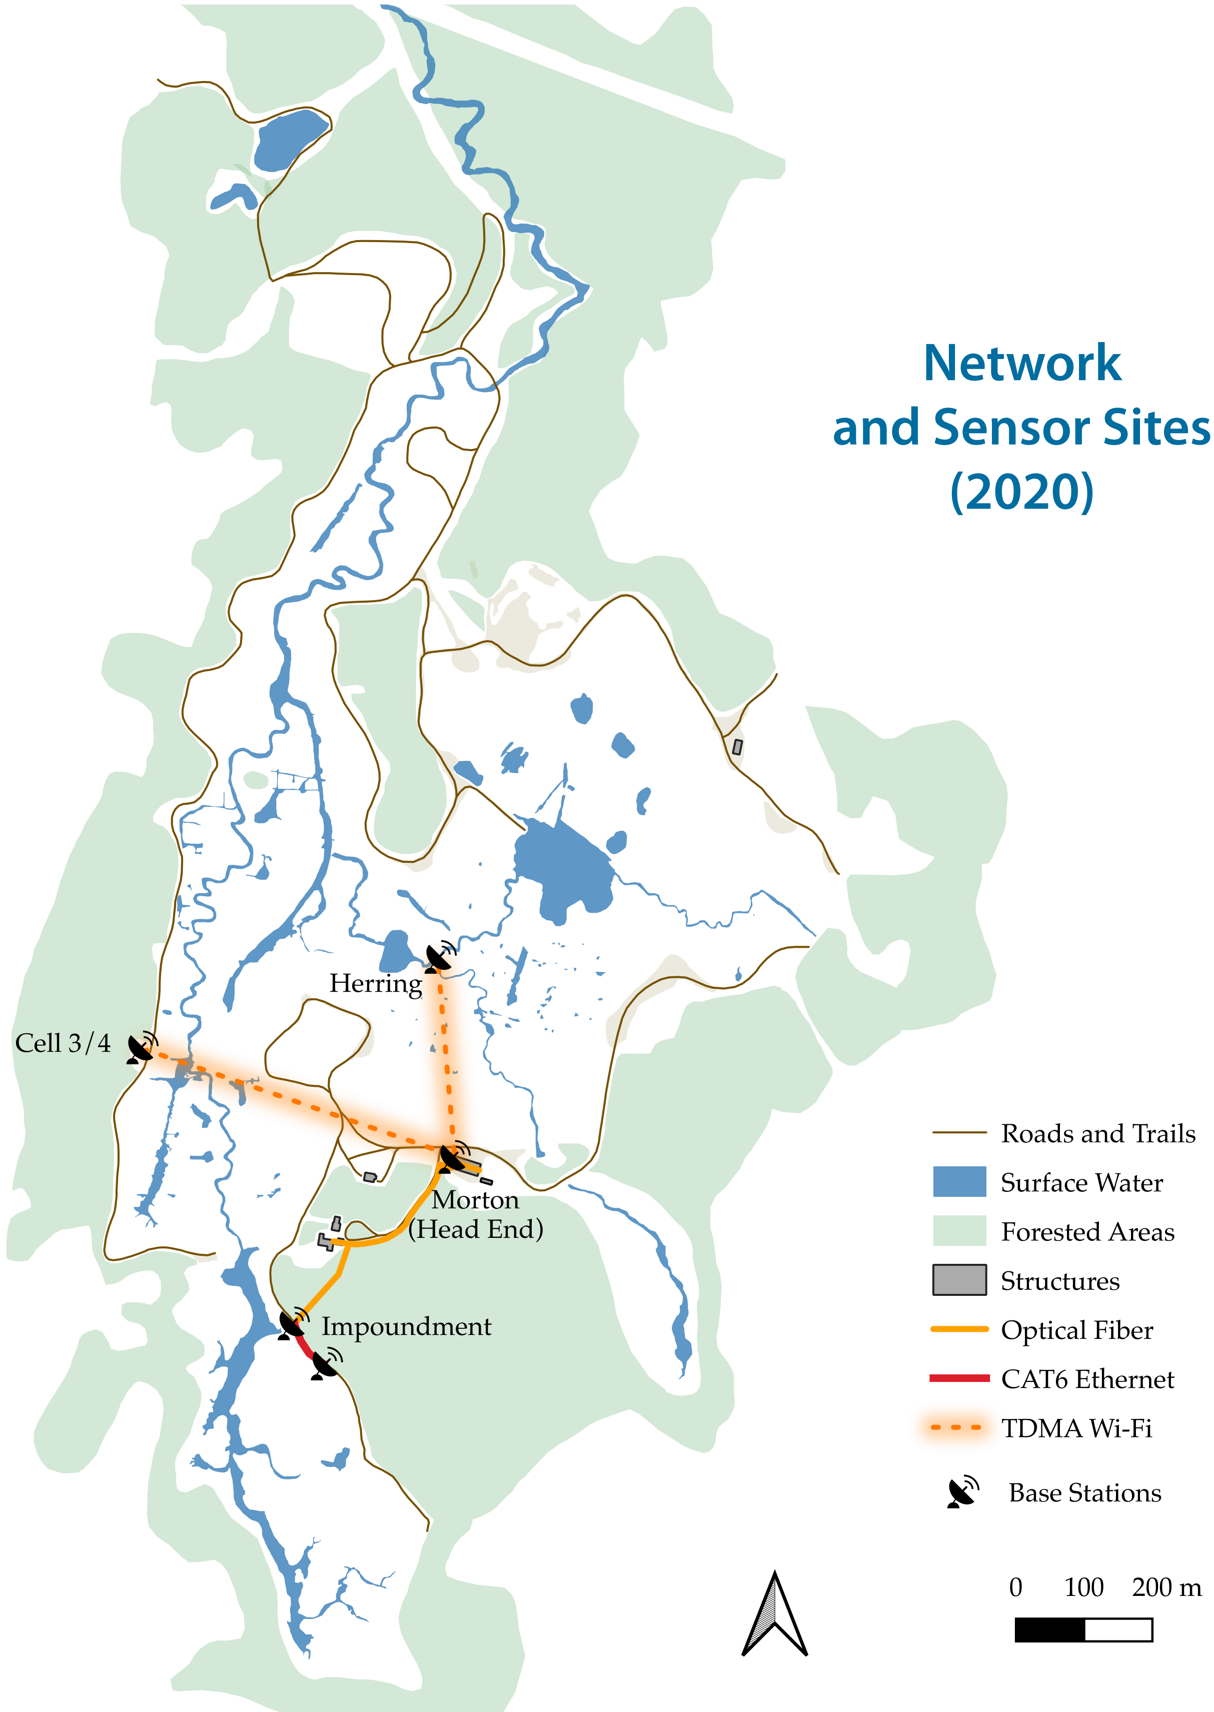

Figure A.6: Network infrastructure and sensor site names (2020).

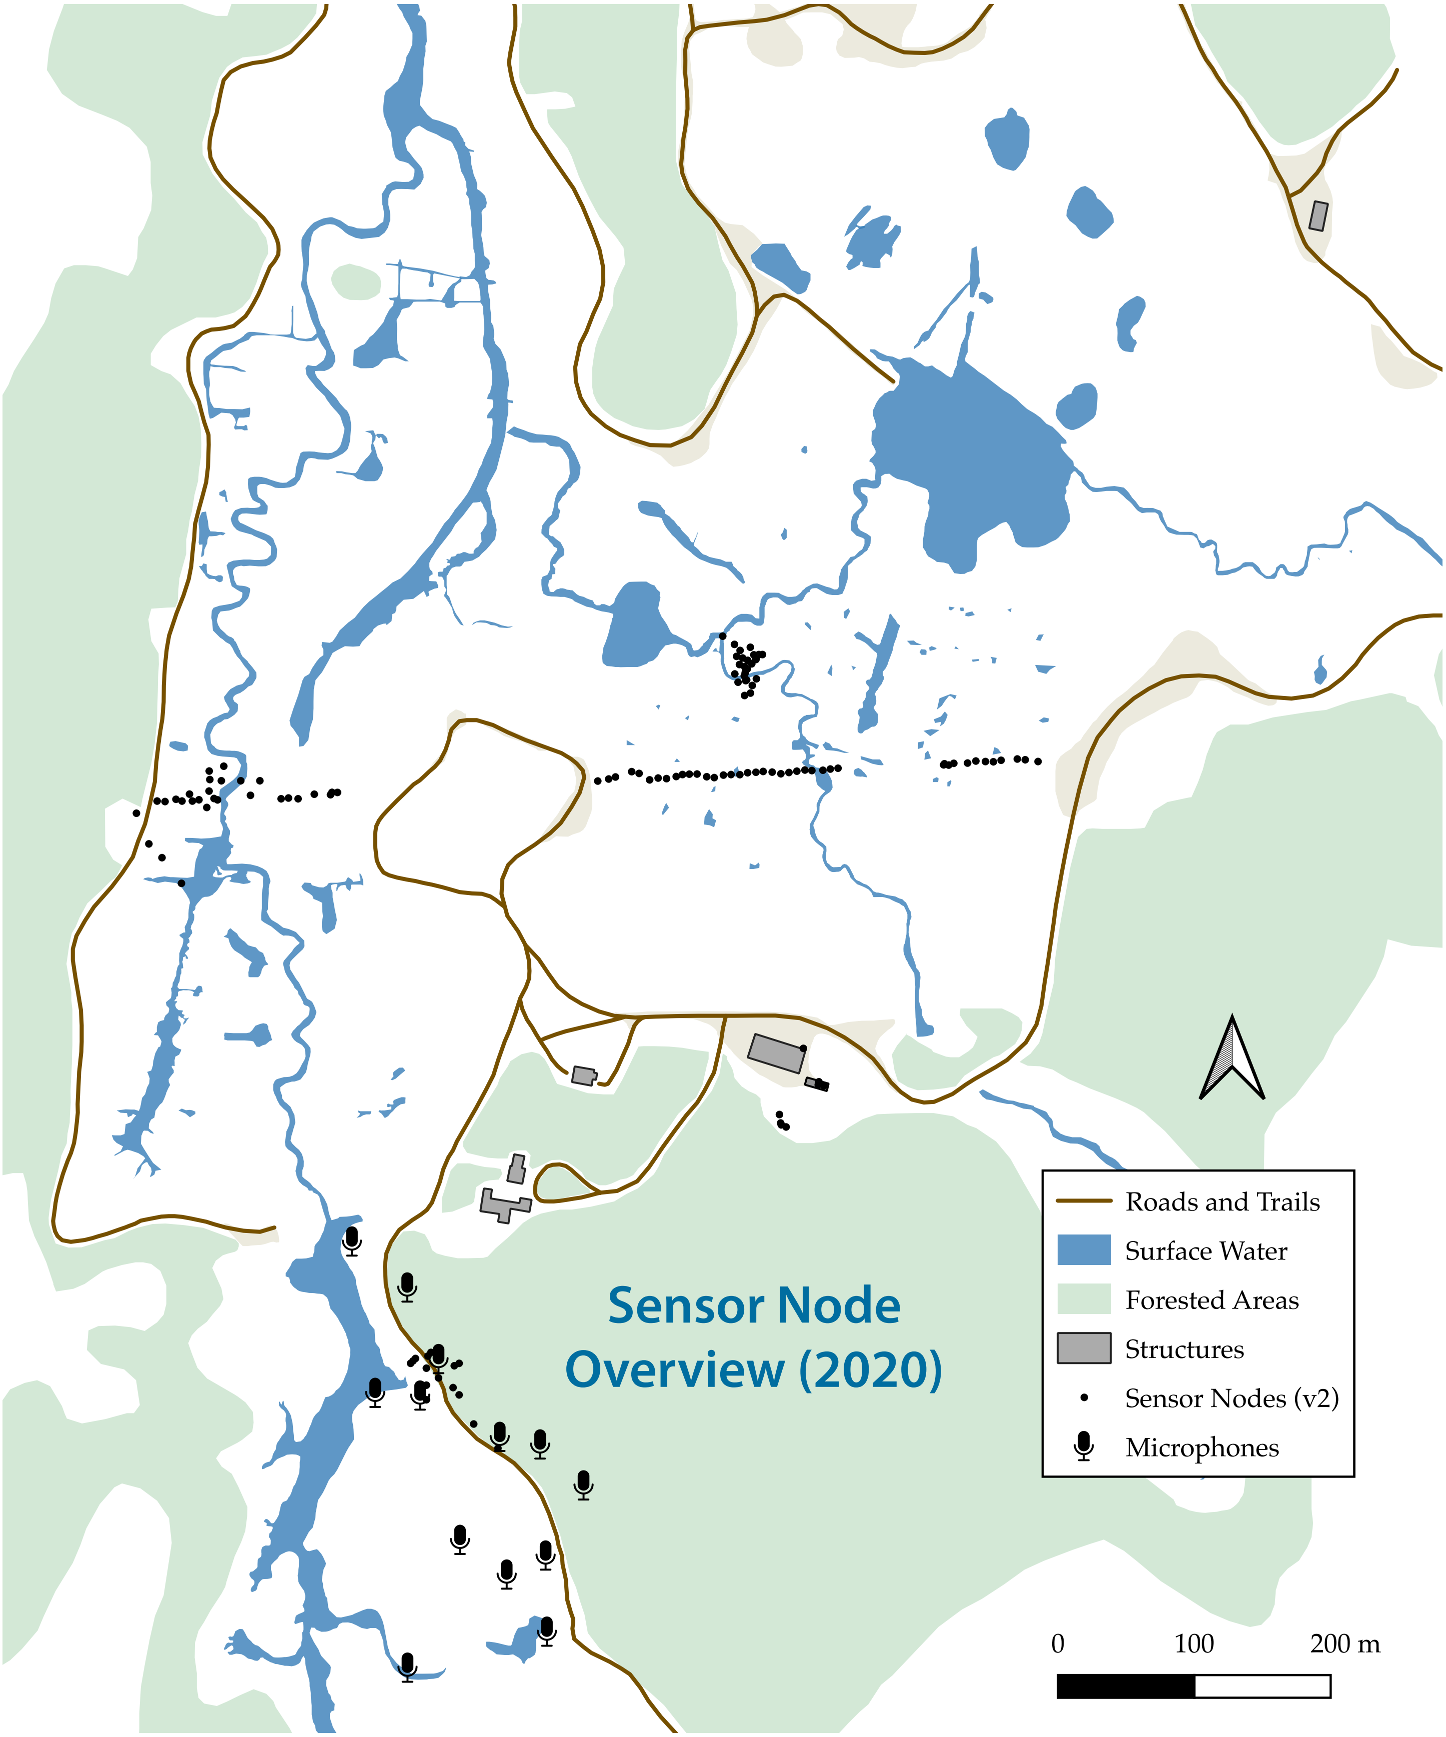

Figure A.7: Sensor nodes and microphones (2020).

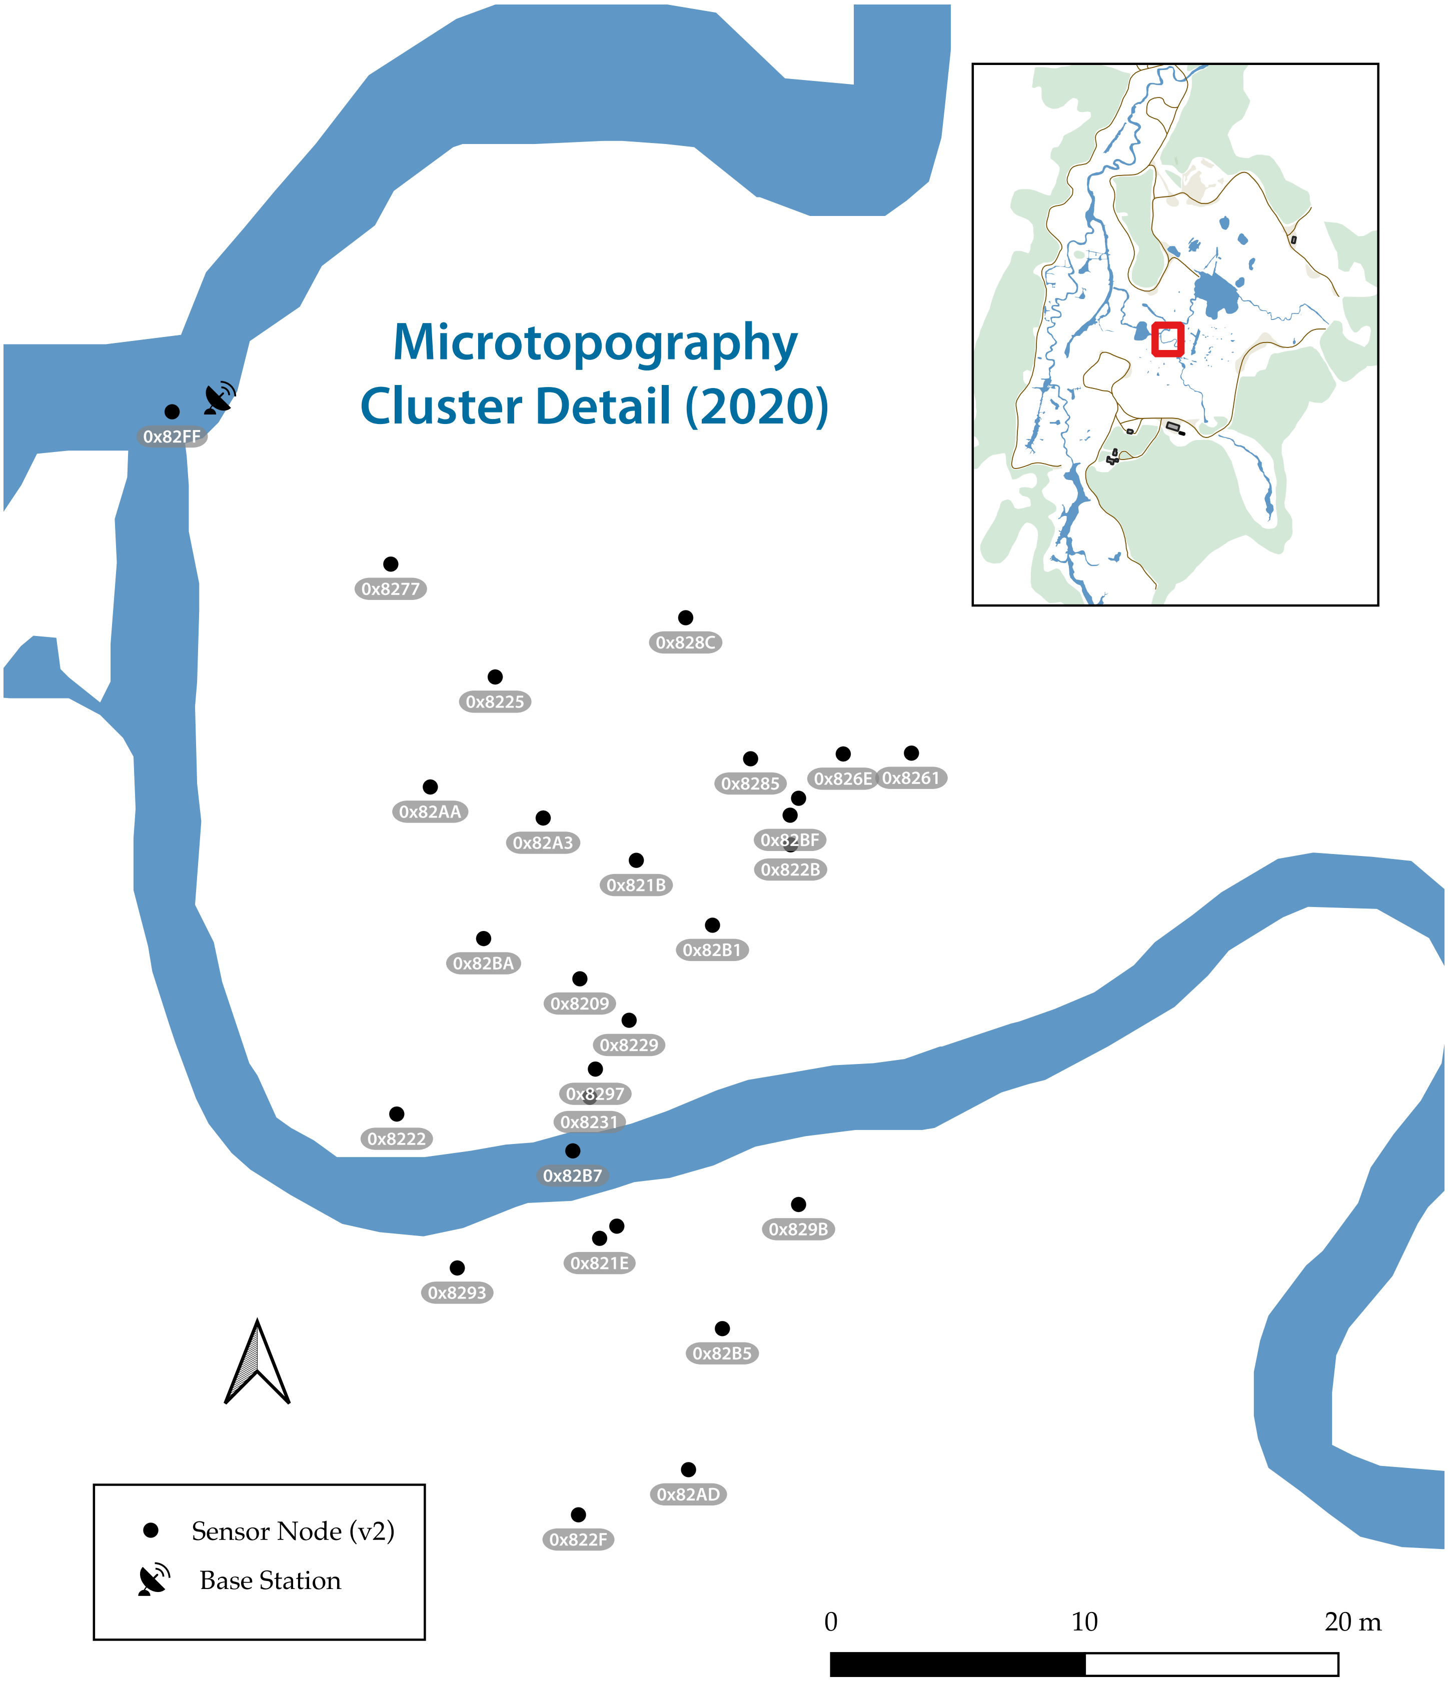

Figure A.8: Microtopography sensor cluster detail (2020).

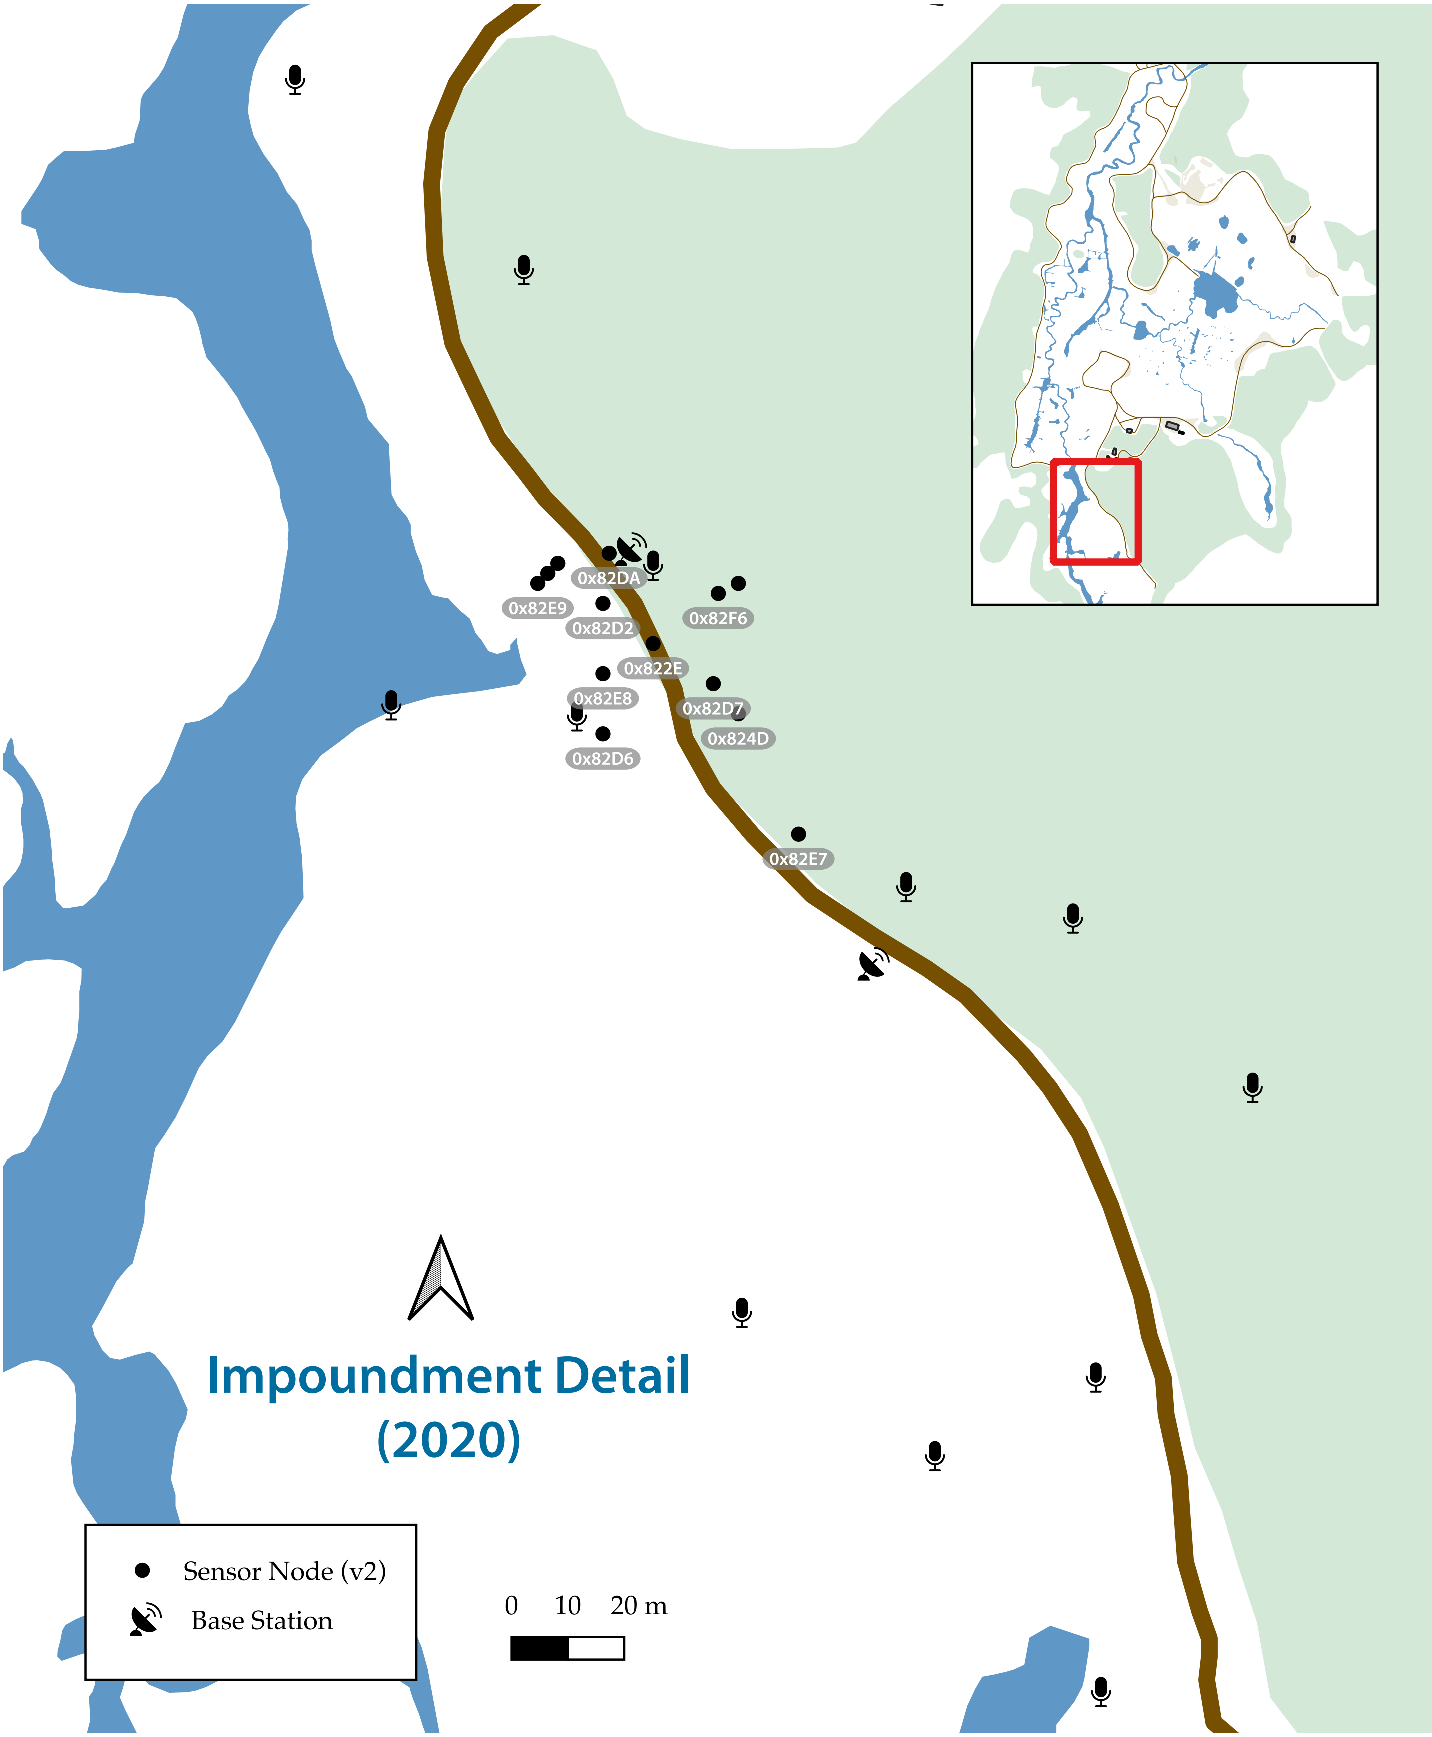

Figure A.9: Impoundment sensors and audio detail (2020).

Appendix B

Assembly, Deployment, and Configuration Instructions

This appendix contains a few of the instructional documents that I prepared for volunteers to use while assembling, deploying, and configuring sensor nodes. This is only a small sample of the extensive technical documentation produced over the course of the project.

Figure B.1 shows a single-page (double-sided) laminated handout that was part of the sensor installation toolkit. I trained volunteers on the deployment process by walking them through the installation of a few sensors. The handout served as a reference and checklist to make sure no steps were missed.

Section B.2 is adapted from a document that describes how to assemble the external probe harness assembly that was attached to sensor nodes in the soil moisture experiments. This work was done by volunteers in the lab prior to deployment, so that a completed probe harness could be plugged in during field deployment.

B.1 Sensor Node Installation Instructions

Tidmarsh Sensor Node Installation Instructions

Rev. A 2018/10/04 Transect Installation

1 Locate Flag

- Flag indicates probe location.
- Place sensor node as close as possible, but leave room to dig the hole.

2 Install Stake

- Hammer into ground until about 70cm remain above ground.
- Stake should stay firmly in place; use a longer pipe if required for stability.

3 Zip Tie Sensor

- The front of the sensor node should face toward the north.
- Tighten zip ties around the pipe with the zip tie puller. Be careful not to pull too hard.
- Be careful not to damage the motion sensor with the puller.
- Cut off excess zip tie with puller.

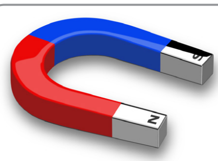

4 Reset Sensor

- Use a magnet to reset the node (slide over enclosure underneath vent).
- The green light on top will blink twice when node boots.

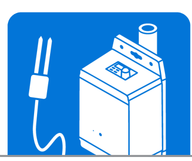

5 Load Configurator

- Open the configurator app and enter the node ID from this sensor.

6 Connect Probes

- If not installing probes, skip to step 12.
- Connect probes to connector on the bottom of the node.
- Select “Probes” in the app.
- You can wake the node immediately by knocking on it.

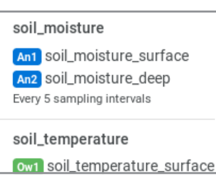

7 Set Probe Channels

- Identify the soil moisture probes by holding one and refreshing to see which has the higher value. Note probe channels.
- Repeat for temperature probes.
- Moisture probes are on fixed channels.
- Drag the temp probes into order.
- Save your changes and return to the menu.

(a) instructions page 1

Figure B.1: Instructions given to volunteers during the large sensor node deployments in the summer of 2019.

8 Dig Hole

- Use a post-hole digger to make the hole.
- Make the sides clean and straight.
- Place dirt in a bucket so it is easy to refill into the hole later.
- Work quickly! Water will start filling the hole as soon as you dig.

9 Install Probes

- Position the jig so the tab is flush with the soil surface (not on top of any plant matter).
- Push the moisture probes into the soil through the notches in the jig.
- Move the jig slightly to the side and repeat with the temperature probes.

10 Photograph Probes

- Using the “Pictures” function in the app, take a photo of the probes in the ground, with the jig in the frame.
- The photo should clearly show the soil where the probes are inserted so that the soil type can be determined.
- Multiple photos are okay.

11 Fill Hole

- Refill the hole with dirt from the bucket.
- Be careful not to disturb the probes.
- Zip tie the probe box to the pole and neaten up any excess wiring.

12 Photograph Node and Surroundings

- Using the “Pictures” section in the app, take as many photos as you want to document the node and its surroundings. (Photos of the installation process are okay too!)
- Feel free to use the “Notes” section with any relevant comments about the installation.

13 Set Sensor Location

- Go to the “Location” section of the app and select “Pick Flag”. Select the ID that was written on the flag. The map will center on this location.
- Save your changes and exit to the menu.

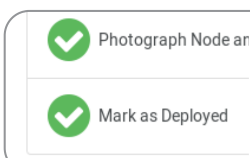

14 Mark Node as Deployed

- Select the “Deploy” function in the app.
- Make sure all relevant steps have check-marks. If not, revisit the appropriate step.
- Tap “Mark as Deployed”.

15 Complete

- The installation is complete!
- Check to make sure that cables are neat, the stake is secure, and that you have all of your tools before moving on.

(b) instructions page 2

Figure B.1: Instructions given to volunteers during the large sensor node deployments in the summer of 2019.

B.2 Probe Harness Assembly Instructions

This document describes how to build and test the soil moisture / temperature probe harness to interface with the Tidmarsh sensor node in 21 easy steps.

After you have assembled several sensor harnesses and are familiar with the complete procedure, you may find the table in Step 16 to be a useful quick reference for the connections that need to be made.

Step 1. Things You Will Need

Gather the following materials and tools (Fig. B.2a):

Materials

- Splice box
- Splice box lid
- Four splice box lid screws
- Cable gland (PG7)
- Two orange 2-wire scotch locks
- Three red 3-wire scotch locks
- Zip tie
- Two EC-5 soil moisture probes
- Two DS18B20 temperature probes
- Switchcraft pigtail cable

Tools

- Wire stripper
- Scotch lock crimping pliers
- PH1 screwdriver
- Soldering iron
- Solder

The sensor cables should be 1 meter long, with the jacket stripped away to expose the inner wires. The individual wires should not yet be stripped—if they have been pre-stripped at the factory, use a wire cutter to trim away the stripped ends.

Step 2. Install the cable gland

Remove the nut from the cable gland and insert it through the hole in the splice box into which it snugly fits. Reinstall and tighten the nut (Fig. B.2b).

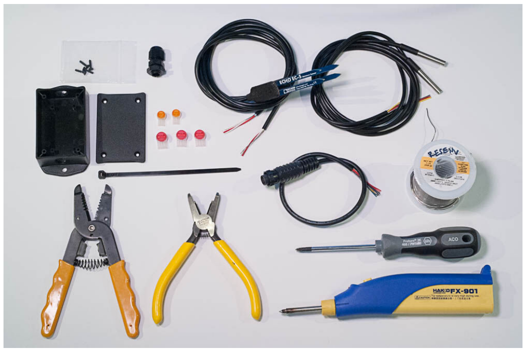

(a) Step 1. Things You Will Need

(b) Step 2. Install the cable gland

(c) Step 3. Insert the pigtail

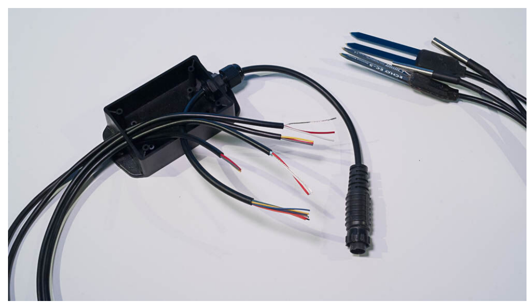

(d) Step 4. Insert the Sensor Cables

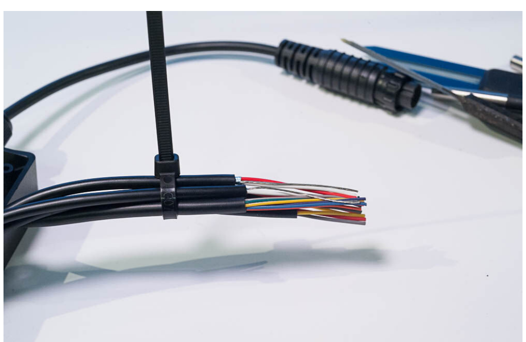

(e) Step 5. Gather the Cable Ends

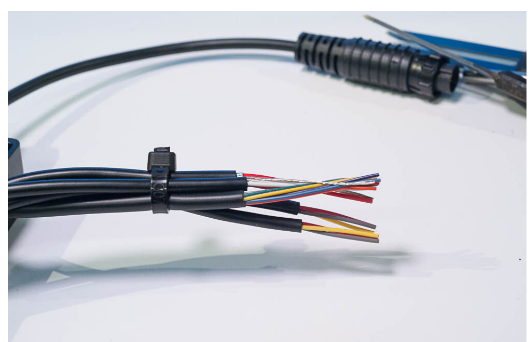

(f) Step 6. Trim the Zip Tie

Figure B.2: Probe harness asseembly instructions, steps 1-6

Step 3. Insert the Pigtail

Loosen the gland closure and insert the pigtail through it, pulling most of the slack into the box to make it easy to work with. Leave the gland loose for now (Fig B.2c).

Step 4. Insert the Sensor Cables

Insert the cables for all four sensor probes through the larger hole in the other side of the splice box, again pulling a bit of slack into the box (Fig B.2d).

Step 5. Gather the Cable Ends

Gather all five wires together with the ends oriented in the same direction, and line up the tips of all of the wires. Use the zip tie to secure the bundle around the cable jackets (Fig. B.2e).

Step 6. Trim the Zip Tie

Make sure the zip tie is tight, then snip the end (Fig B.2f).

Step 7. Soil Moisture Excitation

Locate the blue wire from the pigtail cable and the two white wires from the EC-5 soil moisture probes, and insert them into a red 3-position scotch lock connector (Fig. B.2g). (The wires can be in any order in the connector). Look through the clear back to make sure that all of the wires are fully inserted into the connector, past the metal parts. (If necessary, trim the wires so that they are all the same length and easily insert fully into the connector).

Step 8. Crimp the Scotch Lock

Using the crimping pliers, squeeze the button on the scotch lock until it is completely flush (Fig B.2h).

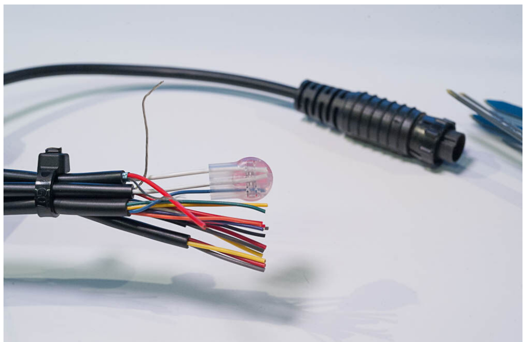

(g) Step 7. Soil Moisture Excitation

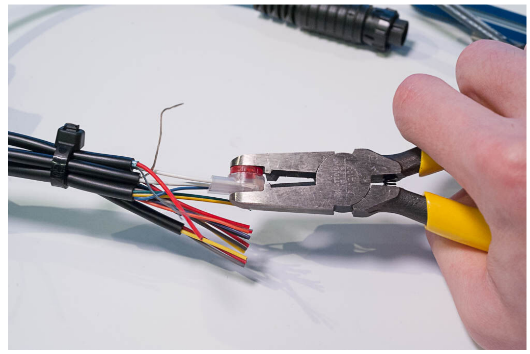

(h) Step 8. Crimp the Scotch Lok

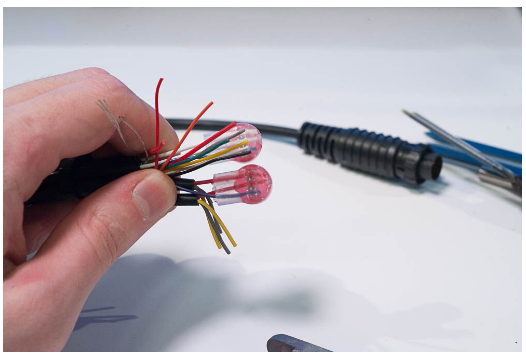

(i) Step 9. Temperature Sensor Excitation

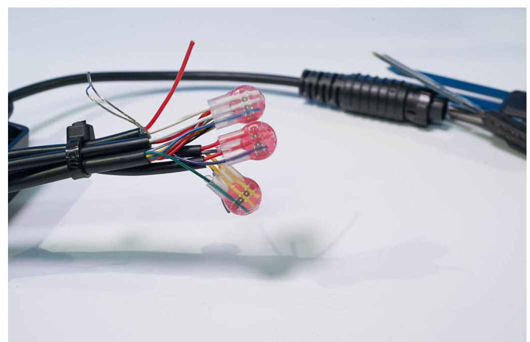

(j) Step 10. Temperature Sensor Onewire Data

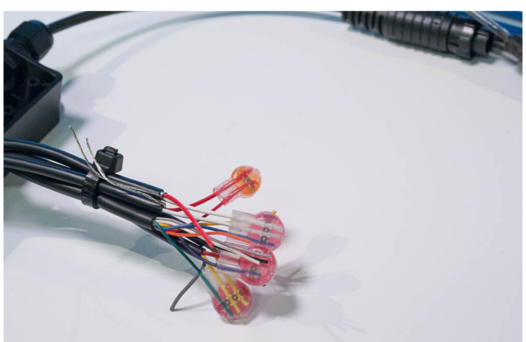

(k) Step 11. Soil Moisture Output 1

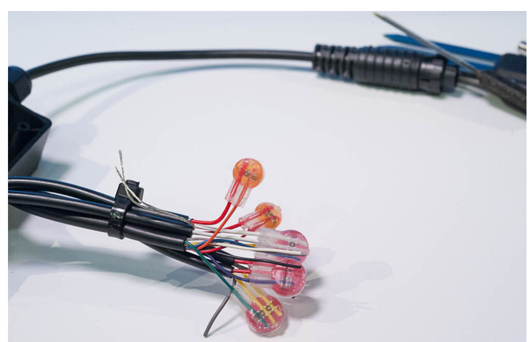

(l) Step 12. Soil Moisture Output 2

Figure B.2: Probe harness assembly instructions, steps 7-12

Step 9. Temperature Sensor Excitation

Locate the purple wire from the pigtail cable and the two red wires from the DS18B20 temperature probes and crimp them together using another red 3-position scotch lock, using the same procedure as steps 7 and 8 (Fig. B.2i).

Step 10. Temperature Sensor Onewire Data

Locate the green wire from the pigtail cable and the two yellow wires from the DS18B20 temperature probes and crimp them together using another red 3-position scotch lock, using the same procedure as steps 7 and 8 (Fig. B.2j).

Step 11. Soil Moisture Output 1

Locate the red wire from the pigtail cable and a red wire from one of the EC-5 soil moisture probes (it does not matter which one) and crimp the two wires together using an orange 2-position scotch lock (Fig. B.2k).

Step 12. Soil Moisture Output 2

Locate the orange wire from the pigtail cable and a red wire from the remaining EC-5 soil moisture probe and crimp the two wires together using an orange 2-position scotch lock (Fig. B.2l).

Step 13. Strip the Ground Wires

Using a 26-gauge wire stripper, strip about 1cm of insulation from the black wire from the pigtail cable and the two black (or gray) wires from the two DS18B20 temperature probes (Fig. B.2m).

Step 14. Twist the Ground Wires

Twist together the ends of the black wire from the pigtail cable, the two black or gray wires from the DS18B20 temperature probes, and the two uninsulated drain wires from the EC-5 soil moisture probes (Fig. B.2n).

(m) Step 13. Strip the Ground Wires

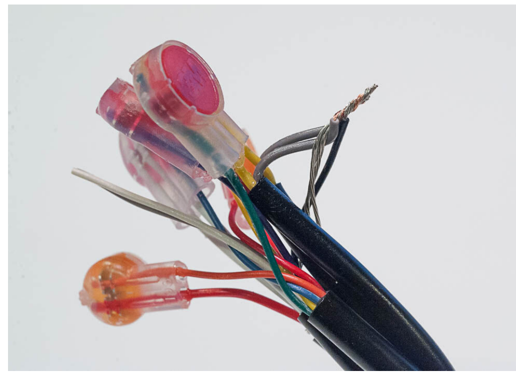

(n) Step 14. Twist the Ground Wires

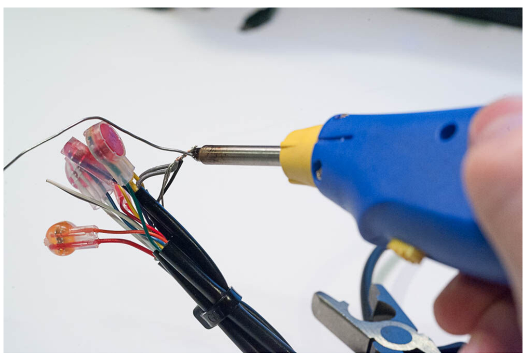

(o) Step 15. Solder the Ground Wires

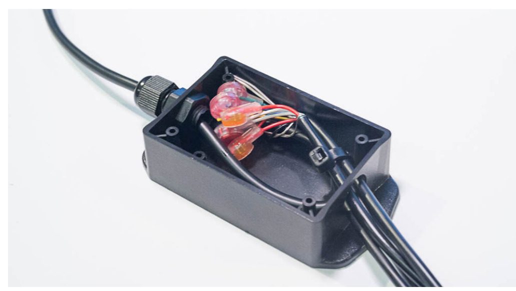

(p) Step 17. Stuff the Box

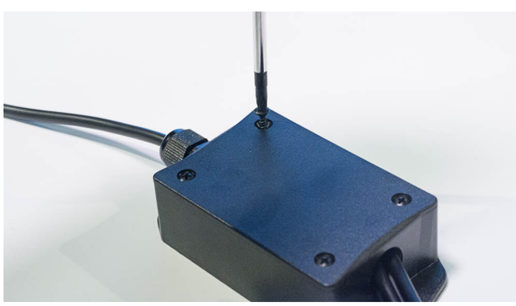

(q) Step 18. Install the Box Lid

(r) Step 19. Connect to a Node

Figure B.2: Probe harness assembly instructions, steps 13-19

Table B.1: External probe harness connections

Function	Pigtail Color	Connection 1	Connection 2
EC-5 Excitation	Blue	EC-5 #1 White	EC-5 #2 White
DS18B20 Excitation	Purple	DS18B20 #1 Red	DS18B20 #2 Red
DS18B20 Data	Green	DS18B20 #1 Yellow	DS18B20 #2 Yellow
EC-5 #1 Data	Red	EC-5 #1 Red	
EC-5 #2 Data	Orange	EC-5 #2 Red	
Ground	Black	EC-5 unshielded wires (2)	DS18B20 black/gray wires (2)

Step 15. Solder the Ground Wires

Complete the connection of the ground wires by soldering them together (Fig. B.2o).

Step 16. Check Connections

Check the completed bundle to confirm that all connections have been made according to Table B.1.

Step 17. Stuff the Box

Pull the slack on the pigtail cable and the sensor cables out of the box, and tuck the bundle of connections into place inside the box. Make sure that none of the wires are blocking any of the screw holes (Fig. B.2p).

Step 18. Install the Box Lid

Use the four lid screws to secure the lid onto the splice box (Fig. B.2q).

Step 19. Connect to a Node

Connect the newly assembled cable harness to a sensor node for testing (Fig. B.2r).

Step 20. Check Temperature Probes

Open the sensor node configurator and go to the **Config** screen. Load the template for a soil moisture node under **External Probe Configuration** (if it has not already been done). Open the **Asssign and Test Probes** section and refresh the **Onewire Probes**.

You should see two probes listed with unique serial numbers (they should appear under the **Excluded** section as they have not yet been assigned to this sensor node). The temperatures should be approximately room temperature (20°C).

Step 21. Check Soil Moisture Probes Probes

Refresh the Analog probes as necessary and look at the values for the first two channels to confirm the following.

- When a probe is in free air (not near the table, other probes, your hand) it should read at around 200 mV.
- When a probe is immersed in water, it should read around 750 mV.

References

[1] *2.4-GHz Inverted F Antenna*. Texas Instruments. 2007. URL: https://www.ti.com/lit/an/swru120d/swru120d.pdf.

[2] *Agriculture | Davis Instruments*. 2019. URL: https://www.davisinstruments.com/usecase/agriculture.

[3] Kemal Akkaya and Mohamed Younis. "A survey on routing protocols for wireless sensor networks". In: *Ad hoc networks* 3.3 (2005), pp. 325–349.

[4] *AVR2130: Lightweight Mesh Developer Guide*. 2016. URL: https://www.microchip.com//wwwAppNotes/AppNotes.aspx?appnote=en591088.

[5] Azby Brown et al. "Safecast: successful citizen-science for radiation measurement and communication after Fukushima". In: *Journal of Radiological Protection* 36.2 (2016), S82.

[6] Halsey Burgund. *Bog People*. Sound art installation. 2016. URL: https://halseyburgund.com/projects/bog.html.

[7] Jenna Burrell, Tim Brooke, and Richard Beckwith. "Vineyard computing: Sensor networks in agricultural production". In: *Pervasive Computing, IEEE* 3.1 (2004), pp. 38–45.

[8] J.B. Campbell. *Introduction to Remote Sensing*. Taylor & Francis, 2002. ISBN: 9780415282949. URL: https://books.google.com/books?id=o-zPxbhgO%5C_AC.

[9] *Campbell Scientific*. 2020. URL: https://www.campbellsci.com/.

[10] Adam Clarke. "Guide to soil moisture sensors: Why use them and the options". In: *Farmers Weekly* (Jan. 22, 2020). URL: https://www.fwi.co.uk/machinery/technology/guide-to-soil-moisture-sensors-why-use-them-and-the-options.

[11] *Cranberry Bog Program*. Division of Ecological Restoration, Massachusetts Department of Fish and Game. 2020. URL: https://www.mass.gov/cranberry-bog-program.

[12] *Cranberry Revitalization Task Force Final Report*. Technical Report. Massachusetts Department of Agricultural Resources, 2016. URL: https://www.mass.gov/files/documents/2016/08/tx/cran-rtfinal-5-24-16.pdf.

[13] *Data Logger for Research | ZL6 Advanced Cloud Data Logger*. 2020. URL: https://www.metergroup.com/environment/products/zl6-data-logger/.

[14] John V Draper, David B Kaber, and John M Usher. "Telepresence". In: *Human Factors: The Journal of the Human Factors and Ergonomics Society* 40.3 (1998), pp. 354–375.

[15] Gershon Dublon. *livingsounds.earth*. slow immediate. 2020. URL: https://livingsounds.earth.

[16] Gershon Dublon. "Sensor(y) Landscapes: technologies for new perceptual sensibilities." Ph.D. thesis. Massachusetts Institute of Technology, 2018.

[17] Gershon Dublon et al. "Doppellab: Tools for exploring and harnessing multimodal sensor network data". In: *Sensors* (2011), pp. 1612–1615.

[18] Clement Duhart et al. "Deep Learning for Wildlife Conservation and Restoration Efforts". In: *36th International Conference on Machine Learning, Long Beach*. Vol. 5. 2019.

[19] Clement Duhart et al. "Deep learning locally trained wildlife sensing in real acoustic wetland environment". In: *International Symposium on Signal Processing and Intelligent Recognition Systems*. Springer. 2018, pp. 3–14.

[20] *EC-5 Manual*. METER Group. 2019. URL: https://www.metergroup.com/environment/products/ec-5-soil-moisture-sensor/.

[21] Tomi Engdahl. *Powering microphones*. 1997-2012. URL: https://www.epanorama.net/circuits/microphone_powering.html.

[22] Deborah Estrin et al. "Instrumenting the world with wireless sensor networks". In: *Acoustics, Speech, and Signal Processing, 2001. Proceedings. (ICASSP'01). 2001 IEEE International Conference on*. Vol. 4. IEEE. 2001, pp. 2033–2036.

[23] *Frequently Asked Questions*. METER Environment. 2020. URL: https://www.metergroup.com/environment/faqs/?faq_category=518#faq-6787.

[24] Steve Graham, Claire Parkinson, and Mous Chahine. *Weather Forecasting Through the Ages*. NASA Earth Observatory. 2002.

[25] *Graphite*. 2020. URL: https://graphiteapp.org.

[26] Don Derek Haddad. "Resynthesizing reality : driving vivid virtual environments from sensor networks." Master's thesis. Massachusetts Institute of Technology, 2018.

[27] Don Derek Haddad et al. "Resynthesizing Reality: Driving Vivid Virtual Environments from Sensor Networks". In: *ACM SIGGRAPH 2017 Talks*. SIGGRAPH '17. Los Angeles, California: Association for Computing Machinery, 2017. ISBN: 9781450350082. DOI: 10.1145/3084363.3085027. URL: https://doi.org/10.1145/3084363.3085027.

[28] Danielle K. Hare. "Hydrogeological control on spatial patterns of groundwater seepage in peatlands". Master's thesis. University of Massachusetts Amherst, 2015.

[29] M. Harvey et al. *Thermal infrared and photogrammetric data collected by small unoccupied aircraft system for the evaluation of wetland restoration design at Tidmarsh Wildlife Sanctuary, Plymouth, Massachusetts, USA*. United States Geological Survey. 2019. URL: https://www.sciencebase.gov/catalog/item/5cd6bb45e4b041d8a7f37963.

[30] Jason L Hill and David E Culler. "Mica: A wireless platform for deeply embedded networks". In: *Micro, IEEE* 22.6 (2002), pp. 12–24.

[31] Benjamin R. Hoekstra, Christopher Neill, and Casey D. Kennedy. "Trends in the Massachusetts cranberry industry create opportunities for the restoration of cultivated riparian wetlands". In: *Restoration Ecology* 28.1 (2020), pp. 185–195. DOI: 10.1111/rec.13037. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/rec.13037. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/rec.13037.

[32] Wen Hu et al. "Design and evaluation of a hybrid sensor network for cane toad monitoring". In: *ACM Transactions on Sensor Networks (TOSN)* 5.1 (2009), p. 4.

[33] Chalermek Intanagonwiwat et al. "Directed diffusion for wireless sensor networking". In: *IEEE/ACM Transactions on Networking (ToN)* 11.1 (2003), pp. 2–16.

[34] Fei Jiang, Don Derek Haddad, and Joseph Paradiso. "Baguamarsh: An Immersive Narrative Visualization for Conveying Subjective Experience". In: *International Conference on Human-Computer Interaction*. Springer. 2020, pp. 596–613.

[35] Thomas U Kampe et al. "NEON: the first continental-scale ecological observatory with airborne remote sensing of vegetation canopy biochemistry and structure". In: *Journal of Applied Remote Sensing* 4.1 (2010), p. 043510.

[36] Navdeep Kaur Kapoor, Biswajit Nandy, and Shikharesh Majumdar. "Dynamic allocation of sensor nodes in wireless sensor networks hosting multiple applications". In: *Performance Evaluation of Computer and Telecom-*

munication Systems (SPECTS), 2013 International Symposium on. IEEE. 2013, pp. 51–57.

[37] Robert Kearns. *Tidmarsh Herring Count Project*. 2020. URL: https://www.massaudubon.org/get-outdoors/wildlife-sanctuaries/tidmarsh/get-involved/herring-count-project.

[38] M. Kelly. *JSON Hypertext Application Lanugage*. May 11, 2016. URL: https://tools.ietf.org/html/draft-kelly-json-hal-08.

[39] Philip Levis et al. "Tinyos: An operating system for sensor networks". In: *Ambient intelligence*. Springer, 2005, pp. 115–148.

[40] Q Li et al. "MarshVis: Visualizing real-time and historical ecological data from a wireless sensor network". In: *IEEE VIS Arts Program (VISAP)* (2015).

[41] *Libelium - Connecting Sensors to the Cloud*. 2015. URL: http://www.libelium.com.

[42] Joshua Lifton et al. "Metaphor and manifestation-cross-reality with ubiquitous sensor/actuator networks". In: Institute of Electrical and Electronics Engineers. 2009.

[43] *Living Observatory - Restoration Sites*. 2018. URL: https://www.livingobservatory.org/restoration-sites.

[44] Jaime Lloret et al. "A wireless sensor network deployment for rural and forest fire detection and verification". In: *sensors* 9.11 (2009), pp. 8722–8747.

[45] Evan F. Lynch. "SensorChimes: musical mapping for sensor networks toward augmented acoustic ecosystem." Master's thesis. Massachusetts Institute of Technology, 2016.

[46] Miklós Maróti et al. "The flooding time synchronization protocol". In: *Proceedings of the 2nd international conference on Embedded networked sensor systems*. ACM. 2004, pp. 39–49.

[47] Brian Mayton. *Tracking the Arc of Change*. Short film. 2019. URL: https://vimeo.com/359648748.

[48] *MEMSIC, Inc - Wireless Sensor Networks*. 2015. URL: http://www.memsic.com/wireless-sensor-networks.

[49] William J. Mitsch and James G. Gosselink. *Wetlands.* Hoboken, NJ : John Wiley and Sons, Inc., [2015], 2015. ISBN: 9781118676820. URL: https://search.ebscohost.com/login.aspx?direct=true&db=cat00916a&AN=mit.002406913&site=eds-live&scope=site.

[50] Robert Moose. *The national geodetic survey gravity network*. Technical Report. Rockville, MD: US Department of Commerce National Oceanic and Atmospheric Administration, 1986.

[51] *MQTT*. 2020. URL: http://mqtt.org.

[52] *NAIP Imagery*. United States Department of Agriculture. 2019. URL: https://www.fsa.usda.gov/programs-and-services/aerial-photography/imagery-programs/naip-imagery/.

[53] *NSF NEON: Open Data to Understand our Ecosystems*. 2020. URL: https://www.neonscience.org.

[54] Nikolaos A Pantazis, Stefanos A Nikolidakis, and Dimitrios D Vergados. "Energy-efficient routing protocols in wireless sensor networks: A survey". In: *Communications Surveys & Tutorials, IEEE* 15.2 (2013), pp. 551–591.

[55] Kristofer S. J. Pister, Joe Kahn, and Bernhard Boser. *Smart Dust*. Proposal BAA 97-43. Berkeley, CA: UC Berkeley, 1997.

[56] Zach Poff. *Alex Rice Piezo Preamplifier*. URL: https://www.zachpoff.com/resources/alex-rice-piezo-preamplifier/.

[57] Edwina Portocarrero. "Networked Playscapes: redefining the playground". Ph.D. thesis. Massachusetts Institute of Technology, 2018.

[58] *Pycno Agriculture Sensors*. 2020. URL: https://pycno.co/.

[59] Vasant Ramasubramanaian. "Quadrasense: immersive UAV-based cross-reality environmental sensor networks". Master's thesis. Massachusetts Institute of Technology, 2015.

[60] P Krishna Rao et al. “Weather satellites: Systems, data, and environmental applications”. In: (1990).

[61] Sara E. Remsen. “Visualizing models driven by real-time, sensor-based data : a new way to experience the inner workings of ecosystems.” Master’s thesis. Massachusetts Institute of Technology, 2017.

[62] Alex Rice. *Piezo Preamp*. 2009. URL: https://web.archive.org/web/20100914091044/www.alexrice.co.uk/content/piezo-preamp.

[63] Spencer Russell, Gershon Dublon, and Joseph A Paradiso. “HearThere: Networked sensory prosthetics through auditory augmented reality”. In: *Proceedings of the 7th Augmented Human International Conference 2016*. 2016, pp. 1–8.

[64] Spencer Russell and Joseph A Paradiso. “Hypermedia APIs for Sensor Data”. In: *Mobiquitous: Mobile and Ubiquitous Systems*. London, 2014.

[65] Spencer F. Russell. “Resynthesizing Volumetric Soundscapes: Low-rank Subspace Methods for Soundfield Estimation and Reconstruction”. Ph.D. thesis. Massachusetts Institute of Technology, 2020.

[66] *Safecast*. 2015. URL: http://blog.safecast.org/.

[67] Ryuichi Sakamoto and YCAM. *Forest Symphony*. Art installation. Yamaguchi, Japan, 2013. URL: http://forestsymphony.ycam.jp.

[68] *Sensoterra*. 2020. URL: https://www.sensoterra.com/.

[69] Valentina Sumini. *Tidmarsh Living Observatory Portal*. Installation piece. 2019. URL: https://www.media.mit.edu/projects/tidmarsh-living-observatory-portal/overview/.

[70] Robert Szewczyk et al. “Habitat monitoring with sensor networks”. In: *Communications of the ACM* 47.6 (2004), pp. 34–40.

[71] Robert Szewczyk et al. “Lessons from a sensor network expedition”. In: *Wireless Sensor Networks*. Springer, 2004, pp. 307–322.

[72] *Teralytic - A Soil Health Company*. 2020. URL: https://teralytic.com/.

[73] *User Manual for the ML3 ThetaProbe*. Delta-T Devices Ltd. URL: http://www.dynamax.com/products/soil-moisture/th2o-portable-soil-moisture-probe.

[74] *USGS EROS Archive - Aerial Photography - High Resolution Orthoimagery (HRO)*. United States Geological Survey. 2016. URL: https://www.usgs.gov/centers/eros/science/usgs-eros-archive-aerial-photography-high-resolution-orthoimagery-hro.

[75] B. Warneke et al. "Smart Dust: communicating with a cubic-millimeter computer". In: *Computer* 34.1 (Jan. 2001), pp. 44–51. ISSN: 0018-9162. DOI: 10.1109/2.895117.

[76] Carl J Watras et al. "Evaluation of wireless sensor networks (WSNs) for remote wetland monitoring: design and initial results". In: *Environmental monitoring and assessment* 186.2 (2014), pp. 919–934.

[77] Geoffrey Werner-Allen et al. "Deploying a wireless sensor network on an active volcano". In: *Internet Computing, IEEE* 10.2 (2006), pp. 18–25.

[78] Joy B Zedler and Suzanne Kercher. "Wetland resources: status, trends, ecosystem services, and restorability". In: *Annu. Rev. Environ. Resour.* 30 (2005), pp. 39–74.