Brian Mayton
Gershon Dublon*
Spencer Russell
Evan F. Lynch
Don Derek Haddad
Vasant Ramasubramanian
Clement Duhart
Responsive Environments Group
MIT Media Lab
75 Amherst St., E14-548
Cambridge, MA 02142

Glorianna Davenport
Living Observatory

Joseph A. Paradiso
MIT Media Lab

The Networked Sensory Landscape: Capturing and Experiencing Ecological Change Across Scales

Abstract

What role will ubiquitous sensing play in our understanding and experience of ecology in the future? What opportunities are created by weaving a continuously sampling, geographically dense web of sensors into the natural environment, from the ground up? In this article, we explore these questions holistically, and present our work on an environmental sensor network designed to support a diverse array of applications, interpretations, and artistic expressions, from primary ecological research to musical composition. Over the past four years, we have been incorporating our ubiquitous sensing framework into the design and implementation of a large-scale wetland restoration, creating a broad canvas for creative exploration at the landscape scale. The projects we present here span the development and wide deployment of custom sensor node hardware, novel web services for providing real-time sensor data to end user applications, public-facing user interfaces for open-ended exploration of the data, as well as more radical UI modalities, through unmanned aerial vehicles, virtual and augmented reality, and wearable devices for sensory augmentation. From this work, we distill the Networked Sensory Landscape, a vision for the intersection of ubiquitous computing and environmental restoration. Sensor network technologies and novel approaches to interaction promise to reshape presence, opening up sensorial connections to ecological processes across spatial and temporal scales.

1 Introduction

Landscape captures the complex exchange between the world we see, the world we make, and the world we imagine. These worlds are often in tension, and perhaps no human endeavor captures this tension more than our pursuit of technology, the most significant driver of our impact on the environment. At the same time, technology provides our primary means of understanding the environment, preserving it, and expressing our relationship to it through art—from cave paintings to audio recordings of melting sea ice (Helmreich, 2016).

What is the role of ubiquitous sensing in the future of landscape? Since Szewczyk et al. (2004) demonstrated the potential of wireless sensor networks (WSNs) as research tools in habitat monitoring, systems like theirs have been used with increasing frequency in primary ecological research and for

Presence, *Vol. 26, No. 2, Spring 2017,* 182–209
doi:10.1162/PRES_a_00292

---

*Correspondence to gershon@media.mit.edu.

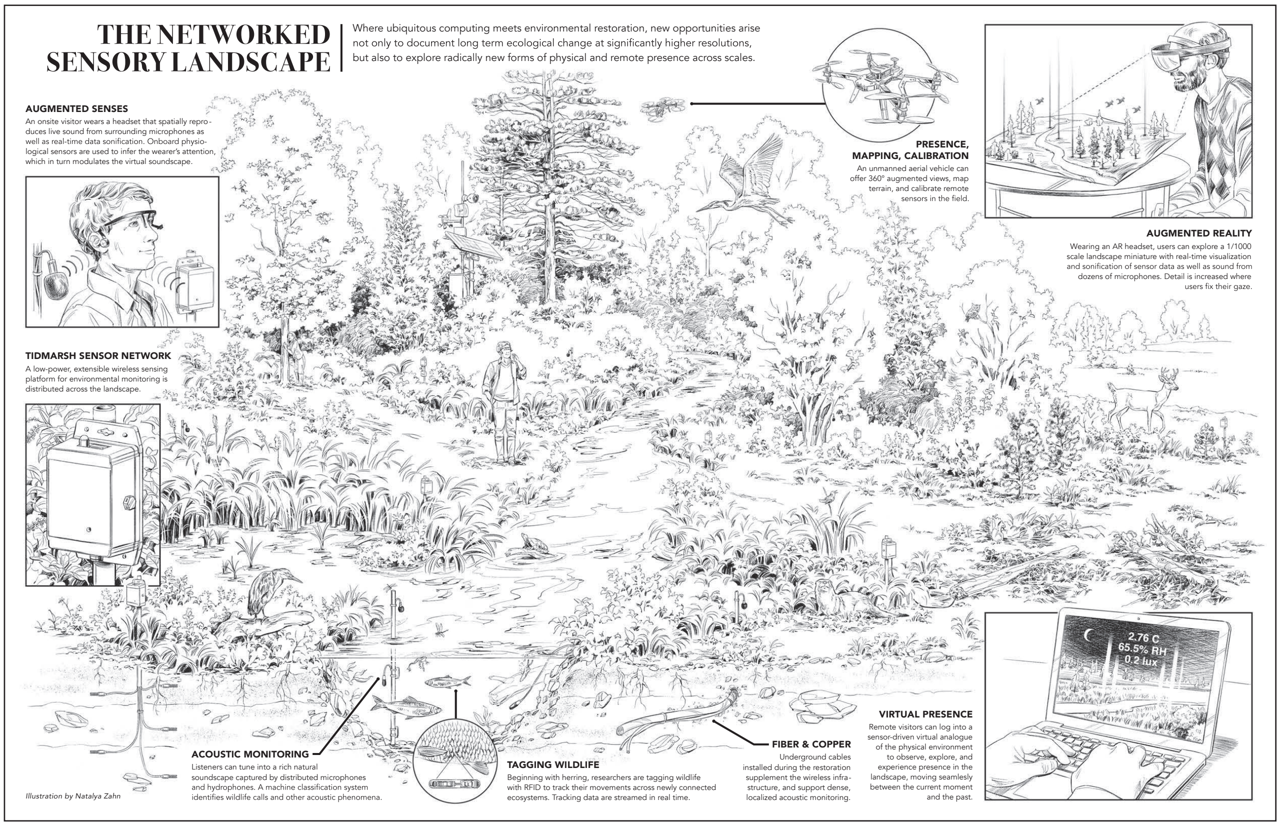

Figure 1. *In this article we introduce our vision for a landscape future at the intersection of ubiquitous computing and environmental restoration. The Networked Sensory Landscape presents opportunities to capture long-term ecological change at higher spatial and temporal resolutions than ever before, and to experiment with radically new forms of physical and remote presence across scales.*

conservation (Hu et al., 2009; Lloret, Garcia, Bri, & Sendra, 2009; Rundel, Graham, Allen, Fisher, & Harmon, 2009; Watras et al., 2014). More than a decade on, in the era of mobile and ubiquitous computing, we are finding that environmental sensor networks embedded in the landscape can serve as a platform for a wide array of applications spanning research, outreach, and art. In general, our work has been focused on the intersection of presence and pervasive sensing, and landscape is a natural site for this broad line of research (Dublon & Paradiso, 2014).

In this article, we call this field the *Networked Sensory Landscape*, and present our own work developing an environmental sensor network along with its diverse applications and outcomes. What opportunities are created by weaving a continuously sampling, geographically dense web of sensors into the natural environment, from the ground up? And how can the data produced by this kind of network extend an enhanced sense of presence to both onsite and remote visitors? Figure 1 captures our vision for the techno-ecological landscape in which this inquiry is rooted, and which we see as a model to carry forward to other sites.

Our field laboratory for this work, called Tidmarsh, is currently the site of the largest freshwater wetland and riverine restoration in the state of Massachusetts. Formerly an industrial-scale cranberry farm, Tidmarsh has recently undergone restoration (Interfluve, 2015) with the goals of reestablishing natural processes, such as the free flow of water, and ecosystem services, such as biodiversity support. At the time of this writing, water control structures have been removed and ditches filled; three miles of new sinuous stream channel have been constructed and connected to a new pond, and riffles have been built to raise the water table. Finally, the flat bog surfaces have been sculpted with microtopography and many tons of fallen trees have been distributed to create a wealth of new microhabitats across the site. Spring 2017 was the nascent wetland's first. In Fall 2017 the land became a public wildlife sanctuary (Massachussetts Audubon Society, 2017). In addition, a nonprofit organization, the Living Observatory (LO), has been formed to bring together the scientists, artists, and wetland restoration practitioners working on the land and engaging through it (Bidgood, 2017).

Over the past four years, as the restoration has gone from the design phase into active construction, we have incorporated our ubiquitous sensing framework at each stage of the process, ultimately spanning sensor nodes, a generalized real-time sensor data API, and novel user experiences. Our continuously evolving sensor network has been in place for more than three years documenting the prerestoration environment, and will continue as nature takes its course. Its data, in conjunction with other, more targeted data collected by our research partners, are available to scientists, restoration engineers, and land managers. But beyond the environmental science and local environmental benefits of the project, the partnership behind the restoration is seeking to understand how future park visitors will interact with this new kind of landscape model. For this reason, our mandate has extended beyond sensing, both to constructing infrastructure and developing user interfaces that bring manifestations of the sensor data to the public. A significant portion of the underlying ecological change is invisible to the naked eye—some of it too slow for us to witness, some passing in the blink of an eye, and much of it occurring where we're not looking. Can we build technologies that allow us to witness, enjoy, and examine landscape-scale change in new ways?

This article will introduce the concept of the Networked Sensory Landscape as we have constructed it at Tidmarsh. Like any end-to-end sensor system, ours has three layers (collection, back-end, and user interface) but in our case, each part is designed with this concept at its core. The sensors themselves and network infrastructure were built for both scientific inquiry and open-ended user experience, the back-end was built to support almost any kind of real-time, user-facing application, and the diverse set of interfaces we present here were built to encompass the myriad ways the public can engage with landscape.

The next sections of the article will detail these layers from the bottom up. In the following section, we describe the sensing layer, which combines a low-power environmental WSN, dense collection of audio from wired and wireless microphones and hydrophones, and wired network cameras. Next, we introduce ChainAPI, a hypermedia web service that links the different parts of our system together, focused on data persistence, a common protocol to link our disparate systems, and the data plumbing required to do so. Finally, we detail the applications and user interfaces that bring the sensed landscape to both researchers and the general public; these include traditional visualization systems, virtual and augmented reality (VR/AR) experiences mirroring the physical landscape, wearable devices, and tools for creative expression—enabling landscape as canvas.

2 Sensing the Landscape

To construct the Networked Sensory Landscape, we have densely instrumented Tidmarsh with numerous data sources and the infrastructure required to make all of the captured data available via the internet in real time. A network of custom low-power wireless sensor nodes provides spatially and temporally dense measurements of basic environmental parameters, and an array of microphones provides many channels of live audio. These two main sources are supplemented by additional data streams in a few locations, such as live camera images. The data sources are concentrated in several different locations on the site, chosen to capture a cross-section of the different conditions and types of habitats across the property. The following section describes each of the instrumented sites and the network infrastructure that connects them. Next, our sensor platforms are introduced, followed by our audio streaming setup. Finally, we describe an example of processing data

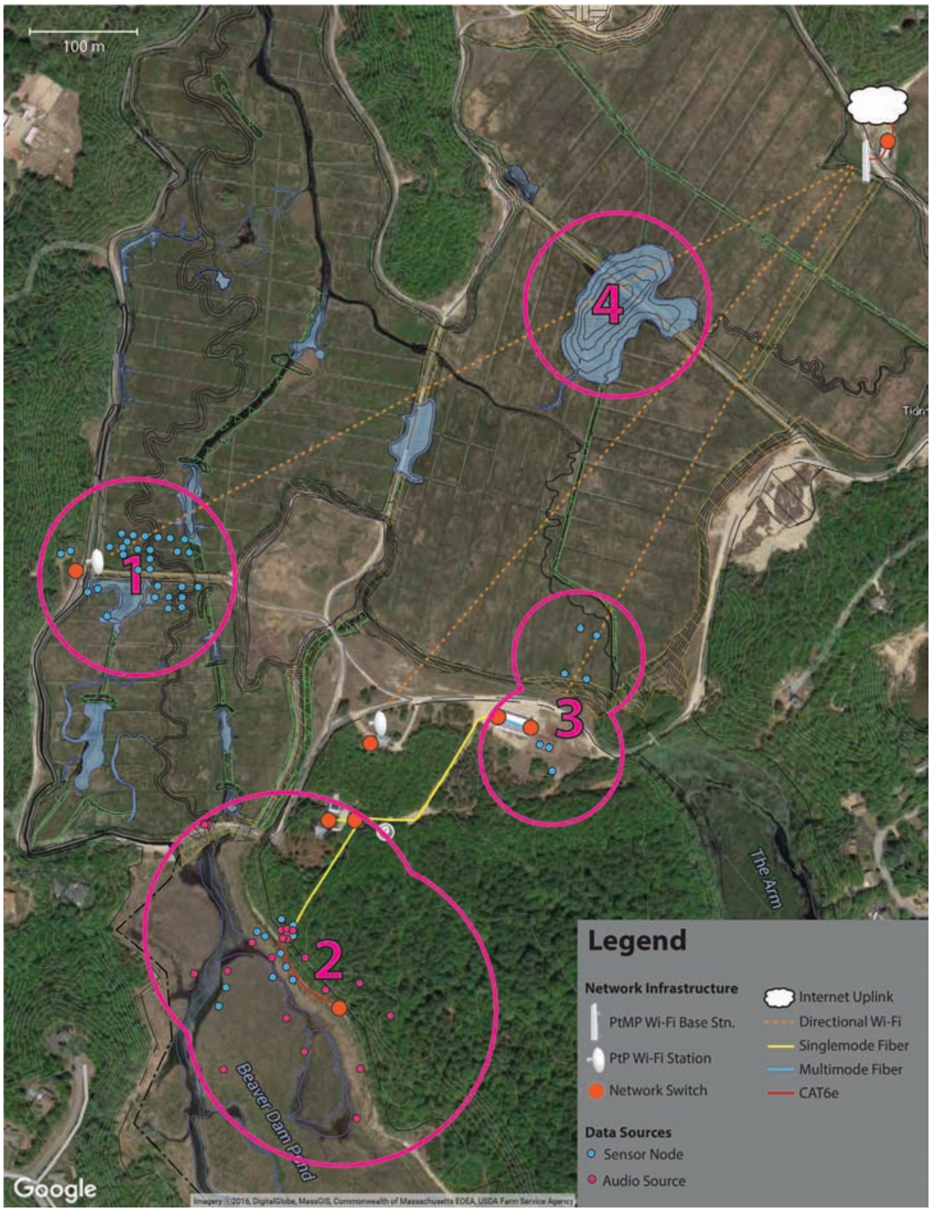

Figure 2. *Tidmarsh network infrastructure and sensor sites.*

streams in real time to produce higher-level data sources that can be used by applications.

2.1 Sensor Sites and Network Infrastructure

Figure 2 shows most of the 600-acre Tidmarsh property, highlighting the locations that we have instrumented. Sensing is deployed across three active sites (numbered 1–3) and will be added to a fourth planned site (4) in the coming spring. These sites are connected together by a low-latency, high-bandwith internet protocol (IP) network, enabling real-time data streaming from the sensors, microphones, and cameras. The IP network also provides internet connectivity for visitors to the site, enabling them to interact with the sensors through portable and wearable devices, including several that will be described later in this article.

Internet connectivity comes in from a local internet service provider at the northeast edge of the property (indicated by the cloud symbol in Figure 2). A barn at this location houses the head-end infrastructure, including the main router and on-site server. A 5.8-GHz 90° sector antenna mounted outside the barn creates a TDMA Wi-Fi network that serves as the wireless backhaul for the property; this network is accessible using inexpensive hardware anywhere with a direct line-of-sight from the bog surface.

The first sensors were installed at the west edge of the property (site 1 in Figure 2 or the "west side") in February 2014, and is now the largest sensor installation at Tidmarsh. This site receives connectivity via the wireless backhaul, and the base station is powered by two 100-watt solar panels and a 2,400-watt-hour battery. The base station also includes a network camera and a stereo audio streaming setup, as well as a weather station measuring wind speed, direction, and rainfall. This site has been planted with Atlantic white cedar seedlings during the summer of 2016, and will become a forest over the coming decades. Sixty-four sensor nodes were in place from October 2014 through August 2015; all but a small handful were then removed to allow the construction and microtopography work to take place. As of autumn 2016, most of the sensor nodes have been re-installed.

The southern end of the property (site 2 in Figure 2, or the former impoundment) was once an artifical pond that provided a source of water to flood the cranberry bogs. The dam impounding the water was removed several years ago and minimal construction work has been done on this part of the site during the restoration. As such, it represents a portion of the property that is several years ahead in its process of transformation. It contains another large sensor node deployment, as well as an extensive array of microphones that stream live audio (see Section 2.3). As there is no line-of-sight to the network head end, connectivity is provided by buried fiber optic cable running through the forest. The fiber connects to the wireless network at the barn near site 3. This site also has wired power that joins the buried fiber at a house partway along the route. This site also includes a camera that has been capturing images for the past year.

Another site along the edge of the bog surface (site 3 in Figure 2, the "greenhouse and test plot") will include sensing in three different environments. A nearby greenhouse contains several sensor nodes and is an example of a controlled environment that has been used to raise seedlings for planting. The nearby bog surface, also in range of the network, will become transitional fen, and sensor nodes will be installed once the restoration work is complete. The nearby bog surface also contains a rectangular test plot where ecologists will conduct soil experiments; this plot has been instrumented with four sensor nodes, each with soil moisture and soil temperature probes at two different depths, in addition to the standard set of sensors inside the node itself. Network connectivity for this site comes from the wireless backhaul, and it has wired electricity from the nearby barn.

A fourth site (site 4 in Figure 2, the "constructed pond") is underway as part of the summer 2017 deployment. At this location, a large pond has been constructed. Despite being less than a year old, the pond is already buzzing with natural activity.

2.2 The Sensor Network

Wireless Sensor Networks have a history going back over a decade, using a variety of hardware and networking protocols. Many projects from the early 2000s, such as the work on Great Duck Island (Szewczyk et al., 2004), made use of Berkeley's Mica motes (Hill & Culler, 2002) and TinyOS (Levis et al., 2005). Several sensor node designs are now commercially available, including some that have grown out of research work (*MEMSIC, Inc—Wireless Sensor Networks*, 2015) and new turnkey product designs (*Libelium—Connecting Sensors to the Cloud*, 2015). While the appearance of these products on the market is exciting for the future of WSNs, we determined that the cost and specialization of existing systems were not compatible with our needs. With a diverse landscape under rapid transformation, we needed a network that was flexible to deploy, customize, and alter as the construction work progressed. In order to support a spectrum of different applications, from traditional ecological research to experiments in presence and perception, frequent sampling and real-time transmission were important considerations. In many cases the requirements of various applications were not well-defined at the time the network and infrastructure were installed. While the engineering costs of developing our own platform have not been insignificant, the lower per-node cost (approximately $80 per assembled node) and flexibility to iterate on the complete hardware and software stack to meet the requirements of different locations and applications as the site changes over time have been of great benefit.

Our low-power wireless sensor network provides general data about the environment at Tidmarsh. The network is built from custom sensor nodes, each consisting of a microcontroller, radio, power source, and sensors in a weatherproof enclosure. This platform is designed to be easily deployed in large numbers across the areas we have chosen to instrument, and the complement of internal sensors provide a standard baseline for the sensing across the site. Expansion capabilities allow sensor nodes to measure additional parameters where it makes sense to do so, allowing many different types of probes to be connected to the network. Two generations of sensor hardware are described in this section. The first-generation sensors, developed in 2013, has been used for most of the sensing to date. The second-generation sensors, preparing for manufacturing at the time of this writing, will feature an expanded set of internal sensors and several improvements that will increase the lifetime of the nodes in the field. The second-generation sensor nodes are expected to be ready for deployment during the summer of 2017.

The first generation of our sensor node platform is built on an 8-bit ATxmega128A4U processor and AT86RF233 2.4-GHz 802.15.4 radio for communication. The on-board sensors include temperature and humidity (SHT21), light (ISL29023), atmospheric pressure (BMP180), and an accelerometer (ADXL362). The accelerometer was added with local wind sensing in mind if the sensor node is installed on a flexible mount; however, this has not yet been thoroughly explored. The PCB and batteries (shown in Figure 3) are designed to fit into a small weatherproof enclosure with ports to expose the sensors to the environment. The assembly mounts to stakes made from PVC pipe that can be rapidly installed in the field and will not decay over time (see Figure 5). The node can be extended through a waterproof connector that connects to an expansion box providing four 12-bit analog input channels and programmable excitation voltages to power the sensor probes.

Figure 3. *The second-generation hardware inside the waterproof enclosure.*

The second generation of the sensor node platform aims to improve the sensing capabilities and extend the lifetime of the node. The new device, diagrammed in Figure 4, uses the same MCU and radio, and most of the same sensors. The single light sensor has been replaced by two digital sensors (VEML6030, providing high dynamic range visible light, and VEML6070, providing separate UVA and UVB measurements). A pair of photodiode amplifiers and 12-bit A/D converters allow two additional photodiodes to be installed, which will be used for IR sensing. Different plant and bacteria species can make use of different wavelengths for photosynthesis; separate measurements of these spectral bands may provide insight into what grows where. A passive IR motion sensor enables the detection of large wildlife, and a microphone with an integrator, peak detector, and FFT capability will allow analysis of audio around the node, including wind noise. Integrated flash memory storage will allow sensor data to be stored

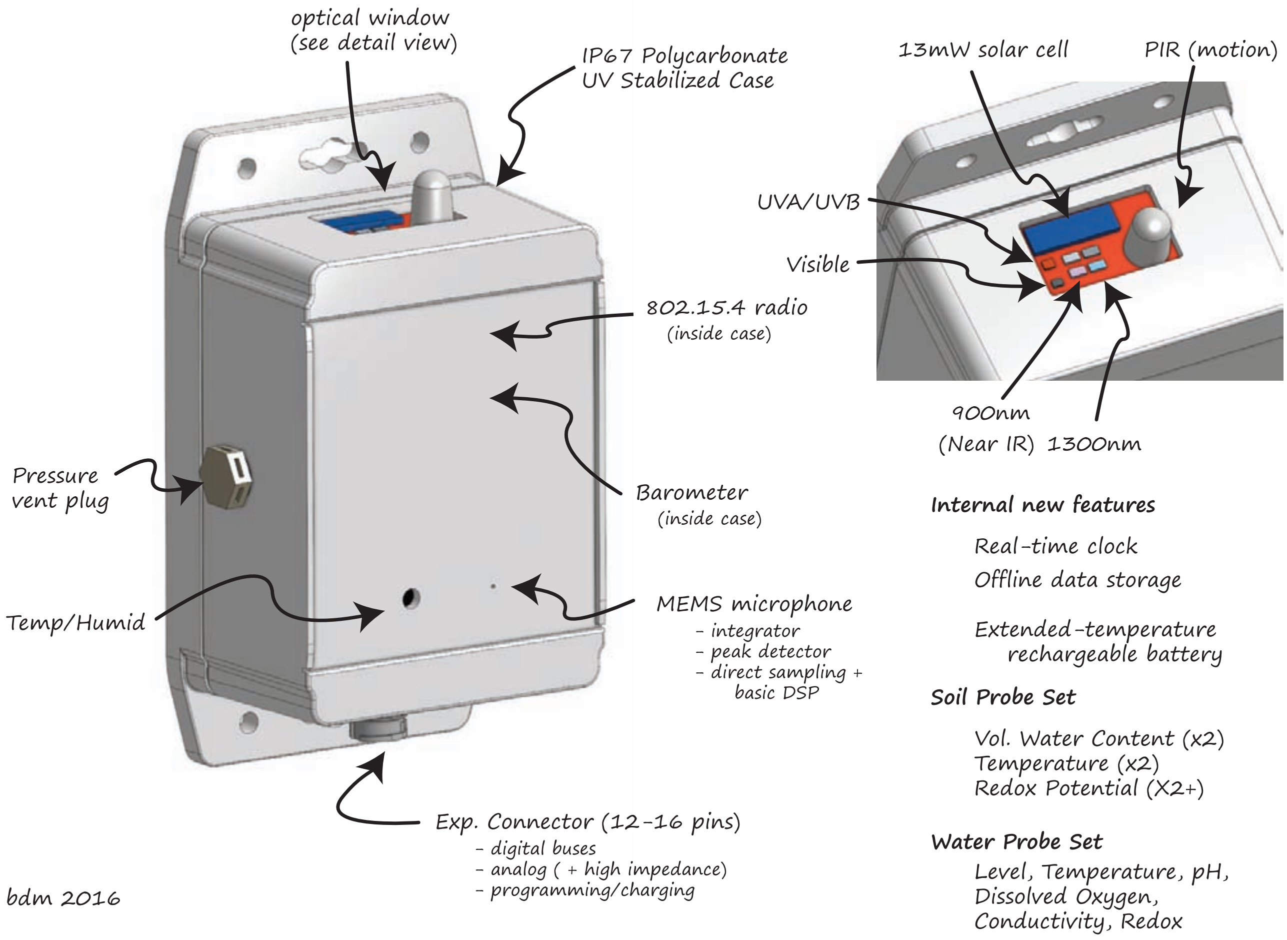

Figure 4. *Second-generation sensor node hardware diagram.*

internally and retrieved later so that data are not lost during network outages. A small solar panel, rechargeable battery, and energy harvesting controller will allow for much longer operation before battery replacement will be required—an important consideration as the site reverts to nature and the nodes become more difficult to access. The second-generation nodes will have the expansion capability built in, allowing external probes to be easily connected without additional electronics. Initial prototypes are currently being evaluated, as shown in Figure 5.

The sensor nodes communicate using a modified version of the Atmel Lightweight Mesh protocol, which is built on top of the IEEE 802.15.4 standard. The protocol provides basic multihop routing support while maintaining a low memory footprint, and is simple to extend and debug. Our extension includes an efficient binary sensor data format that identifies each sensor metric present in the message so that the data can be properly decoded when received by the central server. The network protocol also provides several control commands and an over-the-air firmware update mechanism so that updates to the sensor node software can be installed in the field, which has proven useful as sensing requirements change.

The sensor node has been designed with low-power operation in mind. The first generation nodes are powered by three primary AA cells, while the second generation nodes will have a single rechargeable lithium ion cell. The present generation of the hardware has a quiescent draw of about 25 μA. A custom real-time operating system keeps the node in this low-power mode

Figure 5. *Left: first-generation sensor node in its enclosure, installed in the field. Right: second-generation sensor node prototype, with inset showing optical components on top of the node (detailed in Figure 4).*

most of the time, with only a low-power 32-kHz RC oscillator operating the scheduler task. Every 30 seconds, the scheduler wakes up and triggers all of the sensors to start measurements (the sensing interval is programmable via remote commands). As the A/D conversion for most of the sensors takes some time, the node goes back to sleep for the next second, then wakes up to read the results, assemble a wireless frame, and transmit. The node consumes approximately 2 mA when reading the sensor values and 16 mA while transmitting. The second-generation sensor node reduces the quiescent power draw significantly by correcting several design issues present in the first-generation node.

Soil monitoring will constitute a significant part of our external sensing. Soil hydrology, pH, oxidation-reduction potential, and temperature (Vorenhout, van der Geest, & Hunting, 2011) aid in evaluation of the site's efficacy in providing important ecosystem services (Zedler & Kercher, 2005) and provide a window into normally invisible underground processes. Other external probes include weather stations (monitoring wind speed, wind direction, and rainfall), water quality (dissolved oxygen, pH, conductivity, temperature, redox potential, water level), and air quality.

The low-power wireless sensor network is bridged to the IP backhaul at each site via a gateway node. The gateway platform is an Intel Galileo single-board computer with a custom radio connected to the Galileo's expansion headers. The radio has an external 24-dB power amplifier, LNA, and RF switch with automatic antenna diversity. Software running on the Galileo provides the network discovery and firmware update services, and connects to the central server (running at the Media Lab) through ZeroMQ.

2.3 Audio Streaming

Audio has several qualities which make it compelling as a data source, especially when synthesizing a sense of presence. First, it captures a huge amount of information about the site, including weather conditions and the activity of wildlife, which may be a good indicator of the overall health of the site. It is relatively inexpensive to capture and transmit in high fidelity, compared to other rich data streams such as video. It also possesses a flexibility that video images do not; while a camera presents a single point-of-view through a lens, audio can be spatialized and presented in a partially virtual environment without feeling overly unnatural.

At Tidmarsh, we have implemented audio streaming in a number of different ways, depending on the availability of power and network access at each location. One location (site 1 in Figure 2), which does not have wired power or network, uses a solar-powered computer and audio interface to capture and encode two channels of audio and transmit the result via Wi-Fi. At the former impoundment location (site 2), we have taken advantage of the wired power and network to implement a much larger installation. A rackmount 32-channel mixer acts as an audio interface with DSP (filtering and dynamics processing on each channel reduce the effects of wind noise) and is connected to an x86 computer, colocated with the base station for the wireless sensor network. The computer encodes the audio using the Opus codec and transmits the streams to our central server at the Media Lab.

Sixteen of the inputs are located on the mixer. Cables run underground and connect to microphones up to 500 feet away. Another 16 inputs (and corresponding microphones) are located at a satellite box 300 feet south of the base station box, connected via CAT6 and power cable. As of this writing, there are 24 channels actively being streamed from this installation (including a hydrophone). Microphone locations are shown as pink dots on Figure 2.

As most commercial microphones are unsuitable for installing outdoors where they are exposed to the elements, we have developed our own weatherproof omnidirectional microphones (see Figure 6). These utilize low-noise electret capsules and a circuit for buffering the microphone signals and driving a differential cable. The assembly is protected inside of an aluminum tube and potted with silicone rubber to protect it from water, leaving only the front of the capsule exposed. Foam windscreens placed over the front of the tube reduce wind noise and help repel water from the exposed capsules. The microphones interface directly to standard professional audio equipment with balanced XLR connectors and 48-V phantom power. Several of these microphones have operated in the field for more than two years without issue.

The audio streams are published using the Icecast protocol, commonly used for internet streaming audio. This enables the streams to be directly used in a variety of contexts, such as embedded in a web page or in a virtual representation of the site (see Section 4.7.2). The streams from the former impoundment location are published as a single multichannel stream, which maintains synchronization between the individual microphones, as well as a stereo mix that can be received without specialized software (e.g., on our web site). Various parts of the audio installation, including microphones, the main electrical box, and the satellite box, are shown in Figure 6.

2.4 Virtual Sensors: TidZam

Data sources may also be processed to create additional streams of data useful to applications. These "virtual sensors" may aggregate data from physical sensors, or represent some more abstract data source. TidZam is an example of a virtual sensor that analyzes our microphone audio streams in real time to analyze outdoor ambient sound, identify wildlife, and characterize biodiversity over the course of the restoration.

Bird call recognition has been studied for some time with a number of different approaches (Brandes, 2008), with partial success demonstrated on small subsets of bird species. Bird calls contain complex structures analogous to the phonemes, syllables and even sentences found in human language, which are varied even among the same species. As such, traditional classification

Figure 6. *Parts of the audio streaming installation. Top left: microphones awaiting installation. Bottom left: satellite input box. Center: microphones in trees and in the marsh. Right: main electrical box, showing mixer (top), computer (center), network switch and fiber (bottom left).*

approaches tend to fail because it is difficult to manually select features. Recently, researchers have demonstrated major improvements in bird call recognition using algorithms for unsupervised self-extraction of features (Stowell & Plumbley, 2014).

TidZam is built using convolutional neural networks for self-extraction of hierarchical feature representations that describe the spectral content associated with qualitatively different types of outdoor sound (e.g., rain, wind, frogs, birds, human voices, aircraft, etc.). An expert architecture is used to analyze incoming signals and forward them to more specialized classifiers, which produce the final classification (Jacobs, Jordan, Nowlan, & Hinton, 1991). The classifiers are trained using nonexclusive output classes, which allow class overlapping on the same signal sample, such as birds calling during a rainstorm. The system also generates an ambiguity measure, used for labeling unfamiliar samples or indicating that a final classification (e.g., bird species detection) cannot be made, even if an initial classification (e.g., generic bird detection) has been produced. In this case the signal is extracted dynamically so that an expert may determine the source of the audio (crow, blue jay, frog, rain, etc.) to build a new class or improve an existing one. The framework's interactive interface provides a training feedback mechanism between users/experts and the neural system in order to improve the knowledge of the system and the users. The resulting classifiers form "virtual sensors" associated with the microphone where the audio was recorded. The TidZam web interface is shown in Figure 7.

There remain several open problems, the most pressing of which is developing a process for automatically evaluating system performance on a large scale. The volume of data is significant, and a comparably large ground truth dataset does not yet exist. System performance on hand-labeled data is continuously measured to improve the classification. At the same time, the database of unsupervised labeled examples is being continuously grown as new detections are made and

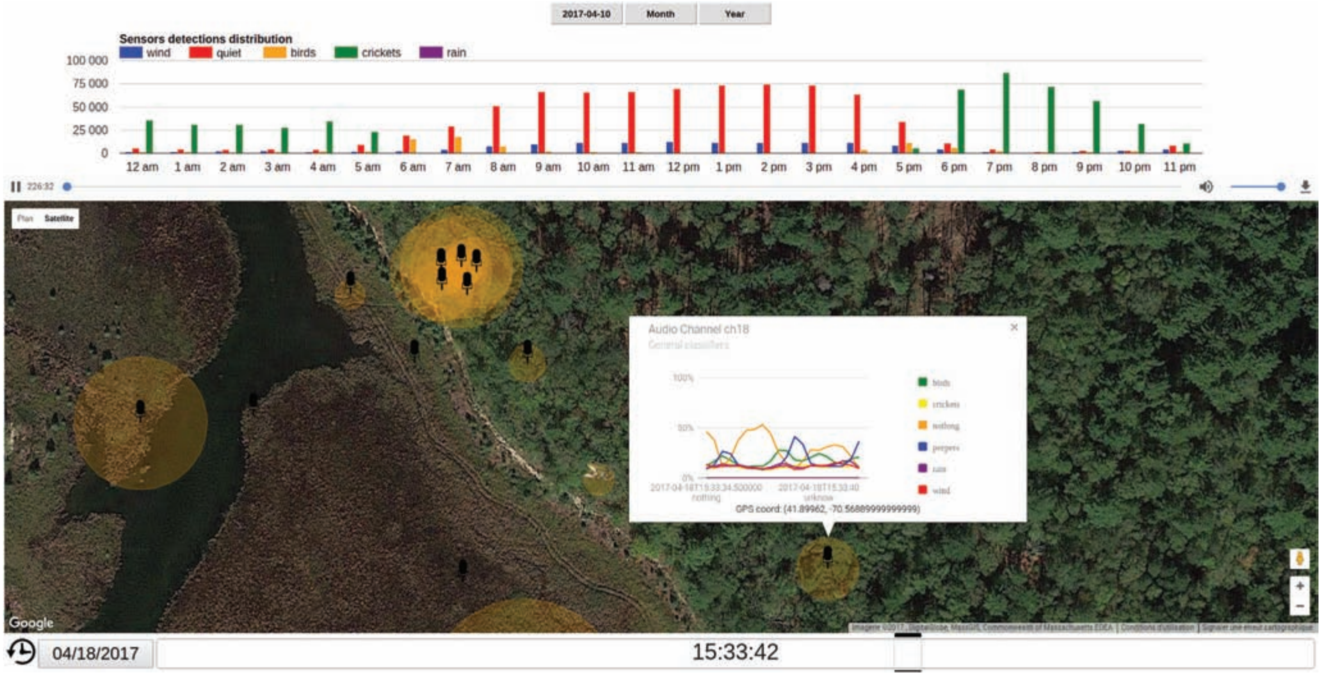

Figure 7. *The TidZam web interface visualizes density of clearly identified outdoor sonic events, such as wildlife and weather conditions, in space and over time. Not shown are labels for "unknown" samples, which are ambient sounds not recognized by the classifiers.*

fed back into the system, and hundreds of ambiguous samples are forwarded to a human labeler each week. This allows us to test new classifiers on recorded data to track their performance as the database grows. Our database currently contains 42,000 classifications, with a system-enforced uniform distribution across the six primary classes (birds, crickets, peepers, rain, wind, ambience/nothing). At the time of writing, the higher-level classifier for bird species is composed of 58,000 examples distributed over 20 species classes, almost entirely originating from the audio streams at Tidmarsh.

3 Linking the Landscape

Distributing data from a heterogeneous collection of sensor systems to many different end-user applications and intermediate systems can create an exponential explosion of integration interfaces. Deploying new sensors or modifying services can require a complex set of changes in multiple pieces of software. To address this concern we built ChainAPI (Russell & Paradiso, 2014), a hypermedia web service that serves a common language linking together the parts of our system. ChainAPI stores sensor data in a database and also acts as a plumbing layer to connect sources of data to sinks in real-time.

Earlier experiments within our group (Dublon et al., 2011) used individualized back-end systems to store sensor data and make it accessible to client software, but after making several of these systems, we started to recognize common patterns and use cases that drove the design of ChainAPI. We based our design on an underlying set of architectural principles first articulated by Roy Fielding known as Representational State Transfer, or REST (Fielding, 2000), which have become foundational ideas in the structure of the Web. Others (Guinard, Trifa, & Wilde, 2010) have begun to explore how these same principles can guide the design of systems for talking to sensors and connected devices, often termed the "Web of Things," to highlight the place of these protocols as an abstraction layer on top of the transport provided by the "Internet of Things."

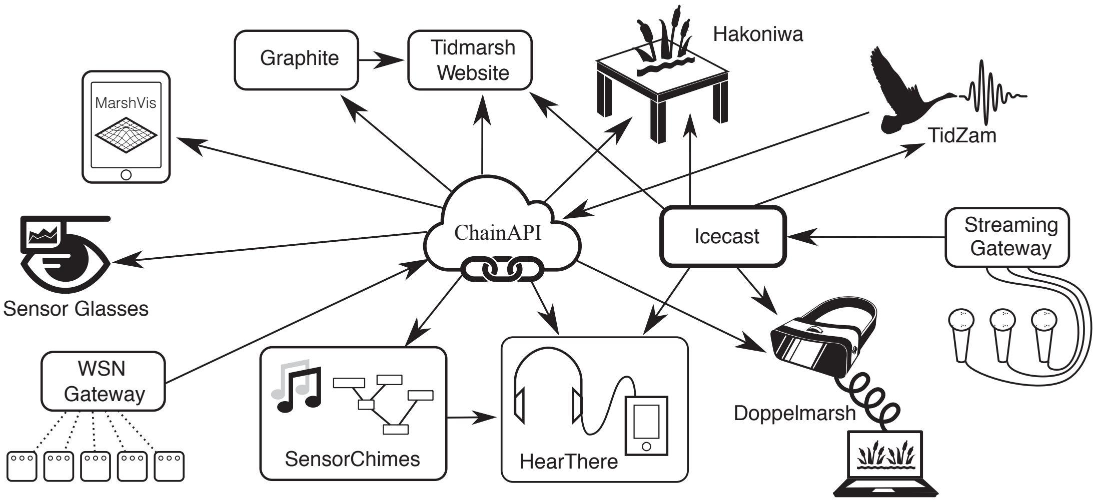

Figure 8. *Data from sensors and other processes flows into ChainAPI, where it is distributed via HTTP and WebSockets to our end-user applications. The Icecast server software distributes Ogg Opus and MP3 audio streams.*

Following the REST architecture, we use hyperlinks as a central organizational principle, allowing the server to add functionality (new link types) without breaking older clients, and also allowing clients to discover what operations and related data are available on any given resource. A device that has a real-time stream available will have a stream link. This hyperlink-oriented architecture enables an ecosystem of search engines and content aggregation crawlers to exist on top of the underlying link substrate, similar to the services that have grown to index the Web.

Our architecture is influenced by Semantic Web technology such as SPITFIRE (Pfisterer et al., 2011) and OpenIoT (Soldatos et al., n.d.), but with a focus on accessibility with mainstream Web technologies including HTTP, WebSockets, and JSON rather than requiring clients to use Semantic Web-specific tooling or conform to complex ontologies. For example, new sensors, or sensor data are added to the system with an HTTP `POST` request of JSON-encoded data. This also contrasts with other IoT standards that focus on end-to-end protocols all the way to the node level, rather than assuming a diversity of protocols at the network edges. COAP (Shelby, Hartke, Bormann, & Frank, 2014) is inspired by HTTP but was redesigned by the IETF to be more appropriate for resource-constrained devices. MQTT (Locke, 2010) is a publish–subscribe model designed for small payloads. Both of these could coexist with ChainAPI with an adapter service.

Our implementation is available under an open-source MIT license and published on GitHub (https://github.com/ResEnv/chain-api).

In Figure 8 we see a wide variety of clients that pull data from ChainAPI, and two sources of data (the Tidmarsh sensor network and animal presence information from the TidZam classifier). The clients shown here are on a diverse set of platforms, including the Unity3D game engine on desktop and mobile, node.js web applications, Python scripts pulling the data into visualization tools, and browser-based JavaScript web applications. To support these applications we have developed client libraries in JavaScript, Python, and C#, each of which is used by multiple applications. In addition to requesting current or historical data over HTTP, clients can use ChainAPI to be notified of new or modified data using WebSockets. These notifications use the same JSON message format as the HTTP payloads. This can be used as a mechanism to set up data processing pipelines. For instance, the sensor data plots on our main website are created by a small script that subscribes to sensor data from ChainAPI and sends it to Graphite, a widely used server application for tracking and plotting data.

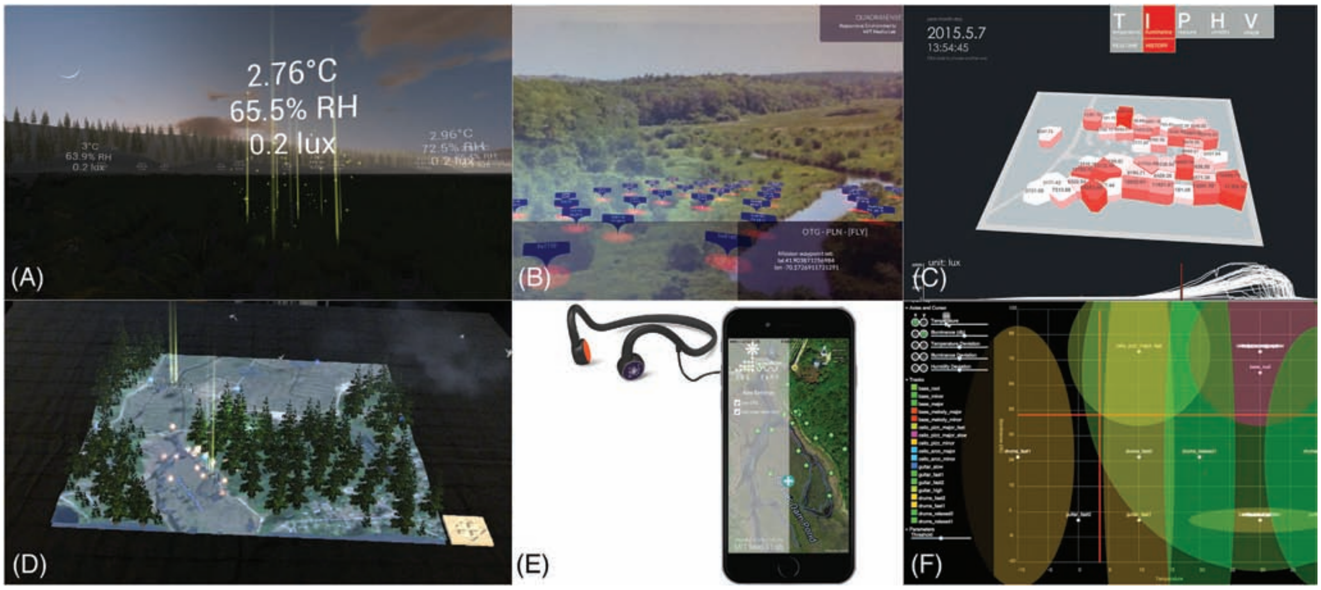

Figure 9. *User Experiences of the Network Sensory Landscape: (A) Doppelmarsh, a sensor-driven virtual world parallel to the physical environment (Section 4.7.2); (B) QuadraSense, a mixed reality view through the lens of a UAV-borne camera (Section 4.7.3); (C) MarshVis, a web-based 3D data visualization, shown here with spatial distribution of illuminance measurements (Section 4.7); (D) Hakoniwa, a miniature living landscape presented on a tabletop through AR glasses (Section 4.7.4); (E) Bog Ears, a wearable auditory AR extension of existing hearing (Section 4.7.5); (F) Sensor Chimes, a musical framework for composing the real-time, location-based sensor data sonifications used by our AR and VR systems (Section 4.5).*

4 Experiencing the Landscape

When we enter a space, some aspects of the environment are obvious, but many phenomena remain imperceptible because we do not have appropriate biological sensors to detect them, they are too large or small, they change on timescales that are too long or short, or they are beyond our reach. How “present” we are in an environment relates not to how much we know abstractly about the environment, but to how much we empathize with it. The modern world is increasingly documented not only by our writing, recording, and collective memory, but by the many sensors that are embedded in ubiquitous devices. Modern sensor technology allows for efficient collection of these data at a large scale.

The concept of a Networked Sensory Landscape is based in our thinking about the ways in which ubiquitous sensing can extend human perception of the environment across spatial and temporal scales. We have explored a diverse set of models for user experience of the sensed landscape, reflecting the range of new opportunities afforded by an open-ended approach to how environmental sensing can leverage ubiquitous computing. Overall, our focus has been on interfaces that provide both onsite and remote visitors with an extended sense of presence through the sensing layer on the landscape.

This section presents the software systems underpinning our approach and shared across the interfaces and tools we have developed. The last part of the section presents the interfaces themselves, catalogued in Figure 9. Among them Doppelmarsh, a virtual world parallel to and driven by sensors in the real one, and QuadraSense, which leverages Doppelmarsh to create first-person mixed reality views through the ultra-wide angle lens of an unmanned aerial vehicle (UAV). On a much smaller scale, Hakoniwa, or Boxed Garden, is a miniaturized living marsh landscape presented in bird’s eye view on a table through an AR head-mounted display. For extended perception onsite, we developed HearThere, a precise head-tracking bone conduction headphone system for auditory AR experiences that overlay sonification and live audio onto the physical environment. Real-time sonifications for our various displays were developed by commissions to outside composers, who designed them using our sensor data sonification and composition tool, called SensorChimes.

Some precedents for our work can be found in the HCI and embedded systems literature around general purpose sensor network user interfaces. Among these, Marquardt et al. (2010) presented the "Visual Environment Explorer," a tool for assessing the states of multiple networked devices and exploring the network in various ways. Michel et al. (2009) demonstrated an end-to-end environmental sensor network and visualization tool comprising environmental monitoring weather stations and various forms of graphical display, including 3D contour map overlays; in its focus on the data plumbing as well as the map-based, physically linked visualization, their system is similar to several of our applications. However, in contrast to these examples, our aim is for *data visceralization* over compact representation (Dobson, 2016). In its encoding in music and abstract forms and presentation in virtual and augmented realities, information in our interfaces can be slower to digest. Our approach privileges user exploration over targeted queries, more akin to a nature walk than a Google search.

To support more open-ended exploration of sensor networks than traditional user interfaces would allow, we proposed Cross-Reality Environments, where ubiquitous sensor networks would interact with pervasively shared virtual worlds (Lifton et al., 2009). In the intervening years, we have focused on evoking experiences of presence by tying physical environments and sensors to immersive virtual counterparts using 3D visualization and spatial sonification. An early example of this work, DoppelLab, is a cross-reality interface to the real-time and recorded data produced by a modern building and its occupants (Dublon et al., 2011). In DoppelLab, representations of sensors are bound to architectural space in a game engine, while users' movements are unconstrained by physical rules, and time is a dynamic parameter (to be traversed, stretched, and compressed). Virtual sensors can be represented in composite visualizations, and detail can be parameterized (representations are zoomable).

Cross-reality environments have allowed us to explore sensor-driven worlds tied to landscape but unencumbered by physical or temporal constraints. Through this process, we have developed a core set of software components, shown in Figure 10, on which most of our user-facing projects depend. Whether targeting traditional devices, AR, or VR, our applications share components that provide location services, produce 3D visualization and sonification, render spatial live audio, process real-time data streams, and include static data resources such as maps, terrains, and other physical datasets (ground penetrating radar data, archived audio, thousands of user-submitted photos, etc.). Sections 4.1 through 4.6 detail these components, and Section 4.7 presents the applications and devices that use them to enhance users' presence in the sensory landscape.

4.1 Static Resources

Our UI applications make use of a large database of recorded data, which we refer to as static resources. These data can be statically linked with the applications or downloaded at runtime from ChainAPI and other file servers. They include recorded audio files, manually acquired landscape datasets such as ground penetrating radar and topographic maps, photographs, and more. For example, our VR applications directly render the physical landscape using LIDAR measurements collected by the United States Geological Survey (USGS).

In another example, our systems access prerecorded audio files from a place-based audio documentary system called Roundware (http://roundware.org). Developed and curated by Halsey Burgund, Roundware allows us to download location-tagged, user-submitted recordings that can be spatially rendered through our applications in AR and VR.

4.2 Real-Time Data

Our applications rely on ChainAPI to keep track of and in many cases directly access real-time and archived data streams; as such, the data handling system is shared

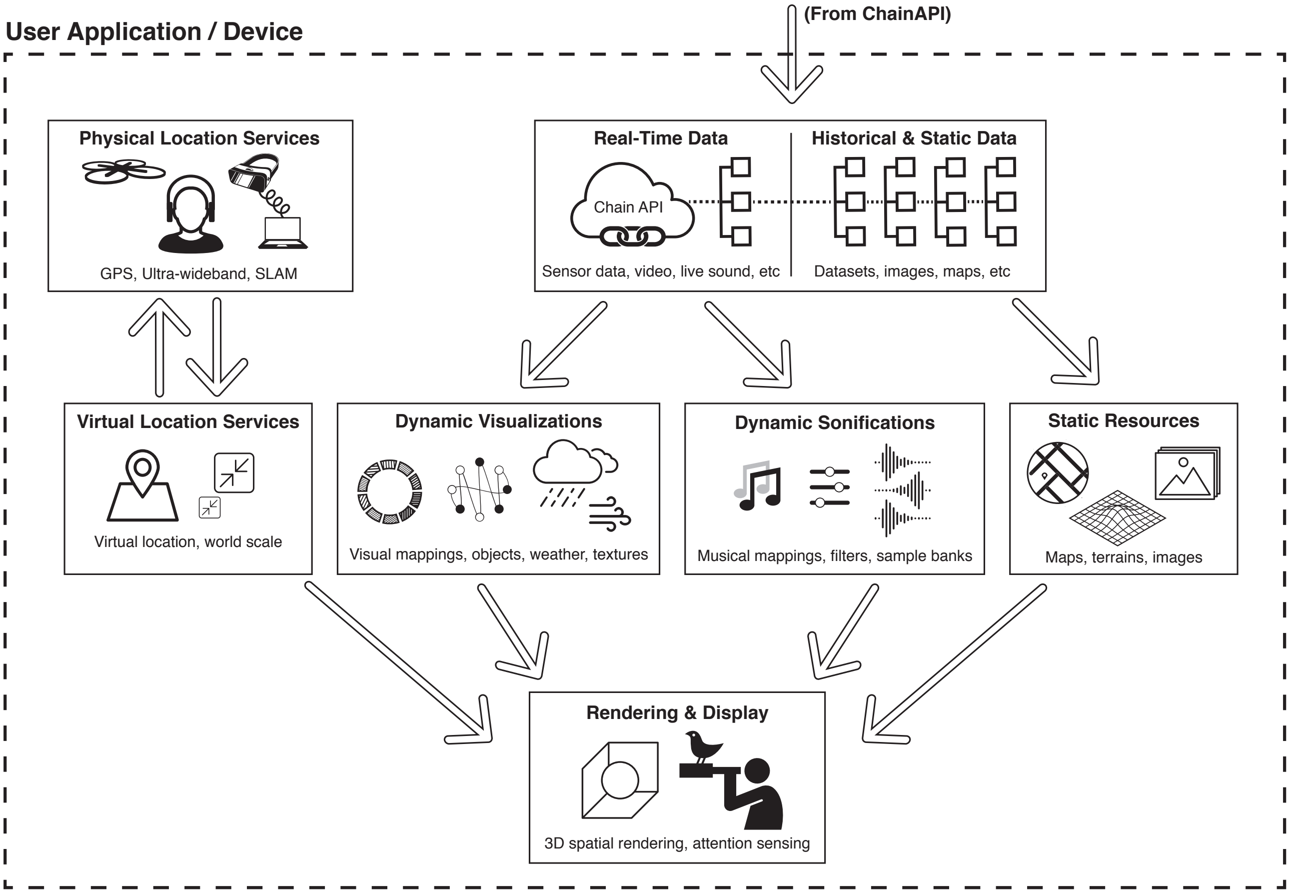

Figure 10. *Components of the UI to Networked Sensory Landscape.*

across them and reflects the design of ChainAPI. This approach lets us decouple the state of the sensor network from any UI code; the application updates dynamically as new nodes come online or others go off.

On startup, our realtime data component, called ChainSync, downloads a summary digest of the sensors on a given top-level ChainAPI “site.” This could be the closest site to the user in an onsite AR application, for example. The summary includes device IDs, geolocations, and a cache of recent data, as well as the unique device URLs. In this way, nodes and sensors can be instantiated in the virtual environment without any hard coding of link URLs, save for the first link to ChainAPI. After parsing the site summary, most of our applications subscribe to ChainAPI’s real-time data stream, using a component we call ChainSocket. Handlers within the application can subscribe to updates for complete sites, sensor nodes, or individual sensors.

As new data comes in, it is forwarded appropriately, keyed on the unique ChainAPI URLs for each device. Some of our systems use ChainAPI to refer to data resources linked on other servers. For example, live audio streams are indexed in ChainAPI but served from other servers.

Some of these data streams require special handling by the application. For example, to handle the special case of multichannel audio streams, we built a cross-platform game engine plugin that decodes audio streams without limits on channel count. Our current systems use the plugin to decode 30-channel Opus-encoded audio streams from the microphone installation at Tidmarsh.

4.3 Location and Navigation

In our mixed reality applications, users are always situated in three-dimensional space and the data are presented around them, responding to their motion and allowing them to explore by moving throughout their environment. In the literature, we find numerous examples of spatially registered data overlay, with applications in geographical information systems (GIS), agriculture, and many other related areas. In one early example, King, Piekarski, and Thomas (2005) visually overlay data such as harvest yield directly onto a grape vineyard in situ. All these systems require the user's real-world motion be tracked so that content can be rendered around them, maintaining alignment between the physical and virtual. These applications also require us to align data associated with physical locations in the real world (generally represented with latitude, longitude, and elevation) with locations in a local virtual world (in x, y, z coordinates).

4.3.1 Real-World Location Tracking. We have explored several technologies for tracking users in physical space. For outdoor use, GPS is an obvious option, and works well for applications in relatively open areas that don't require precision below a few meters. Ultra-wideband radio (UWB) offers much higher precision (to 20 cm) but requires anchor nodes to be installed for localization. In prior work (Russell, Dublon, & Paradiso, 2016), we designed a head-tracking system that uses UWB when it is available and can fall back to GPS otherwise. In both cases, we use an inertial measurement unit (IMU, a combination gyroscope, accelerometer, and magnetometer) to measure the head's orientation, which we combine with the location from UWB or GPS to synthesize the scene around them. The UWB and IMU data are streamed from a custom headset to the user's mobile phone over Bluetooth Low-Energy, where it is combined with the phone's built-in GPS. This application is described in more detail in Section 4.7.5.

We have also used Microsoft's HoloLens platform for our AR applications. The HoloLens uses a depth camera and four standard visible light cameras to perform Simultaneous Location and Mapping (SLAM), meaning that it both maps the environment and tracks its own position.

4.3.2 Virtual Navigation. Each platform and environment provides different navigational affordances to the user. In our desktop software, the user moves with a mouse and keyboard, similar to traditional first-person video games. In VR they can look around freely and move in a limited area, and can travel larger distances by pointing their controller in the direction they want to go and pressing a button to accelerate in that direction. In AR applications, data representations are overlaid onto the physical environment, allowing users full range of motion.

4.3.3 Mapping Geography. Because our data originates in the real world, we have developed a rich toolkit for managing geo-tagged data and mapping it onto virtual environments. Even in our full-scale augmented reality environments, where the virtual and physical worlds are overlaid on top of each other, we need to convert global latitude, longitude, and elevation coordinates into local x, y, and z coordinates in meters. This also sometimes involves scaling up or down into miniature.

Our software library allows designers to place a single reference object into a virtual scene with known geographic coordinates, and the library can use that object as a reference to place the rest of the scene, converting their geographic coordinates to virtual world coordinates. We use the Web Mercator projection so that our virtual world mapping will align with map tiles downloadable from many mapping services (Battersby, Finn, Usery, & Yamamoto, 2014). When designing a scene, it is also possible to place objects directly using their virtual world coordinates, so both coordinate systems can be easily mixed. For example, in Doppelmarsh we get all the sensor coordinates from ChainAPI at runtime and they are placed using our geographic mapping, but other objects such as logs, trees, and terrain are placed in the scene directly using the virtual world coordinates.

4.4 Visualization

Aiming at a more spatial, embodied, visceral interaction with the data, we created a game engine toolkit for modular 3D animated visualizations that can be mixed and matched in different applications. The architecture is oriented around a publish/subscribe model, where data modules (such as ChainSync) can publish to any number of subscribed consumers (such as the visualization modules). This allows us to build self-contained behaviors that can be easily added or removed.

Some visualizations affect the whole application, such as weather or virtual camera effects. Others are local and may vary throughout the site, such as a representation of temperature as measured by each of the sensor nodes. Both use the same publish/subscribe mechanism to subscribe to the data, but the code defining the behavior is instantiated differently. For global visualizations, a single instance of the visualization code can live at the top-level of the scene, and is placed when the scene is designed. For locally varying or per-node visualizations, we don't know ahead of time how many will be needed or where they will be. Game engines provide support for prototype objects that can be defined once and instantiated multiply, in our case at the location of each node. Adding a new representation involves designing a visual or sonic element and exposing one or more parameters to be mapped to a sensor value. For example, properties of a 3D object, such as color or size, could be mapped to humidity. The resulting prototype would then be automatically instantiated across the virtual terrain for each sensor. This framework allows visualizations to incorporate sensor-driven models that viewers can physically move through in virtual reality, such as a kinetic gas model driven by temperature (Remsen, 2017).

4.5 Sonification

Much of our work at Tidmarsh has focused on the sonic experience, which includes both informational auditory display and data-driven musical composition. This work has the parallel goals of improving people's understanding of the environment by extending their perception and building empathetic connections through aesthetic experience. In both cases we seek to augment the natural soundscape rather than replacing it, and prioritize preserving the user's unmediated experience of their natural world.

While most game engines include a wide variety of visual effects and models, they lack the tools we needed to accomplish our vision for pervasive spatial sonification. To fill this gap, we built a framework to allow composers with limited or no knowledge about sensor data processing to compose musical pieces driven by ubiquitous sensing (Lynch & Paradiso, 2016). This framework, called "SensorChimes" and shown in Figure 11, aims to augment our sense of presence by providing additional information about the environment through sound, expanding what we can readily and immediately intuit. The windchime, a prehistoric (and still evolving; Hopkin, 2005) wind sensor that makes music, inspires this project. The windchime is an augmentation to the acoustic environment that mechanically couples wind speed and direction to sound. This project reimagines, generalizes, and augments this concept in the digital domain with electronic sensors that measure many parameters, electronic music composition, and virtual reality.

SensorChimes is implemented as a library for the graphical programming environment Max/MSP which makes it easy to route real-time and historical data from a sensor network into a Max patch. It has also been ported to the PureData environment and made fully embeddable within the game engine, where it can be used to create rich and responsive real-time musical experiences that are coupled with the 3D visualizations. This allows composers and sound designers to work within a familiar musical programming environment. SensorChimes provides an interface for data from each device in the network as well as aggregate metrics over many devices, allowing for quick realization of innumerable musical mapping ideas.

Four compositions have been written using SensorChimes. Each uses a different mapping strategy and explores a different part of the potential of sensor network-driven music as a new canvas for artists. In one example, by coauthor Evan Lynch, real-time data modulates a live mix of looping acoustic tracks,

Selected Composition: Clarinet Loops - Evan Ziporyn and Evan Lynch

Composition Settings: array6 Granular Buffer

Position Source: mouse

'0x8138'

accel_x

0.

Absolute/Deviation:

30. array1

30.

30.

Temp. (°F) 16.620001

Ill. (lux) 1.091553

Temp. Dev. (std) -0.437321

Ill. Dev. (std) -0.351031

Hum. Dev. (std) 0.

Figure 11. *SensorChimes is a composition environment for writing sensor-driven music. Composers can use Max/MSP or PureData to build sonification programs, which can then be embedded within our mixed-reality applications to provide portable music and sound design elements for 3D experiences.*

where the mix parameters are determined by the environmental conditions at Tidmarsh. A second piece, by commissioned composer Evan Ziporyn, explores using real-time data to modulate a shifting timbral space of instrumental textures. A third piece, by commissioned composer Ricky Graham, uses granular synthesis to realize data on multiple timescales. These musical works are presented as immersive virtual explorations of the wetland with spatialized musical mappings keyed to the user's position in the virtual landscape. Several examples of SensorChimes-based compositions are available at http://resenv.media.mit.edu/sensorchimes.

4.6 Devices and Interaction

Using the game engine Unity3D as a basis, we have built modular applications targeting a number of different devices. Virtual environments may be rendered on a flat desktop display, virtual reality displays such as the HTC Vive and Oculus Rift, augmented reality displays such as the Microsoft HoloLens, or our own custom hardware, such as HearThere (see Section 4.7.5). Components can be included depending on the platform needs; for example, our auditory AR systems depend on the real-time data and head-tracking components but omit the visual rendering systems.

4.6.1 User Attention. Our interfaces make large amounts of real-time data perceptually available to users in both remote settings and in situ. To manage this flow, we have explored various ways of sensing user intent to offer more implicit control over these displays. For example, as a user concentrates on a body of water, our auditory AR display offers more information from the subsurface, such as hydrophone audio streams and water quality sonification. Similarly, a visual AR display would render more detail in subsurface animations. Currently, we use focused gaze as a proxy for attention, an approach we have found intuitive to users in AR scenarios.

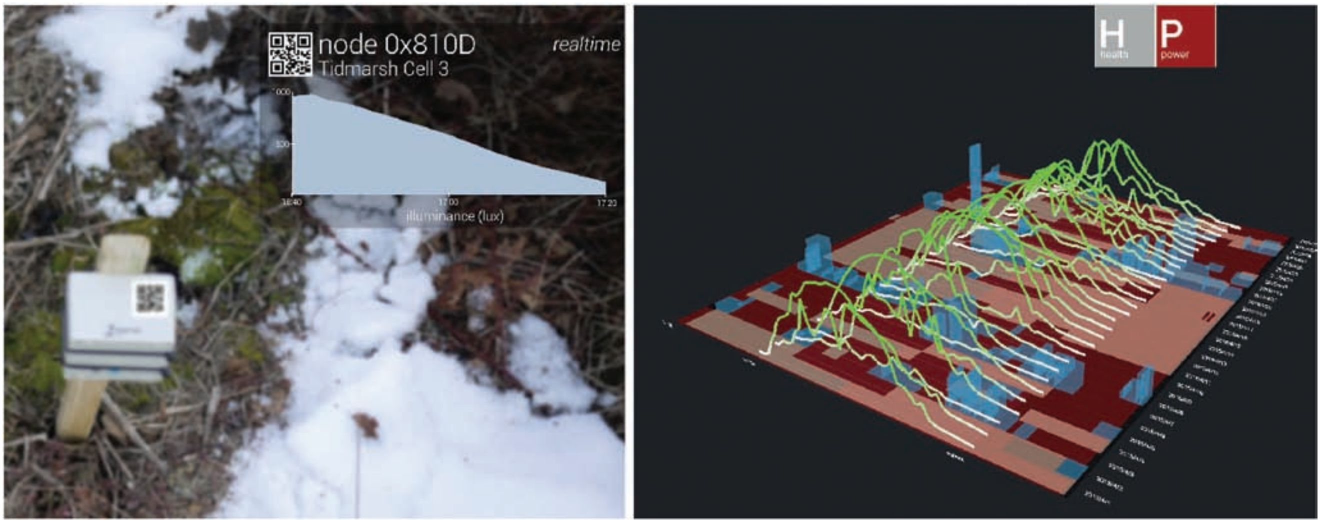

Figure 12. *Left: Sensor Glasses allows sensor data to be accessed by a Google Glass user. Right: A web interface developed by a visiting researcher explores the use of web technologies to create an interactive visualization of the sensor network and data. The application pulls data from ChainAPI and presents it in several different ways.*

In addition, we are conducting psychoacoustic studies investigating methods of sensing the wearer's auditory attention, and adjusting the source levels in response. In addition to gaze, we collect physiological signals such as heart rate variability and facial expression as subjects shift their attention from source to source in a spatial sound presentation. As future work, we intend to train models of attention for individuals. In our vision, the system would modulate the distribution of detail in the display, reacting to the user's mental attention instead of relying on stationary gaze as a proxy.

4.7 Applications

In this section, we describe several user-facing applications that we have built. These applications bring together components from the preceding sections to provide the user with various means of experiencing and interacting with the Networked Sensory Landscape.

4.7.1 Visualization Tools. While most of our work has focused on virtual environments, we have developed a number of other visualization tools for data, network performance, and diagnostics. Figure 12 shows a web-based tool plotting the charging state of our solar-powered backhaul against weather conditions. The same application provides several different views of the sensor network including a visualization of packets received from each sensor node over time, and an overview of when each node was active on the network. Yet another view renders data from each sensor on a stylized map of the site, using color and height to show different values (shown in Figure 9[c]). It also presents an interface that allows the relationships between nearby sensor nodes to be explored, and allows the user to scrub through historical data in addition to the live view (Li, Dublon, Mayton, & Paradiso, 2015).

Using Google's Glass platform, we developed an application called Sensor Glasses (also shown in Figure 12) that uses the built-in camera to recognize QR codes printed on the sensor nodes and display plots of recent sensor data. While the visualization is location-dependent, the displayed information doesn't maintain registration with physical space (the display floats at a fixed position relative to the user's head).

4.7.2 Doppelmarsh. *Doppelmarsh* is a virtual world based on the state of the physical environment at Tidmarsh. The virtual site is visually and sonically rendered and displayed to the user, re-synthesizing reality

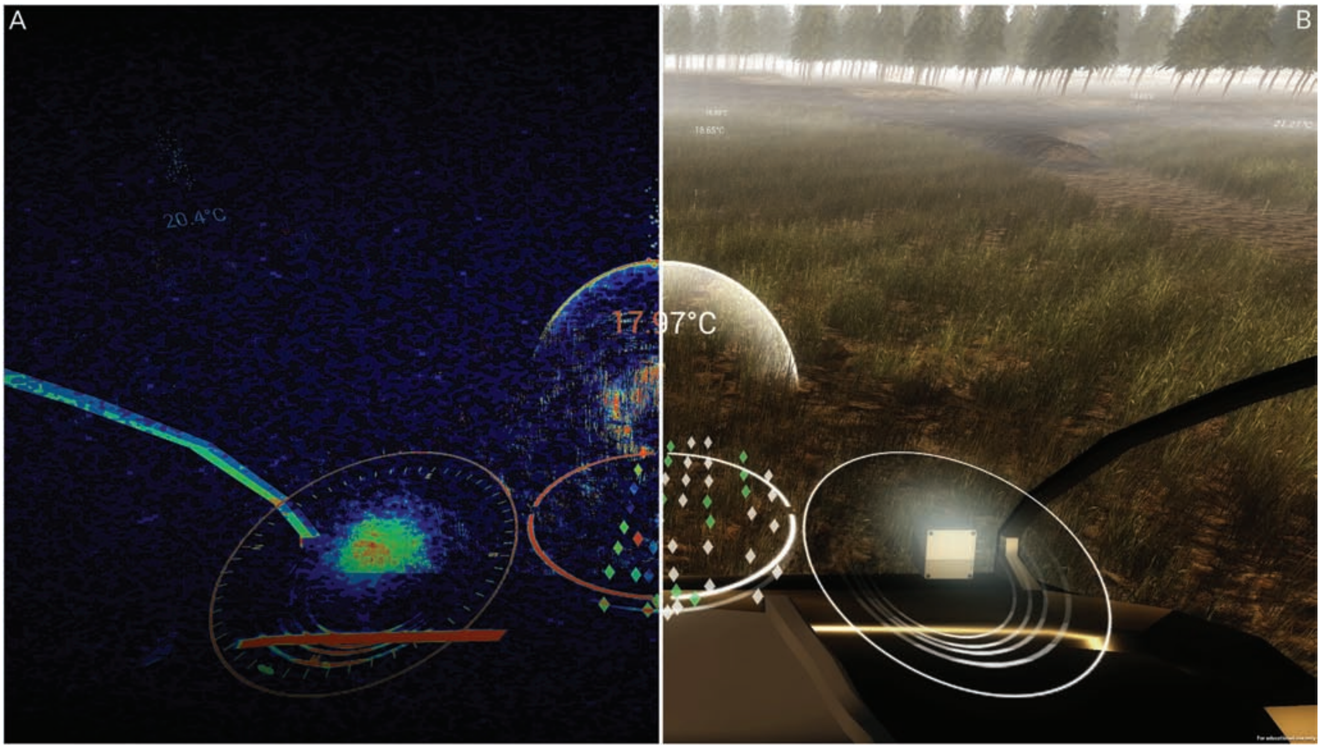

Figure 13. *Different virtual "lenses" highlight various aspects of the sensory world in Doppelmarsh: at left (A) with heatmaps on the terrain, and at right (B) with sensor-driven simulated microclimates. A "mini-map" overlay shows the statuses and relative locations of sensor nodes around the user.*

but free from physical constraints. The development of Doppelmarsh has opened up new avenues for research, such as dynamic automated scene painting from live video and real-time rendering of environmental conditions (Haddad et al., 2017).

The Doppelmarsh terrain forms the virtual world onto which data are rendered. Our terrain models originated as topographic maps and LIDAR scans that we converted and imported into the game engine as heightmaps. As the construction work on the site has transformed the terrain, we have modified the models to reflect the current landscape using the restoration engineering plans as a guide. The terrain contour itself will continue to be updated to reflect changes from the restoration as new data are collected, including from UAV-based imaging and LIDAR flyovers.

We have developed a variety of interaction modalities for user exploration of the virtual environment. The primary one is a first-person view, shown in Figure 13, that allows users to see and hear the site as they would if they were there in person. The first-person view is well-suited to virtual reality displays, where the user's head movements are tracked and used to position the virtual camera. This mode provides a realistic sense of scale on the ground, but traversing large distances in this way is slow. For this reason, a top-down mode allows users to move more quickly, as well as take in larger sections of topography, wildlife, and sensor data. Finally, a position-tracking VR headset with accompanying handheld controllers (HTC Vive) can be used to fly in first-person view. These modes are also used to drive and explore other systems on the site, such as a remotely operated UAV (detailed in Section 4.7.3).

Doppelmarsh relies on our real-time data framework introduced in Section 4.2. As such, the application does not contain any hard coding of nodes or sensors, but rather crawls ChainAPI to build up the state of the

Figure 14. *Doppelmarsh virtual weather patterns: fog is controlled by nearby humidity sensors, and rain intensity by precipitation meters. Sensor data are combined with machine vision on the camera images to control rendering of snowpack and vegetation.*

sensor network. New data is pushed to the application over the ChainAPI WebSocket stream.

Real-time sensor data visualizations are associated with 3D representations of each sensor node automatically placed on the terrain using the metadata received from ChainAPI. In contrast to our previous work (Dublon et al., 2011), where the sensor visualizations were strongly animated by the sensor values and dominated how the building model looked, our current visualizations for Tidmarsh blend more subtly with the environment. Sensor data is seen as white text floating above the sensor location, which can be configured to show particular metrics of interest. Whenever a sensor updates, an illuminated "burst" is animated. As we add nodes and sensors to the landscape, we will be adding many different types of sensor-driven animation, but will still strive to keep a pleasing aesthetic guided by the realistic wetland depiction.

Weather conditions are extracted from the sensor data, which are then synthesized into experiential dimensions of the virtual landscape, shown in Figure 14. For instance, wind data from sensor nodes are used to control the in-game wind speed and direction, which animate the movements of the grass and the trees in the virtual marsh. From humidity sensors, we are able to extract whether it is foggy or rainy, which we then manifest as virtual fog and rain, with intensity determined by an on-site rain gauge. For slower transformations in the environment, such as seasonal changes, we extract features from images captured daily from cameras on-site. For example, the grass tint in Doppelmarsh is taken from a palette recognition algorithm and then rendered dynamically in the scene. This technique is also used to detect and render conditions such as snow accumulation.

In keeping with the relaxing natural ethos suggested by our visualization, we have created a default accompanying sensor-driven soundscape in our virtual marsh experience. Rendered in SensorChimes (see Section 4.5), the soundscape is inspired in part by Wendy Carlos' *Sonic Seasonings* (Carlos, 1972). It is based on a combination of recorded plucked string

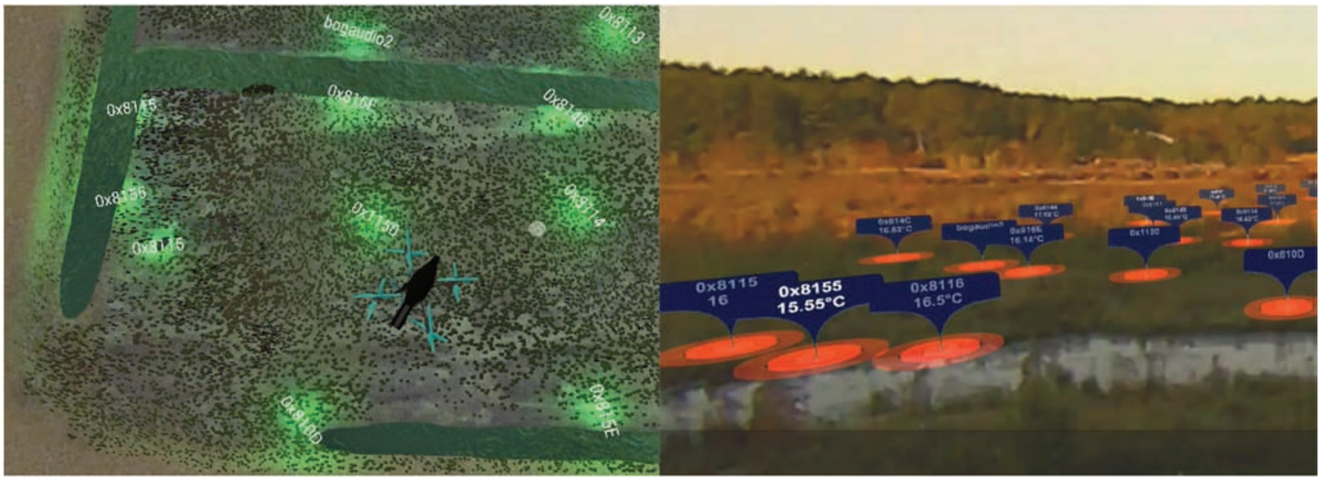

Figure 15. *Left: The QuadraSense control view, showing a top down rendering of the UAV above Tidmarsh. Clicking the map commands the UAV to fly to the indicated position. Right: A live video feed from the UAV, with overlaid sensor data.*

and singing bowl samples mixed with live audio from proximate microphones; both are virtually spatialized to be responsive to the user's movement and respond to sensed parameters. For the data sonification, the sample's pitch is determined by temperature and timbre by humidity; the musical scale changes from day to night, from a pastoral daytime experience to a more mysterious nighttime musical setting. Other soundscape programs described in Section 4.5 can be loaded by the user at runtime.

We are taking Doppelmarsh in a number of directions; already, it has been extended as a control interface for some of the systems on site. To that end, we are also working to integrate the 3D world with previous work developing more traditional 2D graphical interfaces, enabling users to jump back and forth between an experience of presence on the ground and a more symbolic, data-centered view. We are also working to enable more fluid time traversal similar to the interface we provided in DoppelLab, requiring new features in ChainAPI that would pre-compute and cache averages over intervals and other longer-term metrics, such as variance and range.

Most recently, we have integrated our systems with LO collaborator Rob Vincent's pit tagging system, which will allow us to render herring movements in real time. We are also working on ways of representing information from below the surface, such as soil and water temperatures at depth, and related hydrologic models. One approach may be to adjust the terrain opacity and place animations below ground. Finally, we are also investigating more automation in scene painting and terrain updates; as new LIDAR data or even geotagged public photographs become available, we would like to be able to integrate information, such as plant textures in the virtual environment, without human intervention.

4.7.3 QuadraSense.

Camera-equipped drones are now readily available, and using sensor-laden drones with fixed sensor network deployments is a fertile area explored in recent research (Valente et al., 2011; Dong et al., 2014). The Doppelmarsh framework provides a natural way to accommodate mobile sensor agents that live in both real and virtual worlds. QuadraSense (Ramasubramanian, 2015) extends the virtual view in Doppelmarsh with a similarly unconstrained physical camera, using a UAV controlled by the remote user as a semi-autonomous agent in the Tidmarsh airspace. The system provides two interfaces, shown in Figure 15. Navigation is controlled from within the Doppelmarsh application running on a traditional screen: as the user clicks on a location, commands are sent to the UAV and it navigates using GPS to the corresponding

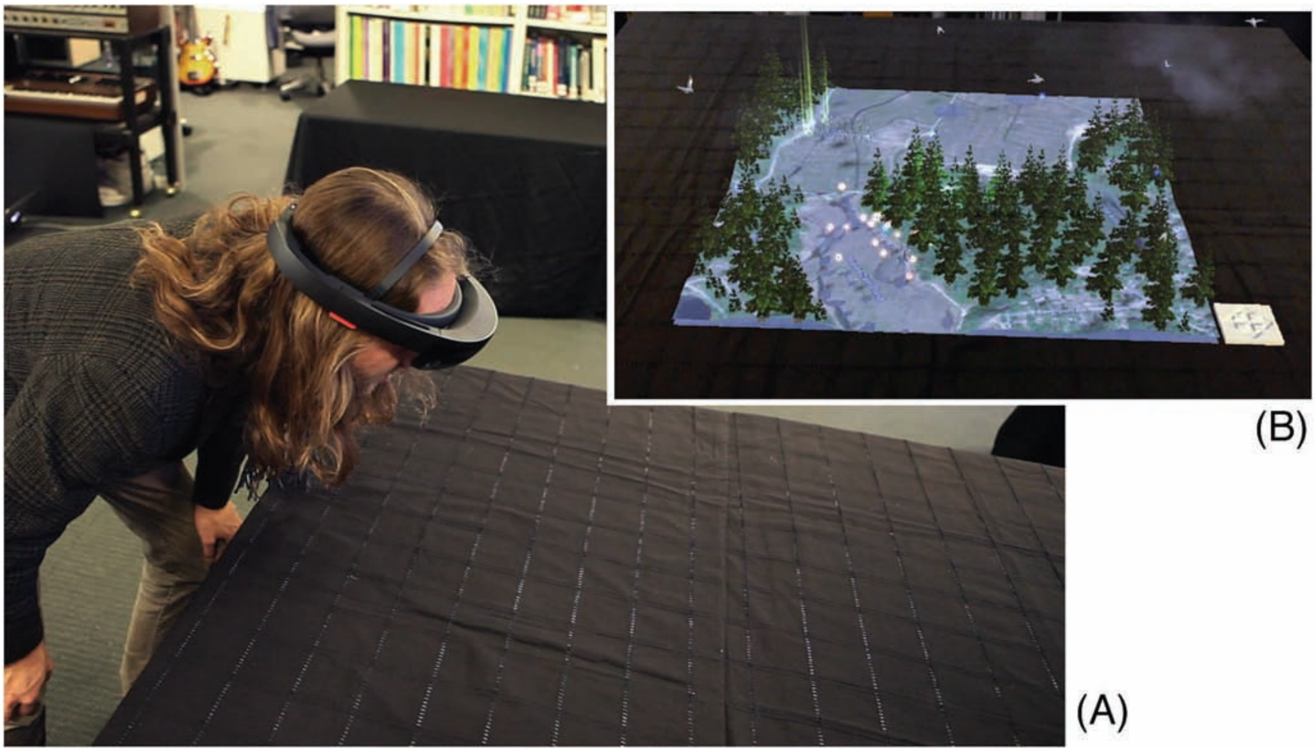

Figure 16. *Hakoniwa uses a head-mounted augmented reality display (A) to render a real-time sensor-driven miniature marsh landscape on a tabletop (B).*

position in the physical airspace. Simultaneously, near-omnidirectional real-time video imagery from the UAV's 180° camera is presented to the user through a head-mounted display (HMD). Users see sensor data rendered on top of the real-time imagery of the landscape, and can switch seamlessly between the aligned virtual and physical cameras. Turning one's head in the VR headset dewarps the omnidirectional video in hardware to produce a natural and responsive field of view for the user.

4.7.4 Hakoniwa.

Hakoniwa, from the Japanese word for a miniature garden, is our most recent work with augmented reality, combining aspects of Doppelmarsh and HearThere. Using Microsoft's HoloLens, a miniature version of the terrain is rendered atop a physical table in the user's environment. Representations of real-time data (sensor nodes, live microphones) as well as static data (audio recordings) are shown as objects on the terrain. As shown in Figure 16, the user is free to walk around the table and lean in to inspect the miniature Tidmarsh. Sensor readings from the area are shown as floating text, and nearby audio sources (both live and prerecorded) are played as spatialized sound, relative to the user's head position.

Bird and insect calls detected by TidZam (see Section 2.4) are rendered on the scene as icons. As the user's attention (as indicated by the cursor in the center of the view) shifts to various parts of the scene, relevant information is displayed, culled from sensors within the user's field-of-view. The user's gaze not only determines the sensor data they see, but also enhances the sound around that location, affecting the spatial roll-off curves for the nearby sources to allow the user to hear those sources more clearly. Multiple users can share a viewing experience, both in colocated scenarios (where the mini-landscape is rendered on the same table for all participants) and in remote scenarios (where two or more users can look at the same activity on the site from afar).

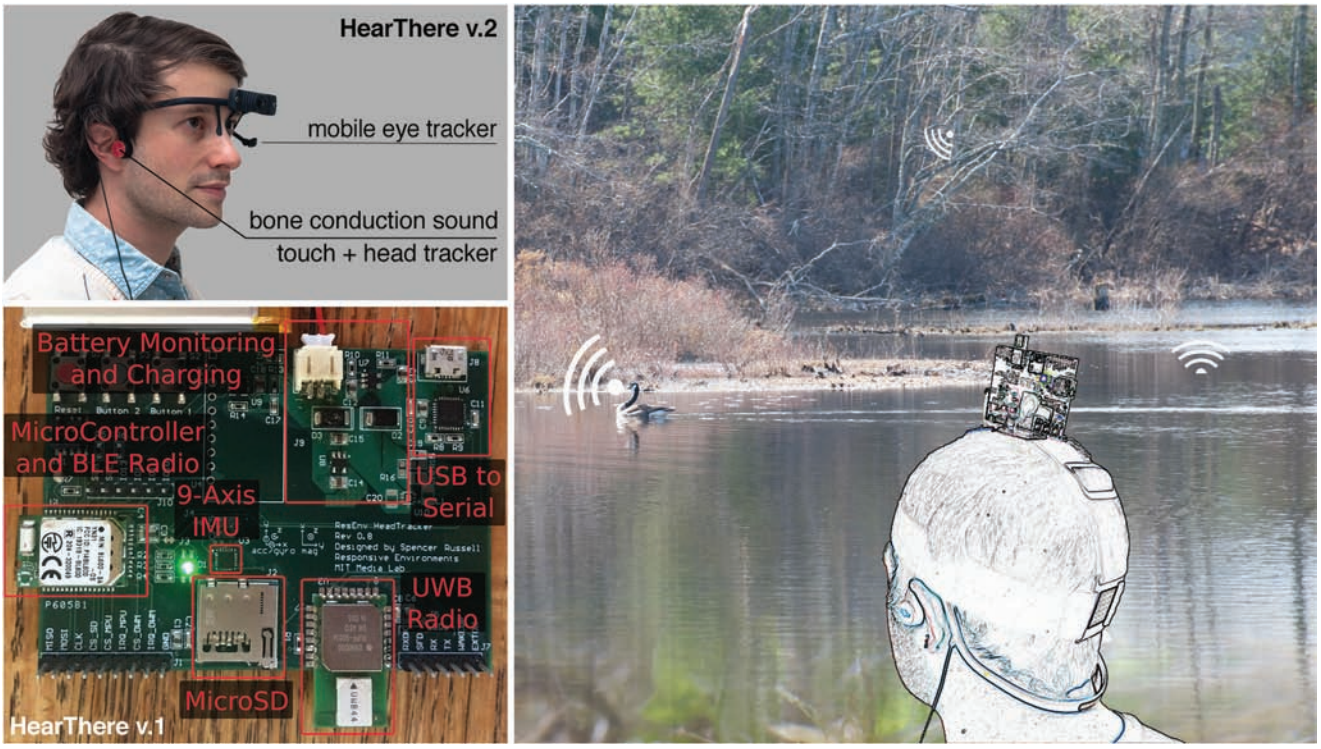

Figure 17. *The HearThere head-mounted wearable device for auditory augmented reality: the wearer gazing across this marsh instrumented with microphones is able to hear spatial live sound as well as spatial sonification of the sensor data. Insets: compact version 2 with eye tracker (top left), precision tracking development version 1 (bottom left).*

4.7.5 HearThere: Bog Ears. HearThere (Russell et al., 2016) is an auditory AR display we developed to allow dynamic auditory objects to be placed in a user's real environment (see Figure 17). HearThere uses head tracking to preserve spatial registration between the real and virtual, and bone-conduction headphones to present the audio without occluding the natural soundscape. The system includes both the head-tracking hardware and software libraries for the mobile iOS platform that allow developers to design auditory AR into a variety of application scenarios. In our prior work, we validated that bone-conduction is a viable technology with which to present spatialized audio. Despite challenges with bandwidth and volume, users still report a convincing sense of spatialization and are able to accurately locate virtual audio sources both indoors and outdoors.

Bog Ears, shown in Figure 18, is an application built for HearThere that extends its wearer's sensory perception on site, during walks through the physical Tidmarsh landscape. Our aim is to provide a transparent experience of extended hearing. This augmented soundscape includes situated sensor data sonifications to give visitors immediate sensory access to otherwise imperceptible properties like soil pH or aquatic oxygen levels. It also allows the user to hear underwater or at great distances by tapping into the audio streams from deployed microphones and hydrophones. These spatialized sounds are directly and immediately attributable by the human perceptual system to their real-world locations, adding to our dynamic perception of the world around us.

Built as an iOS application that pairs with the hardware, Bog Ears tracks the user's location using GPS and head orientation using the HearThere IMU, and presents dynamic spatial sound through bone conduction headphones. All sources are spatialized relative to the user and thus appear to originate in the landscape.

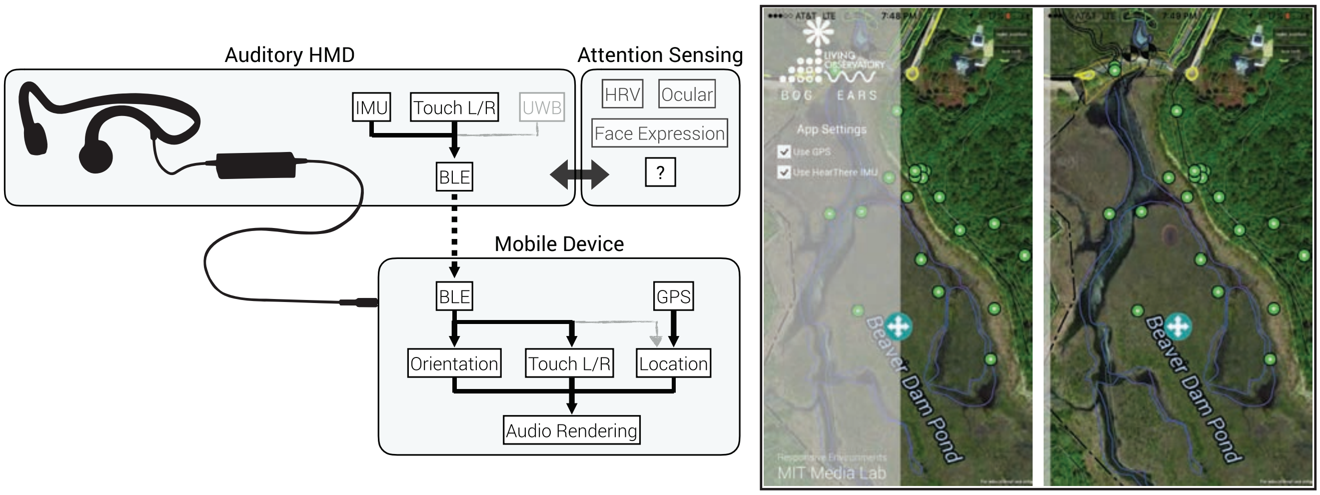

Figure 18. *System diagram and iOS screen captures of Bog Ears, an auditory AR application that provides in situ access to live microphones and sonification. Green circles represent nearby audio sources, and the user's head is shown on the map (center of screen).*

Our current development efforts are focused on making the Bog Ears system responsive to the user's auditory attention, initially through a gaze-based proxy similar to that of Hakoniwa, and later through the attention modeling approach introduced in Section 4.6.1. To that end, we have added mobile eye-tracking glasses to the head-tracking bone conduction system, and are beginning to test the combined system in the field.

4.8 User Experience Takeaways

In our work developing and exploring the Networked Sensory Landscape, we've developed a set of reusable software components and game engine extensions as well as custom hardware, and used those components to build a collection of end-user applications providing real-time visualizations and sonifications on a variety of platforms. These building blocks include modules to manage the users' location and maintain alignment with virtual content, store and retrieve historical data as well as provide real-time updates, and build a richly augmented audiovisual world that can be overlaid on the users' physical environment without mediating their surroundings.

Through these explorations we've solved engineering problems in developing and maintaining a distributed system with many interoperating components, as well as developing new auditory display technologies and localization systems. We've mapped out a design space of sensory augmentation, investigating questions of scale, symbolic and sensory representations of the sensed world, and effective attention-driven interfaces.

As we refine these components, systems, and design methodologies, we envision a set of technologies that can connect the public to their world through networks of sensors and unobtrusive sensory prostheses. These tools also empower artists and composers to craft and curate these experiences to reveal hidden concepts and connections, in the same way that the technology itself can make the imperceptible perceptible, whether it is too small, too quiet, or too far away.

5 A Networked Sensory Landscape Future

At the time of this writing, thousands of decomposing logs and felled trees have been strewn across the nascent wetland landscape of Tidmarsh. Each one is a microhabitat, teeming with life at a spatial and temporal scale that we can barely perceive. Each one presents an opportunity to capture and transmit an important story about ecological function and the process of restoration. In our vision, ubiquitous sensing and computation weave into the natural environment to form the Networked Sensory Landscape that would make that story possible. Inspired by work in acoustic ecology (Dunn & Crutchfield, 2006), we are developing vibration pickups and other sensors to capture the insect life and natural processes within. At the same time, leveraging the interface components introduced here, we are investigating new ways of experiencing these processes. For example, a virtual log could be presented at full scale on a tabletop, allowing users to peer in and listen, or be magnified substantially to allow users to walk through their buzzing interiors. On the site itself, a passer-by wearing HearThere might stop and sit on the log to listen to its rich history condensed into a few minutes. Our design process around this single example reflects the approach we are taking overall: meaningful sensing that both advances the science and offers transformative experiences of presence within ecological processes and across scales.

Through virtual world "browsers," we use sensor information to resynthesize reality. In these physically linked virtual environments, users can float across a landscape that maintains a natural aesthetic while being augmented by animation, spatial sonification, and live sound that together bring invisible phenomena into view. Our experiments in the user experience of ubiquitous sensing point at different ways people can be connected to a dense sensor infrastructure. These include wearable devices that extend the perception of on-site scientists and visitors through attention-driven manifestations of embedded sensor data, a framework for mapping the data to music and sound, and mobile agents such as UAVs. To accommodate the diverse requirements of these different applications, we present our hypermedia framework for flexibly organizing and accessing sensor data. Demonstrating the effectiveness of this approach for building interconnected sensor systems, we introduce our ongoing work in recognizing birds and wildlife from linked audio streams and posting the results as virtual sensors. Each of these applications opens up a plethora of exciting future work.

Broadly, while the world is captivated by the Internet of Things, far less attention has been paid to the ways in which our senses of presence can be altered by seamless connection to the sensors appearing all around us. In "Beyond Being There," Hollan and Stornetta (1992) posit that telepresence would not meet its full potential until users could do *more* remotely than on location. While still in its infancy, the Networked Sensory Landscape points to something even beyond Hollan's imagination, as we see a parallel connected world and hear its music develop from day to night, season to season, and year to year. That world will both transform physical presence and enable new forms of telepresence. As large sensor networks become increasingly commonplace, their data are forming a rich medium upon which artists can build creations that continually evolve and grow, driven by underlying activity that reveals both patterns and exceptions in compelling ways. Today we can "tune in" to a wetland, hearing it through the ears of an artist of our choice; tomorrow we could experience a resynthesized city as a symphonic sensory landscape, heralding a new art form of interpretive presence enabled by ubiquitous sensing.

More information on the work presented in this article, as well as live data, project videos, and downloadable applications, are available at http://tidmarsh.media.mit.edu. More information about the Living Observatory, a learning collaborative formed around the Tidmarsh wetland restoration, is available at http://livingobservatory.org.

Acknowledgments

This article was made possible with the support of many individual and institutional partners. Thanks to Tidmarsh Farms for providing the restoration landscape that contains and inspires our work; to the Living Observatory (LO) for underwriting sensor node hardware and other physical infrastructure; to Alex Hackman, MA Department of Ecological Restoration, for his passionate work and continuing insights; to Brian Spires; to Jessica Norriss and Charis Durrance of LO; and to Paula Aguilera and Jonathan Williams. Thanks to Mary Murphy-Hoye, to Intel and EMC for their in-kind and financial support, and to the sponsors of the MIT Media Lab.

Thanks to our affiliates at MIT: Qiansheng Li, for his work visualizing sensor data (MarshVis); Sara Remsen, for her thesis on environmental models in VR; Halsey Burgund, for

Roundware and Bog People, from which we draw recorded audio content; undergraduate researcher Ben Eck; and Robert Vincent, of MIT Sea Grant, for his work on wildlife tracking systems that interface in real time with our network.

Thanks to our natural sciences partners: Kate Ballantine, Assistant Professor of Environmental Studies at Mount Holyoke and Christine E. Hatch, Assistant Professor of Geosciences at UMass, for their indispensable guidance in sensor development and deployment. Thanks to the US EPA for their critical work on the site and across the region, and to our landscape partners: USDA-NRCS, USFWS, NOAA, MA Environmental Trust (MET), MA Fish and Game, and the Town of Plymouth. Thanks to Mass Audubon and Bob Wilber, its Director of Land Conservation, for establishing a wildlife sanctuary at Tidmarsh, and for their stewardship of the nature of Massachusetts.

References

Battersby, S. E., Finn, M. P., Usery, E. L., & Yamamoto, K. H. (2014). Implications of Web Mercator and its use in online mapping. *Cartographica: The International Journal for Geographic Information and Geovisualization*, *49*(2), 85–101.

Bidgood, J. (2017). The rewilding of a century-old cranberry bog. *The New York Times*, A9. Retrieved from https://nyti.ms/2tHX8hX

Brandes, T. S. (2008). Automated sound recording and analysis techniques for bird surveys and conservation. *Bird Conservation International*, *18*(S1), S163–S173.

Carlos, W. (1972). *Sonic seasonings.* Columbia Records.

Dobson, K. (2016). *The energy propagated.* Retrieved from http://cargocollective.com/TheEnergyPropagated

Dong, M., Ota, K., Lin, M., Tang, Z., Du, S., & Zhu, H. (2014). UAV-assisted data gathering in wireless sensor networks. *The Journal of Supercomputing*, *70*(3), 1142–1155.

Dublon, G., & Paradiso, J. A. (2014). Extra Sensory Perception. *Scientific American*, *311*(1), 36–41.

Dublon, G., Pardue, L. S., Mayton, B., Swartz, N., Joliat, N., Hurst, P., & Paradiso, J. A. (2011). DoppelLab: Tools for exploring and harnessing multimodal sensor network data. *IEEE International Conference on Sensors*, 1612–1615.

Dunn, D., & Crutchfield, J. P. (2006). Insects, trees, and climate: The bioacoustic ecology of deforestation and entomogenic climate change. arXiv preprint q-bio/0612019

Fielding, R. T. (2000). *Architectural styles and the design of network-based software architectures* (Unpublished doctoral dissertation). University of California, Irvine.

Guinard, D., Trifa, V., & Wilde, E. (2010). A resource oriented architecture for the web of things. *Internet of things (IOT)* (pp. 1–8).

Haddad, D. D., Dublon, G., Mayton, B., Russell, S., Xiao, X., Perlin, K., & Paradiso, J. A. (2017). Resynthesizing reality: Driving vivid virtual environments from sensor networks. *SIGGRAPH Talks*, *51*, 1–2.

Helmreich, S. (2016). Melt. Theorizing the contemporary. *Cultural anthropology website.* Retrieved from https://culanth.org/fieldsights/801-melt

Hill, J. L., & Culler, D. E. (2002). Mica: A wireless platform for deeply embedded networks. *IEEE Micro*, *22*(6), 12–24.

Hollan, J., & Stornetta, S. (1992). Beyond being there. *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, 119–125.

Hopkin, B. (2005). *Wind chimes design and construction: A book and audio CD.* Experimental Musical Instruments.

Hu, W., Bulusu, N., Chou, C. T., Jha, S., Taylor, A., & Tran, V. N. (2009). Design and evaluation of a hybrid sensor network for cane toad monitoring. *ACM Transactions on Sensor Networks*, *5*(1), 1–28.

Interfluve. (2015). *Tidmarsh Restoration Project* [Engineering drawings].

Jacobs, R. A., Jordan, M. I., Nowlan, S. J., & Hinton, G. E. (1991). Adaptive mixtures of local experts. *Neural Computation*, *3*(1), 79–87.

King, G. R., Piekarski, W., & Thomas, B. H. (2005). ARVino: Outdoor augmented reality visualisation of viticulture GIS data. *IEEE International Symposium on Mixed and Augmented Reality.* doi:10.1109/ISMAR.2005.14

Levis, P., Madden, S., Polastre, J., Szewczyk, R., Whitehouse, K., Woo, A., … Culler, D. (2005). Tinyos: An operating system for sensor networks. In W. Weber, J. Rabaey, & E. Aarts (Eds.), *Ambient intelligence* (pp. 115–148). Berlin: Springer.

Li, Q., Dublon, G., Mayton, B., & Paradiso, J. A. (2015). *MarshVis: Visualizing real-time and historical ecological data from a wireless sensor network.* Tech. Rep. MIT Media Lab. Cambridge, MA.

Libelium—Connecting sensors to the cloud. (2015). Retrieved from http://www.libelium.com

Lifton, J., Laibowitz, M., Harry, D., Gong, N. W., Mittal, M., & Paradiso, J. A. (2009). *Metaphor and manifestation:*

Cross-reality with ubiquitous sensor/actuator networks. IEEE Pervasive Computing, *8*(3).

Lloret, J., Garcia, M., Bri, D., & Sendra, S. (2009). A wireless sensor network deployment for rural and forest fire detection and verification. *IEEE Sensors*, *9*(11), 8722–8747.

Locke, D. (2010). *MQ Telemetry Transport (MQTT) V3.1 Protocol Specification*. IBM Developerworks Technical Library. Retrieved from http://www.ibm.com/developerworks/webservices/library/ws-mqtt/index.html

Lynch, E. F., & Paradiso, J. A. (2016). SensorChimes: Musical mapping for sensor networks. *Conference on New Interfaces for Musical Expression*, 137–142.

Marquardt, N., Gross, T., Carpendale, S., & Greenberg, S. (2010). Revealing the invisible: Visualizing the location and event flow of distributed physical devices. *International Conference on Tangible, Embedded, and Embodied Interaction*, 41–48.

Massachussetts Audubon Society. (2017). Discover Tidmarsh: Take a walk on the wildflower side. *Mass Audubon Explore Newsletter*. Retrieved from http://massaudubon.org/news-events/publications/explore/spring-2017

Memsic, Inc.—Wireless Sensor Networks. (2015). Retrieved from http://www.memsic.com/wireless-sensor-networks

Michel, S., Salehi, A., Luo, L., Dawes, N., Aberer, K., Barrenetxea, G., … Zhou, Y. (2009). Environmental monitoring 2.0. *International Conference on Data Engineering*, 1507–1510.

Pfisterer, D., Romer, K., Bimschas, D., Kleine, O., Mietz, R., Truong, C., … Richardson, R. (2011). SPITFIRE: Toward a semantic web of things. *IEEE Communications Magazine*, *49*(11), 40–48.

Ramasubramanian, V. (2015). *Quadrasense: Immersive UAV-based cross-reality environmental sensor networks*. Master's thesis, Massachusetts Institute of Technology.

Remsen, S. (2017). *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.

Rundel, P. W., Graham, E. A., Allen, M. F., Fisher, J. C., & Harmon, T. C. (2009). Environmental sensor networks in ecological research. *New Phytologist*, *182*(3), 589–607.

Russell, S., Dublon, G., & Paradiso, J. A. (2016). HearThere: Networked sensory prosthetics through auditory augmented reality. *Proceedings of the ACM Augmented Human International Conference*, *20*, 1–8.

Russell, S., & Paradiso, J. A. (2014). Hypermedia APIs for sensor data: A pragmatic approach to the web of things. *Proceedings of the 11^th^ International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, 30–39.

Shelby, Z., Hartke, K., Bormann, C., & Frank, B. (2014). *Rfc 7252: The Constrained Application Protocol (CoAP)*. Internet Engineering Task Force.

Soldatos, J., Kefalakis, N., Hauswirth, M., Serrano, M., Calbimonte, J. P., Riahi, M., … Herzog, R. (n.d.). Open IoT: Open Source Internet-of-Things in the Cloud, 13–25.

Stowell, D., & Plumbley, M. D. (2014). Automatic large-scale classification of bird sounds is strongly improved by unsupervised feature learning. *PeerJ*, *2*, e488.

Szewczyk, R., Mainwaring, A., Polastre, J., Anderson, J., & Culler, D. (2004). An analysis of a large scale habitat monitoring application. *International Conference on Embedded Networked Sensor Systems*, *2*, 214–226.

Valente, J., Sanz, D., Barrientos, A., Cerro, J. D., Ribeiro, Á., & Rossi, C. (2011). An air-ground wireless sensor network for crop monitoring. *MDPI Sensors*, *11*(6), 6088–6108.

Vorenhout, M., van der Geest, H. G., & Hunting, E. R. (2011). An improved datalogger and novel probes for continuous redox measurements in wetlands. *International Journal of Environmental and Analytical Chemistry*, *91*(7–8), 801–810.

Watras, C. J., Morrow, M., Morrison, K., Scannell, S., Yaziciaglu, S., Read, J. S., … Kratz, T. (2014). Evaluation of wireless sensor networks (WSNs) for remote wetland monitoring: Design and initial results. *Environmental Monitoring and Assessment*, *186*(2), 919–934.

Zedler, J. B., & Kercher, S. (2005). Wetland resources: Status, trends, ecosystem services, and restorability. *Annual Review of Environment and Resources*, *30*, 39–74.