Chapter 18

ANALYSIS OF WIRELESS SENSOR NETWORKS FOR HABITAT MONITORING

Joseph Polastre[1], Robert Szewczyk[1], Alan Mainwaring[2], David Culler[1,2] and John Anderson[3]

[1] *University of California at Berkeley*
Computer Science Department
Berkeley, CA 94720
{polastre,szewczyk,culler}@cs.berkeley.edu

[2] *Intel Research Lab, Berkeley*
2150 Shattuck Ave. Suite 1300
Berkeley, CA 94704
{amm,dculler}@intel-research.net

[3] *College of the Atlantic*
105 Eden Street
Bar Harbor, ME 04609
jga@ecology.coa.edu

Abstract We provide an in-depth study of applying wireless sensor networks (WSNs) to real-world habitat monitoring. A set of system design requirements were developed that cover the hardware design of the nodes, the sensor network software, protective enclosures, and system architecture to meet the requirements of biologists. In the summer of 2002, 43 nodes were deployed on a small island off the coast of Maine streaming useful live data onto the web. Although researchers anticipate some challenges arising in real-world deployments of WSNs, many problems can only be discovered through experience. We present a set of experiences from a four month long deployment on a remote island. We analyze the environmental and node health data to evaluate system performance. The close integration of WSNs with their environment provides environmental data at densities previously impossible. We show that the sensor data is also useful for predicting system operation and network failures. Based on over one million data readings, we analyze the node and network design and develop network reliability profiles and failure models.

Keywords: Wireless Sensor Networks, Habitat Monitoring, Microclimate Monitoring, Network Architecture, Long-Lived Systems

18.1 Introduction

The emergence of wireless sensor networks has enabled new classes of applications that benefit a large number of fields. Wireless sensor networks have been used for fine-grain distributed control [27], inventory and supply-chain management [25], and environmental and habitat monitoring [22].

Habitat and environmental monitoring represent a class of sensor network applications with enormous potential benefits for scientific communities. Instrumenting the environment with numerous networked sensors can enable long-term data collection at scales and resolutions that are difficult, if not impossible, to obtain otherwise. A sensor's intimate connection with its immediate physical environment allows each sensor to provide localized measurements and detailed information that is hard to obtain through traditional instrumentation. The integration of local processing and storage allows sensor nodes to perform complex filtering and triggering functions, as well as to apply application-specific or sensor-specific aggregation, filtering, and compression algorithms. The ability to communicate not only allows sensor data and control information to be communicated across the network of nodes, but nodes to cooperate in performing more complex tasks. Many sensor network services are useful for habitat monitoring: localization [4], tracking [7, 18, 20], data aggregation [13, 19, 21], and, of course, energy-efficient multihop routing [9, 17, 32]. Ultimately the data collected needs to be meaningful to disciplinary scientists, so sensor design [24] and in-the-field calibration systems are crucial [5, 31]. Since such applications need to run unattended, diagnostic and monitoring tools are essential [33].

In order to deploy dense wireless sensor networks capable of recording, storing, and transmitting microhabitat data, a complete system composed of communication protocols, sampling mechanisms, and power management must be developed. We let the application drive the system design agenda. Taking this approach separates actual problems from potential ones, and relevant issues from irrelevant ones from a biological perspective. The application-driven context helps to differentiate problems with simple, concrete solutions from open research areas.

Our goal is to develop an effective sensor network architecture for the domain of monitoring applications, not just one particular instance. Collaboration with scientists in other fields helps to define the broader application space, as well as specific application requirements, allows field testing of experimental systems, and offers objective evaluations of sensor network technologies. The impact of sensor networks for habitat and environmental monitoring will be measured by their ability to enable new applications and produce new, otherwise unattainable, results.

Few studies have been performed using wireless sensor networks in long-term field applications. During the summer of 2002, we deployed an outdoor habitat monitoring application that ran unattended for four months. Outdoor applications present an additional set of challenges not seen in indoor experiments. While we made many simplifying assumptions and engineered out the need for many complex services, we were able to collect a large set of environmental and node diagnostic data. Even though the collected data was not high enough quality to make scientific conclusions, the fidelity of the sensor data yields important observations about sensor network behavior. The data analysis discussed in this paper yields many insights applicable to most wireless sensor deployments. We examine traditional quality of service metrics such as packet loss; however, the sensor data combined with network metrics provide a deeper understanding of failure modes including those caused by the sensor node's close integration with its environment. We anticipate that with system evolution comes higher fidelity sensor readings that will give researchers an even better better understanding of sensor network behavior.

In the following sections, we explain the need for wireless sensor networks for habitat monitoring in Section 18.2. The network architecture for data flow in a habitat monitoring deployment is presented in Section 18.3. We describe the WSN application in Section 18.4 and analyze the network behaviors deduced from sensor data on a network and per-node level in Section 18.5. Section 18.6 contains related work and Section 18.7 concludes.

18.2 Habitat Monitoring

Many research groups have proposed using WSNs for habitat and microclimate monitoring. Although there are many interesting research problems in sensor networks, computer scientists must work closely with biologists to create a system that produces useful data while leveraging sensor network research for robustness and predictable operation. In this section, we examine the biological need for sensor networks and the requirements that sensor networks must meet to collect useful data for life scientists.

18.2.1 The Case For Wireless Sensor Networks

Life scientists are interested in attaining data about an environment with high fidelity. They typically use sensors on probes and instrument as much of the area of interest as possible; however, densely instrumenting any area is expensive and involves a maze of cables. Examples of areas life scientists currently monitor are redwood canopies in forests, vineyard microclimates, climate and occupancy patterns of seabirds, and animal tracking. With these applications in mind, we examine the current modes of sensing and introduce wireless sensor networks as a new method for obtaining environmental and habitat data at scales and resolutions that were previously impractical.

Traditional data loggers for habitat monitoring are typically large in size and expensive. They require that intrusive probes be placed in the area of interest and the corresponding recording and analysis equipment immediately adjacent. Life scientists typically use these data loggers since they are commercially available, supported, and provide a variety of sensors. Probes included with data loggers may create a "shadow effect"–a situation that occurs when an organism alters its behavioral patterns due to an interference in their space or lifestyle [23]. Instead, biologists argue for the miniaturization of devices that may be deployed on the surface, in burrows, or in trees. Since interference is such a large concern, the sensors must be inconspicuous. They should not disrupt the natural processes or behaviors under study [6]. One such data logger is the Hobo Data Logger [24] from Onset Corporation. Due to size, price, and organism disturbance, using these systems for fine-grained habitat monitoring is inappropriate.

Other habitat monitoring studies install one or a few sophisticated weather stations an "insignificant distance" (as much as tens of meters) from the area of interest. A major concern with this method is that biologists cannot gauge whether the weather station actually monitors a different microclimate due to its distance from the organism under study [12]. Using these readings, biologists make generalizations through coarse measurements and sparsely deployed weather stations. A revolution for biologists would be the ability to monitor the environment on the scale of the organism, not on the scale of the biologist [28].

Life scientists are increasingly concerned about the potential impacts of direct human interaction in monitoring plants and animals in field studies. Disturbance effects are of particular concern in a small island situation where it may be physically impossible for researchers to avoid impacting an entire population. Islands often serve as refugia for species that cannot adapt to the presence of terrestrial mammals. In Maine, biologists have shown that even a 15 minute visit to a seabird colony can result in up to 20% mortality among eggs and chicks in a given breeding year [2]. If the disturbance is repeated, the entire colony may abandon their breeding site. On Kent Island, Nova Scotia, researchers found that nesting Leach's Storm Petrels are likely to abandon their burrows if disturbed during their first 2 weeks of incubation. Additionally, the hatching success of petrel eggs was reduced by 56% due to investigator disturbance compared to a control group that was not disturbed for the duration of their breeding cycle [3].

Sensor networks represent a significant advance over traditional, invasive methods of monitoring. Small nodes can be deployed prior to the sensitive period (*e.g.,* breeding season for animals, plant dormancy period, or when the ground is frozen for botanical studies). WSNs may be deployed on small islets where it would be unsafe or unwise to repeatedly attempt field studies. A key difference between wireless sensor networks and traditional probes or data loggers is that WSNs permit real-time data access without repeated visits to sensitive habitats. Probes provide real-time data, but require the researcher to be present on-site, while data in data loggers is not accessible until the logger is collected at a point in the future.

Deploying sensor networks is a substantially more economical method for conducting long-term studies than traditional, personnel-rich methods. It is also more economical than installing many large data loggers. Currently, field studies require substantial maintenance in terms of logistics and infrastructure. Since sensors can be deployed and left, the logistics are reduced to initial placement and occasional servicing. Wireless sensor network may organize themselves, store data that may be later retrieved, and notify operates that the network needs servicing. Sensor networks may greatly increase access to a wider array of study sites that are often limited by concerns about disturbance or lack easy access for researchers.

18.2.2 Great Duck Island

Great Duck Island (GDI), located at (44.09N, 68.15W), is a 237 acre island located 15 km south of Mount Desert Island, Maine. The Nature Conservancy, the State of Maine, and the College of the Atlantic (COA) hold much of the island in joint tenancy. Great Duck contains approximately 5000 pairs of Leach's Storm Petrels, nesting in discrete "patches" within the three major habitat types (spruce forest, meadow, and mixed forest edge) found on the island [1]. COA has ongoing field research programs on several remote islands with well established on-site infrastructure and logistical support. Seabird researchers at COA study the Leach's Storm Petrel on GDI. They are interested in four major questions [26]:

1 What is the usage pattern of nesting burrows over the 24-72 hour cycle when one or both members of a breeding pair may alternate incubation duties with feeding at sea?

2 What environmental changes occur inside the burrow and on the surface during the course of the seven month breeding season (April-October)?

3 What is the variation across petrel breeding sites? Which of these conditions yield an optimal microclimate for breeding, incubation, and hatching?

4 What are the differences in the micro-environments between areas that contain large numbers of nesting petrels and those areas that do not?

Petrels nest in underground burrows that provide relatively constant conditions in terms of temperature and humidity. Burrows are usually within 2–6 cm of the surface and range from 40 cm to over one meter in length with an internal diameter of approximately 6 cm. One sensor node per burrow is sufficient for data sampling but it must be small enough in size such that the sensor and petrel can coexist without interfering with the petrel's activities and does not obstruct the passage. Burrows occur in discrete "patches" around the island that may be hundreds of meters from locations that can support network and power infrastructure. Each patch may contain over 50 burrows; a large number of these burrows should be populated with sensors. Some should be left unpopulated to evaluate if there are disturbance effects caused by wireless sensors. Sensors should cover as many petrel burrows as possible.

Above ground, the environmental conditions vary widely, depending on vegetation type, density, exposure, and location. Humidity readings at a given point in time will vary with vegetation type; a forested area will have higher humidity due to moisture retained by trees and an open meadow will have lower humidity due to direct sunlight and evaporation. Monitoring the environment above each burrow, biologists can examine differences between the above-ground and in-burrow microclimates. Variations in local microclimates may provide clues to nest site selection and overall fitness.

Overall, the petrel cycle lasts approximately 5 months [16]. The deployed system must efficiently manage its power consumption through low duty cycle operation in order to operate for an entire field season. In order to adequately monitor the habitat, it must be monitored on the spatial scale of the organism at frequencies that match environmental changes and organism behavior. By increasing the size of the area monitored and the number of sampling locations, we can obtain data at resolutions and densities not possible using traditional methods. Temporally, sensors should collect data at a rate equal or greater to changing environmental conditions that the organism experiences (on the order of 5-10 times per hour). Traditional data collection systems calculate the average, minimum, and maximum over 24 hour periods as well as time series data. This methodology runs the risk of missing significant but short-term variations in conditions and presents an overly uniform view of micro-environmental conditions. Data analysis must be able to capture duration of events in addition to environmental changes.

It is unlikely that any one parameter or sensor reading could determine why petrels choose a specific nesting site. Predictive models will require multiple measurements of many variables or sensors. These models can then be used to determine which conditions seabirds prefer. To link organism behavior with environmental conditions, sensors may monitor micro-environmental conditions (temperature, humidity, pressure, and light levels) and burrow occupancy (by detecting infrared radiation) [22].

Finally, sensor networks should be run alongside traditional methods to validate and build confidence in the data. The sensors should operate reliably and predictably.

18.3 Network Architecture

In order to deploy a network that satisfies the requirements of Section 18.2, we developed a system architecture for habitat monitoring applications. Here, we describe the architecture, the functionality of individual components, and the interoperability between components.

The system architecture for habitat monitoring applications is a tiered architecture. Samples originate at the lowest level that consists of *sensor nodes*. These nodes perform general purpose computing and networking in addition to application-specific sensing. Sensor nodes will typically be deployed in dense *sensor patches* that are widely separated. Each patch encompasses a particular geographic region of interest. The sensor nodes transmit their data through the sensor patch to the sensor network *gateway*. The gateway is responsible for transmitting sensor data from the *sensor patch* through a local transit network to the remote *base station* that provides WAN connectivity and data logging. The base station connects to database replicas across the Internet. Finally, the data is displayed to scientists through any number of user interfaces. Mobile devices may interact with any of the networks–whether it is used in the field or across the world connected to a database replica. The full architecture is depicted in Figure 18.1.

Sensor nodes are small, battery-powered devices are placed in areas of interest. Each sensor node collects environmental data about its immediate surroundings. The sensor node computational module is a programmable unit that provides computation, storage, and bidirectional communication with other nodes in the system. It interfaces with analog and digital sensors on the sensor module, performs basic signal processing (*e.g.,* simple translations based on calibration data or threshold filters), and dispatches the data according to the application's needs. Compared with traditional data logging systems, it offers two major advantages: it can *communicate* with the rest of the system in real time and can be *retasked* in the field. WSNs may coordinate to deliver data and be reprogrammed with new functionality.

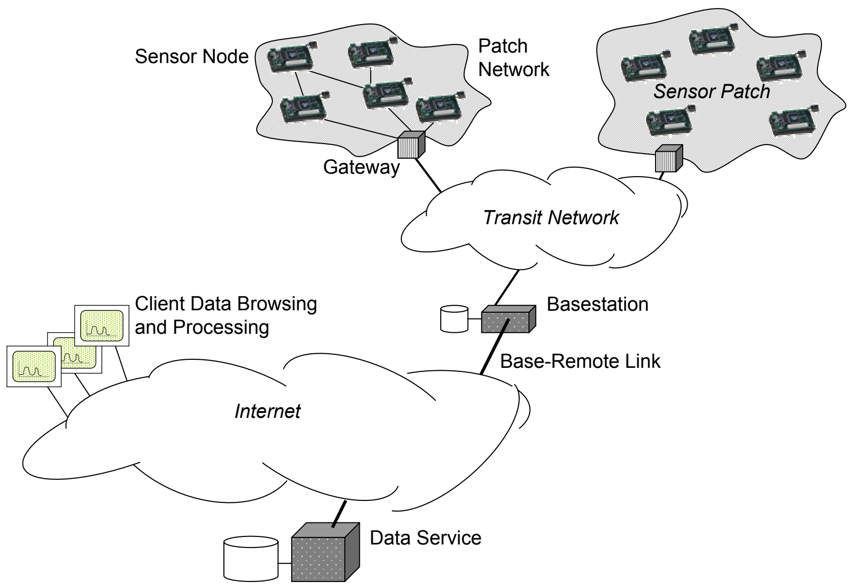

Figure 18.1. System architecture for habitat monitoring

Individual sensor nodes communicate and coordinate with one another in the same geographic region. These nodes make up a *sensor patch*. The sensor patches are typically small in size (tens of meters in diameter); in our application they correspond to petrel nesting patches.

Using a multi-tiered network is particularly advantageous since each habitat involves monitoring several particularly interesting areas, each with its own dedicated sensor patch. Each sensor patch is equipped with a *gateway*. The gateway provides a bridge that connects the sensor network to the base station through a transit network. Since each gateway may include more infrastructure (*e.g.,* solar panels, energy harvesting, and large capacity batteries), it enables deployment of small devices with low capacity batteries. By relying on the gateway, sensor nodes may extend their lifetime through extremely low duty cycles. In addition to providing connectivity to the base station, the gateway may coordinate the activity within the sensor patch or provide additional computation and storage. In our application, a single repeater node served as the transit network gateway. It retransmitted messages to the base station using a high gain Yagi antenna over a 350 meter link. The repeater node ran at a 100% duty cycle powered by a solar cell and rechargeable battery.

Ultimately, data from each sensor needs to be propagated to the Internet. The propagated data may be raw, filtered, or processed. Bringing direct wide area connectivity to each sensor patch is not feasible–the equipment is too costly, it requires too much power and the installation of all required equipment is quite intrusive to the habitat. Instead, the wide area connectivity is brought to a *base station*, where adequate power and housing for the equipment is provided. The base station communicates with the sensor patches using the transit network. To provide data to remote end-users, the *base station* includes wide area network (WAN) connectivity and persistent data storage for the collection of sensor patches. Since many habitats of interest are quite remote, we chose to use a two-way satellite connection to connect to the Internet.

Data reporting in our architecture may occur both spatially and temporally. In order to meet the network lifetime requirements, nodes may operate in a phased manner. Nodes primarily sleep; periodically, they wake, sample, perform necessary calculations, and send readings through the network Data may travel spatially through various routes in the sensor patch, transit network, or wide area network; it is then routed over long distances to the wide area infrastructure.

Users interact with the sensor network data in two ways. Remote users access the replica of the base station database (in the degenerate case they interact with the database directly). This approach allows for easy integration with data analysis and mining tools, while masking the potential wide area disconnections with the base stations. Remote control of the network is also provided through the database interface. Although this control interface is is sufficient for remote users, on-site users may often require a more direct interaction with the network. Small, PDA-sized devices enables such interaction. The data store replicates content across the WAN and its efficiency is integral to live data streams and large analyses.

18.4 Application Implementation

In the summer of 2002, we deployed a 43 node sensor network for habitat monitoring on Great Duck Island. Implementation and deployment of an experimental wireless sensor network platform requires engineering the application software, hardware, and electromechanical design. We anticipated contingencies and feasible remedies for the electromechanical, system, and networking issues in the design of the application that are discussed in this section.

18.4.1 Application Software

Our approach was to simplify the system design wherever possible, to minimize engineering and development efforts, to leverage existing sensor network

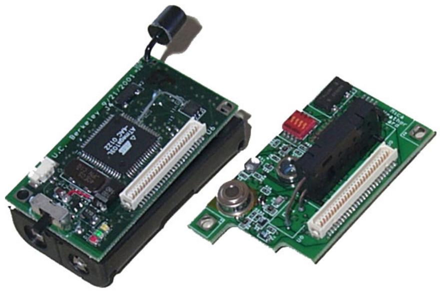

Figure 18.2 Mica mote (left) with Mica Weather Board sensor board for habitat monitoring includes sensors for light, temperature, humidity, pressure, and infrared radiation.

platforms and components, and to use off-the-shelf products where appropriate to focus attention upon the sensor network itself. We chose to use the Mica mote developed by UC Berkeley [14] running the TinyOS operating system [15].

In order to evaluate a long term deployment of a wireless sensor network, we installed each node with a simple periodic application that meets the biologists requirements defined in Section 18.2. Every 70 seconds, each node sampled each of its sensors and transmitted the data in a single 36 byte data packet. Packets were timestamped with 32-bit sequence numbers kept in flash memory. All motes with sensor boards were transmit-only devices that periodically sampled their sensors, transmitted their readings, and entered their lowest-power state for 70 seconds. We relied on the underlying carrier sense MAC layer protocol in TinyOS to prevent against packet collisions.

18.4.2 Sensor board design

To monitor petrel burrows below ground and the microclimate above the burrow, we designed a specialized sensor board called the Mica Weather Board. Environmental conditions are measured with a photoresistive sensor, digital temperature sensor, capacitive humidity sensor, and digital pressure sensor. To monitor burrow occupancy, we chose a passive infrared detector (thermopile) because of its low power requirements. Since it is passive, it does not interfere with the burrow environment. Although a variety of surface mount and probe-based sensors were available, we decided to use surface mount components because they were smaller and operated at lower voltages. Although a probe-based approach has the potential to allow precise co-location of a sensor with its phenomenon, the probes typically operated at higher voltages and currents than equivalent surface mount parts. Probes are more invasive since they puncture the walls of a burrow. We designed and deployed a sensor board with all of the sensors integrated into a single package. A single package permitted miniaturization of the node to fit in the size-constrained petrel burrow.

Even if this initial generation of devices were larger than a highly engineered, application-specific platform would be, we wanted to push in the direction of miniaturized wireless sensors. All sensors resided on a single sensor board, one per mote. This preserved the form factor of the underlying mote platform and limited the circuit board design and simplified manufacturing. The board includes a separate 12-bit ADC to maximize sensor resolution isolate analog noise, and allow concurrent sensor processing and node operation. One consequence of a single integrated design is the amount of shared fate between sensors; a failure of one sensor is likely affects all other sensors. The design did not consider fault isolation among independent sensors or controlling the effects of malfunctioning sensors on shared hardware resources.

18.4.3 Packaging strategy

The environmental conditions on offshore islands are diverse. In-situ instrumentation experiences rain, often with pH readings of less than 3, dew, dense fog, and flooding. They could experience direct sunlight and extremes of warm and cold temperatures. Waterproofing was a primary packaging concern.

Sealing electronics from the environment could be done with conformal coating, packaging, or combinations of the two. Since our sensors were surface mounted and needed to be exposed to the environment, we sealed the entire mote with a parylene sealant. Through successful tests in our lab, we concluded the mote's electronics could be protected from the environment with sealant. A case provides mechanical protection but would not be required for waterproofing. Our survey of off-the-shelf enclosures found many that were slightly too small for the mote or too large for tunnels. Custom enclosures were too costly. Above ground motes were placed in ventilated acrylic enclosures. In burrows, motes were deployed without enclosures.

Of primary concern for the packaging was the effect it has on RF propagation. We decided to use board-mounted miniature whip antennas. There were significant questions about RF propagation from motes inside burrows, above ground on the surface, within inches of granite rocks, tree roots and low, dense vegetation. When we deployed the motes we noted the ground substrate, distance into the burrow, and geographic location of each mote to assist in the analysis of the RF propagation for each mote.

18.4.4 Experiment goals

Since our deployment was the first long term use of the mote platform, we were interested in how the system would perform. Specifically, this deployment served to prove the feasibility of using a miniature low-power wireless sensor network for long term deployments where robustness and predictable operation are essential. We set out to evaluate the efficacy of the sealant, the

Figure 18.3. Acrylic enclosures used at different outdoor applications.

radio performance in and out of burrows, the usefulness of the data for biologists including the occupancy detector, and the system and network longevity. Since each hardware and software component was relatively simple, our goal was to draw significant conclusions about the behavior of wireless sensor networks from the resulting data.

After 123 days of the experiment, we logged over 1.1 million readings. During this period, we noticed abnormal operation among the node population. Some nodes produced sensor readings out of their operating range, others had erratic packet delivery, and some failed. We sought to understand why these events had occurred. By evaluating these abnormalities, future applications may be designed to isolate problems and provide notifications or perform self-healing. The next section analyzes node operation and identifies the causes of abnormal behavior.

18.5 System Analysis

Before the disciplinary scientists perform the analysis of the sensor data, we need convincing evidence that the sensor network is functioning correctly. We look at the performance of the network as a whole in Section 18.5.1 as well as the failures experienced by individual nodes in Section 18.5.2.

In order to look more closely at the network and node operation, we would like to introduce you to the node community that operated on Great Duck Island. Shown in Table 18.1, each of the nodes is presented with their node ID and lifetime in days. Some of the nodes had their batteries replaced and ran for a second "life". Of importance is that some of the nodes fell victim to raw humidity readings of zero or significant clock skew. The number of days after the first sign of either abnormality is referred to as the amount of time on "death row". We discuss individual nodes highlighted in Table 18.1 throughout our analysis in this section and explain their behavior.

Table 18.1. The node population and their properties. RH indicates whether the node experienced raw relative humidity readings of zero during its lifetime (if RH=1). CS indicates that the node experienced excessive clock skew (if CS=1). After the first sign of abnormal humidity readings or clock skew, the node's remaining lifetime (in days) is given in the "death row" (DR) column. The lifetime (in days) of each node on their first and second set of batteries is listed. The total RH and CS counts how many nodes exhibited those properties. The total DR and lifetime is the average time on "death row" and average lifetime over the entire population. Shaded nodes appear in our analysis and graphs.

Node	*RH*	*CS*	*DR*	*Life 1*	*Life 2*
2	1	1	4	14	-
3	1	1	12	14	-
4	1	1	2	2	-
5	1	0	1	13	-
9	1	1	12	12	-
10	1	0	1	1	-
12	0	0	0	25	-
13	0	0	0	31	40
15	0	0	0	31	40
16	1	0	1	1	-
17	0	1	1	27	-
18	0	1	6	44	-
19	1	1	6	2	-
22	1	0	1	1	-
24	1	0	1	14	35
25	1	1	1	1	-
26	0	0	6	6	-
29	0	1	0	56	66
30	0	1	0	51	28
32	1	1	1	44	-
35	0	0	0	54	33
38	0	0	0	35	-
39	0	0	41	44	-
40	1	0	6	6	-
41	1	1	60	67	-
42	0	1	1	6	-
43	1	1	11	12	-
44	1	1	1	1	-
45	0	0	11	13	-
46	1	1	7	67	-
47	0	0	0	16	-
48	1	1	12	16	-
49	1	1	1	1	-
50	1	1	8	8	-
51	1	1	2	2	-
52	1	1	5	6	-
53	0	0	2	8	-
54	1	1	2	4	-
55	0	0	1	54	-
57	0	1	0	67	-
58	1	1	6	6	-
59	1	1	2	2	-
90	1	1	1	1	-
Total	26	26	5.5	20.7	-

18.5.1 Network Analysis

We need evaluate the behavior of the sensor network to establish convincing evidence that the system is operating correctly. Our application was implemented as a single hop network, however the behavior in a single hop is equivalent to what occurs in any WSN radio cell. We begin by examining WSN operation and its performance over time in order to evaluate network cell characteristics.

Packet loss. A primary metric of network performance is packet loss in the network over time. Packet loss is a quality-of-service metric that indicates the effective end-to-end application throughput [8]. The average daily packet delivery is shown in Figure 18.4. Two features of the packet delivery plot de-

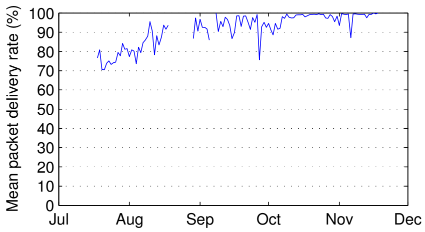

Figure 18.4. Average daily packet delivery in the network throughout the deployment. The gap in the second part of August corresponds to a database crash.

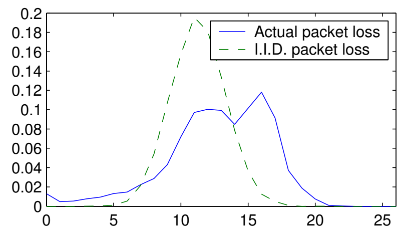

Figure 18.5. Distribution of packet losses in a time slot. Statistically, the losses are not independently distributed.

mand explanation: (1) why was the initial loss rate high and (2) why does the network improve with time? Note that the size of the sensor network is declining over time due to node failures. Either motes with poor packet reception die quicker or the radio channel experiences less contention and packet collisions as the number of nodes decreases. To identify the cause, we examine whether a packet loss at a particular node is dependent on losses from other nodes.

The periodic nature of the application allows us to assign virtual time slots to each data packet corresponding with a particular sequence number from each node. After splitting the data into time slices, we can analyze patterns of loss within each time slot. Figure 18.6 shows packet loss patterns within the network during the first week of August 2002. A black line in a slot indicates that a packet expected to arrive was lost, a white line means a packet was successfully received. If all packet loss was distributed independently, the graph would contain a random placement of black and white bars appearing as a gray square. We note that 23 nodes do not start to transmit until the morning of August 6; that reflects the additional mote deployment that day. Visual inspection reveals patterns of loss: several black horizontal lines emerge, spanning almost all nodes, *e.g.*, midday on August 6, 7, and 8. Looking closer at the packet loss on August 7, we note it is the only time in the sample window when motes 45 and 49 transmit packets successfully; however, heavy packet loss occurs at most other nodes. Sequence numbers received from these sensors reveal they transmitted data during every sample period since they were deployed even though those packets were not received.

More systematically, Figure 18.5 compares the empirical distribution of packet loss in a slot to an independent distribution. The hypothesis that the two distributions are the same is rejected by both parametric (χ^2 test yields 10^8) and non-parametric techniques (rank test rejects it with 99% confidence). The empirical distribution appears a superposition of two Gaussian functions: this is not particularly surprising, since we record packet loss at the end of the

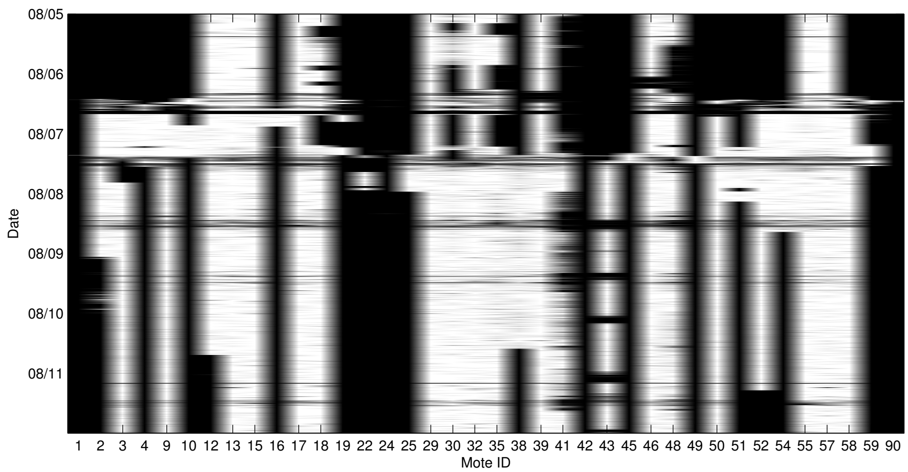

Figure 18.6. Packet loss patterns within the deployed network during a week in August. Y-axis represents time divided into virtual packet slots (note: time increases downwards). A black line in the slot indicates that a packet expected to arrive in this time slot was missed, a white line means that a packet was successfully received.

path (recall network architecture, Section 18.4). This loss is a combination of potential losses along two hops in the network. Additionally, packets share the channel that varies with the environmental conditions, and sensor nodes are likely to have similar battery levels. Finally, there is a possibility of packet collisions at the relay nodes.

Network dynamics. Given that the expected network utilization is very low (less than 5%) we would not expect collisions to play a significant role. Conversely, the behavior of motes 45 and 49 implies otherwise: their packets are only received when most packets from other nodes are lost. Such behavior is possible in a periodic application: in the absence of any backoff, the nodes will collide repeatedly. In our application, the backoff was provided by the CSMA MAC layer. If the MAC worked as expected, each node would backoff until it found a clear slot; at that point, we would expect the channel to be clear. Clock skew and channel variations might force a slot reallocation, but such behavior should be infrequent.

Looking at the timestamps of the received packets, we can compute the phase of each node, relative to the 70 second sampling period. Figure 18.7 plots the phase of selected nodes from Figure 18.6. The slope of the phase corresponds to a drift as a percentage of the 70-second cycle. In the absence of clock drift and MAC delays, each mote would always occupy the same time

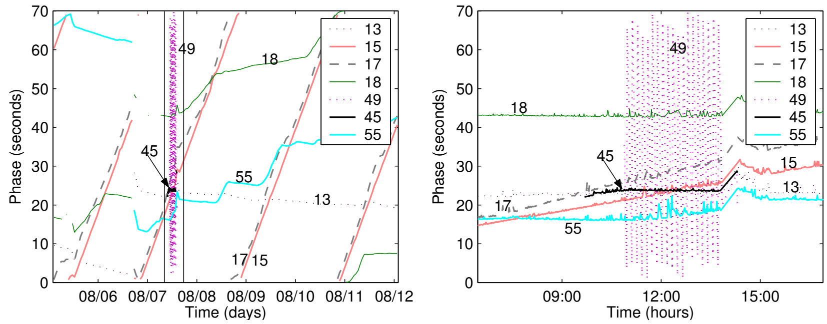

Figure 18.7. Packet phase as a function of time; the right figure shows the detail of the region between the lines in the left figure.

slot cycle and would appear as a horizontal line in the graph. A 5 ppm oscillator drift would result in gently sloped lines, advancing or retreating by 1 second every 2.3 days. In this representation, the potential for collisions exists only at the intersections of the lines.

Several nodes display the expected characteristics: motes 13, 18, and 55 hold their phase fairly constant for different periods, ranging from a few hours to a few days. Other nodes, *e.g.*, 15 and 17 appear to delay the phase, losing 70 seconds every 2 days. The delay can come only from the MAC layer; on average they lose 28 msec, which corresponds to a single packet MAC backoff. We hypothesize that this is a result of the RF automatic gain control circuits: in the RF silence of the island, the node may adjust the gain such that it detects radio noise and interprets it as a packet. Correcting this problem may be done by incorporating a signal strength meter into the MAC that uses a combination of digital radio output and analog signal strength. This additional backoff seems to capture otherwise stable nodes: *e.g.*, mote 55 on August 9 transmits in a fixed phase until it comes close to the phase of 15 and 17. At that point, mote 55 starts backing off before every transmission. This may be caused by implicit synchronization between nodes caused by the transit network.

We note that potential for collisions does exist: the phases of different nodes do cross on several occasions. When the phases collide, the nodes back off as expected, *e.g.*, mote 55 on August 9 backs off to allow 17 to transmit. Next we turn to motes 45 and 49 from Figure 18.6. Mote 45 can collide with motes 13 and 15; collisions with other nodes, on the other hand, seem impossible. In contrast, mote 49, does not display any potential for collisions; instead it shows a very rapid phase change. Such behavior can be explained either though a clock drift, or through the misinterpretation of the carrier sense (*e.g.*, a mote determines it needs to wait a few seconds to acquire a channel). We associate such behavior with faulty nodes, and return to it in Section 18.5.2.

18.5.2 Node Analysis

Nodes in outdoor WSNs are exposed to closely monitor and sense their environment. Their performance and reliability depend on a number of environmental factors. Fortunately, the nodes have a local knowledge of these factors, and they may exploit that knowledge to adjust their operation. Appropriate notifications from the system would allow the end user to pro-actively fix the WSN. Ideally, the network could request proactive maintenance, or self-heal. We examine the link between sensor and node performance. Although the particular analysis is specific to this deployment, we believe that other systems will be benefit from similar analyses: identifying outliers or loss of expected sensing patterns, across time, space or sensing modality. Additionally, since battery state is an important part of a node's self-monitoring capability [33], we also examine battery voltage readings to analyze the performance of our power management implementation.

Sensor analysis. The suite of sensors on each node provided analog light, humidity, digital temperature, pressure, and passive infrared readings. The sensor board used a separate 12-bit ADC to maximize the resolution and minimize analog noise. We examine the readings from each sensor.

Light readings. The light sensor used for this application was a photoresistor that we had significant experience with in the past. It served as a confidence building tool and ADC test. In an outdoor setting during the day, the light value saturated at the maximum ADC value, and at night the values were zero. Knowing the saturation characteristics, not much work was invested in characterizing its response to known intensities of light. The simplicity of this sensor combined with an *a priori* knowledge of the expected response provided a valuable baseline for establishing the proper functioning of the sensor board. As expected, the sensors deployed above ground showed periodic patterns of day and night and burrows showed near to total darkness. Figure 18.8 shows light and temperature readings and average light and temperature readings during the experiment.

The light sensor operated most reliably of the sensors. The only behavior identifiable as failure was disappearance of diurnal patterns replaced by high value readings. Such behavior is observed in 7 nodes out of 43, and in 6 cases it is accompanied by anomalous readings from other sensors, such as a 0^oC temperature or analog humidity values of zero.

Temperature readings. A Maxim 6633 digital temperature sensor provided the temperature measurements While the sensor's resolution is 0.0625^oC, in our deployment it only provided a 2^oC resolution: the hardware always supplied readings with the low-order bits zeroed out. The enclosure was IR

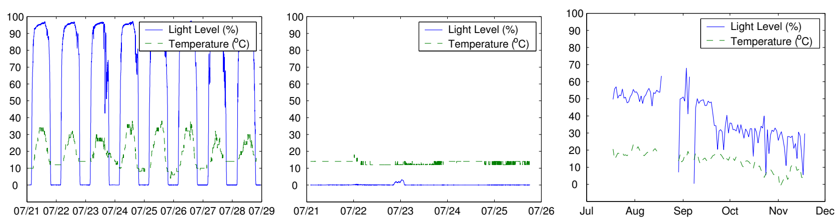

Figure 18.8. Light and temperature time series from the network. From left: outside, inside, and daily average outside burrows.

transparent to assist the thermopile sensor; consequently, the IR radiation from direct sunlight would enter the enclosure and heat up the mote. As a result, temperatures measured inside the enclosures were significantly higher than the ambient temperatures measured by traditional weather stations. On cloudy days the temperature readings corresponded closely with the data from nearby weather buoys operated by NOAA.

Even though motes were coated with parylene, sensor elements were left exposed to the environment to preserve their sensing ability. In the case of the temperature sensor, a layer of parylene was permissible. Nevertheless the sensor failed when it came in direct contact with water. The failure manifested itself in a persistent reading of 0^oC. Of 43 nodes, 22 recorded a faulty temperature reading and 14 of those recorded their first bad reading during storms on August 6. The failure of temperature sensor is highly correlated with the failure of the humidity sensor: of 22 failure events, in two cases the humidity sensor failed first and in two cases the temperature sensor failed first. In remaining 18 cases, the two sensors failed simultaneously. In all but two cases, the sensor did not recover.

Humidity readings. The relative humidity sensor was a capacitive sensor: its capacitance was proportional to the humidity. In the packaging process, the sensor needed to be exposed; it was masked out during the parylene sealing process, and we relied on the enclosure to provide adequate air circulation while keeping the sensor dry. Our measurements have shown up to 15% error in the interchangeability of this sensor across sensor boards. Tests in a controlled environment have shown the sensor produces readings with 5% variation due to analog noise. Prior to deployment, we did not perform individual calibration; instead we applied the reference conversion function to convert the readings into SI units.

In the field, the protection afforded by our enclosure proved to be inadequate. When wet, the sensor would create a low-resistance path between the power supply terminals. Such behavior would manifest itself in either abnor-

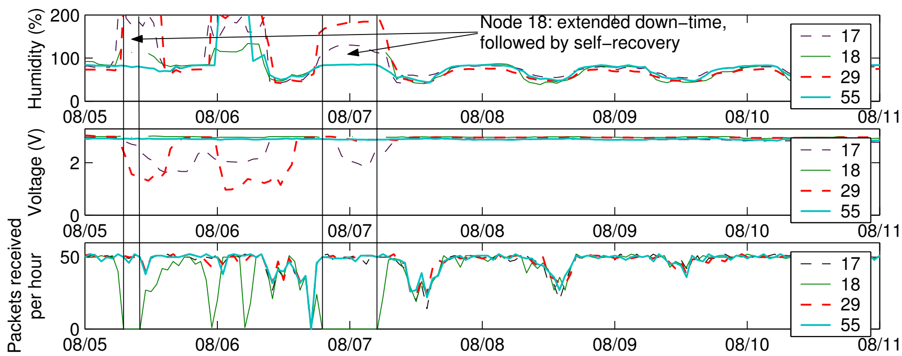

Figure 18.9. Sensor behavior during the rain. Nodes 17 and 29 experience substantial drop in voltage, while node 55 crashes. When the humidity sensor recovers, the nodes recover.

mally large (more than 150%) or very small humidity readings (raw readings of 0V). Figure 18.9 shows the humidity and voltage readings as well as the packet reception rates of selected nodes during both rainy and dry days in early August. Nodes 17 and 29 experienced a large drop in voltage while recording an abnormally high humidity readings on Aug 5 and 6. We attribute the voltage drop to excessive load on the batteries caused by the wet sensor. Node 18 shows an more severe effect of rain: on Aug 5, it crashes just as the other sensors register a rise in the humidity readings. Node 18, on the other hand, seems to be well protected: it registers high humidity readings on Aug 6, and its voltage and packet delivery rates are not correlated with the humidity readings. Nodes that experienced the high humidity readings typically recover when they dried up; nodes with the unusually low readings would fail quickly. While we do not have a definite explanation for such behavior, we evaluate that characteristics of the sensor board as a failure indicator below.

Thermopile readings. The data from the thermopile sensor proved difficult to analyze. The sensor measures two quantities: the ambient temperature and the infrared radiation incident on the element. The sum of thermopile and thermistor readings yields the object surface temperature, *e.g.*, a bird. We would expect that the temperature readings from the thermistor and from the infrared temperature sensor would closely track each other most of the time. By analyzing spikes in the IR readings, we should be able to deduce the bird activity.

The readings from the thermistor do, in fact, track closely with the temperature readings. Figure 18.10 compares the analog thermistor with the digital maxim temperature sensor. The readings are closely correlated although dif-

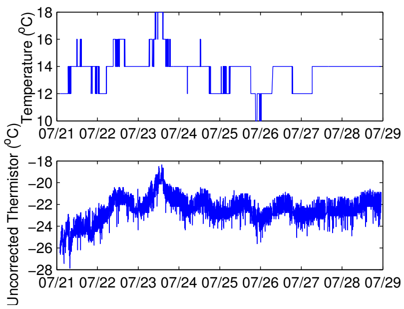

Figure 18.10. The digital temperature sensor (top) and analog thermistor (bottom), though very different on the absolute scale, are closely correlated: a linear fit yields a mean error of less than 0.8^{o}C.

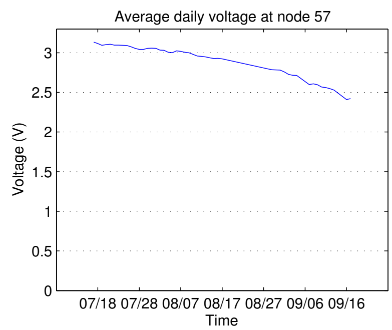

Figure 18.11. Voltage readings from node 57. Node 57 operates until the voltage falls below 2.3V; at this point the alkaline cells can not supply enough current to the boost converter.

ferent on an absolute scale. A best linear fit of the temperature data to the thermistor readings on a per sensor per day basis yields a mean error of less than 0.9^{o}C, within the half step resolution of the digital sensor. The best fit coefficient varies substantially across the nodes.

Assigning biological significance to the infrared data is a difficult task. The absolute readings often do not fall in the expected range. The data exhibits a lack of any periodic daily patterns (assuming that burrow occupancy would exhibit them), and the sensor output appears to settle quickly in one of the two extreme readings. In the absence of any ground truth information, *e.g.*, infrared camera images corresponding to the changes in the IR reading, the data is inconclusive.

Power Management. As mentioned in Section 18.4, one of the main challenges was sensor node power management. We evaluate the power management in the context of the few nodes that did not exhibit other failures. Motes do not have a direct way of measuring the energy they consumed, instead we use battery voltage as an indirect measure. The analysis of the aggregate population is somewhat complicated by in-the-field battery replacements, failed voltage indicators, failed sensors and gaps in the data caused by the database crashes. Only 5 nodes out of 43 have clearly exhausted their original battery supply. This limited sample makes it difficult to perform a thorough statistical analysis. Instead we examine the battery voltage of a single node without other failures. Figure 18.11 shows the battery voltage of a node as a function of time. The batteries are unable to supply enough current to power the node once the voltage drops below 2.30V. The boost converter on the Mica mote is able to extract only 15% more energy from the battery once the voltage drops below 2.5V (the lowest operating voltage for the platform without the voltage regulation). This fell far short of our expectations of being able to drain the batteries down to 1.6V, which represents an extra 40% of energy stored in a cell [10]. The periodic, constant power load presented to the batteries is ill suited to extract the maximum capacity. For this class of devices, a better solution would use batteries with stable voltage, *e.g.*, some of the lithium-based chemistries. We advocate future platforms eliminate the use of a boost converter.

Node failure indicators. In the course of data analysis we have identified a number of anomalous behaviors: erroneous sensor readings and application phase skew. The humidity sensor seemed to be a good indicator of node health. It exhibited 2 kinds of erroneous behaviors: very high and very low readings. The high humidity spikes, even though they drained the mote's batteries, correlated with recoverable mote crashes. The humidity readings corresponding to a raw voltage of 0V correlated with permanent mote outage: 55% of the nodes with excessively low humidity readings failed within two days. In the course of packet phase analysis we noted some motes with slower than usual clocks. This behavior also correlates well with the node failure: 52% of nodes with such behavior fail within two days.

These behaviors have a very low false positive detection rate: only a single node exhibiting the low humidity and two nodes exhibiting clock skew (out of 43) exhausted their battery supply instead of failing prematurely. Figure 18.12 compares the longevity of motes that have exhibited either the clock skew or a faulty humidity sensor against the survival curve of mote population as a whole. We note that 50% of motes with these behaviors become inoperable within 4 days.

18.6 Related Work

As described in Section 18.2, traditional data loggers are typically large and expensive or use intrusive probes. Other habitat monitoring studies install weather stations an "insignificant distance" from the area of interest and make coarse generalizations about the environment. Instead, biologists argue for the miniaturization of devices that may be deployed on the surface, in burrows, or in trees.

Habitat monitoring for WSNs has been studied by a variety of other research groups. Cerpa et. al. [7] propose a multi-tiered architecture for habitat monitoring. The architecture focuses primarily on wildlife tracking instead of habitat monitoring. A PC104 hardware platform was used for the implementation with future work involving porting the software to motes. Experimentation using a hybrid PC104 and mote network has been done to analyze acoustic signals [30], but no long term results or reliability data has been published. Wang

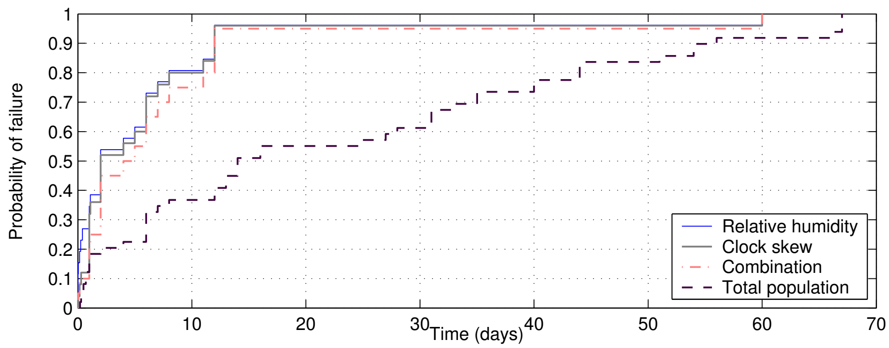

Figure 18.12. Cumulative probability of node failure in the presence of clock skew and anomalous humidity readings compared with the entire population of nodes.

et. al. [29] implement a method to acoustically identify animals using a hybrid iPaq and mote network.

ZebraNet [18] is a wireless sensor network design for monitoring and tracking wildlife. ZebraNet uses nodes significantly larger and heavier than motes. The architecture is designed for an always mobile, multi-hop wireless network. In many respects, this design does not fit with monitoring the Leach's Storm Petrel at static positions (burrows). ZebraNet, at the time of this writing, has not yet had a full long-term deployment so there is currently no thorough analysis of the reliability of their sensor network algorithms and design.

The number of deployed wireless sensor network systems is extremely low. There is very little data about long term behavior of sensor networks, let alone wireless networks used for habitat monitoring. The Center for Embedded Network Sensing (CENS) has deployed their Extensible Sensing System [11] at the James Mountain Reserve in California. Their architecture is similar to ours with a variety of sensor patches connected via a transit network that is tiered. Intel Research has recently deployed a network to monitor Redwood canopies in Northern California and a second network to monitor vineyards in Oregon. Additionally, we have deployed a second generation multihop habitat monitoring network on Great Duck Island, ME. As of this writing, these systems are still in their infancy and data is not yet available for analysis.

18.7 Conclusion

We have presented the need for wireless sensor networks for habitat monitoring, the network architecture for realizing the application, and the sensor network application implementation. We have shown that much care must be taken when deploying a wireless sensor network for prolonged outdoor operation keeping in mind the sensors, packaging, network infrastructure, application software. We have analyzed environmental data from one of the first outdoor deployments of WSNs. While the deployment exhibited very high node failure rates and failed to produce meaningful data for the disciplinary sciences, it yielded valuable insight into WSN operation that could not have been obtained in simulation or in an indoor deployment. We have identified sensor features that predict a 50% node failure within 4 days. We analyzed the application-level data to show complex behaviors in low levels of the system, such as MAC-layer synchronization of nodes.

Sensor networks do not exist in isolation from their environment; they are embedded within it and greatly affected by it. This work shows that the anomalies in sensor readings can be used to predict node failures with high confidence. Prediction enables pro-active maintenance and node self-maintenance. This insight will be very important in the development of self-organizing and self-healing WSNs.

Notes

Data from the wireless sensor network deployment on Great Duck Island can be view graphically at `http://www.greatduckisland.net`. Our website also includes the raw data for researchers in both computer science and the biological sciences to download and analyze.

This work was supported by the Intel Research Laboratory at Berkeley, DARPA grant F33615-01-C1895 (Network Embedded Systems Technology "NEST"), the National Science Foundation, and the Center for Information Technology Research in the Interest of Society (CITRIS).

References

[1] Julia Ambagis. Census and monitoring techniques for Leach's Storm Petrel (Oceanodroma leucorhoa). Master's thesis, College of the Atlantic, Bar Harbor, ME, USA, 2002.

[2] John G. T. Anderson. Pilot survey of mid-coast maine seabird colonies: An evaluation of techniques. In *Report to the State of Maine Department of Inland Fisheries and Wildlife*, Bangor, ME, USA, 1995.

[3] Alexis L. Blackmer, Joshua T. Ackerman, and Gabrielle A. Nevitta. Effects of investigator disturbance on hatching success and nest-site fidelity in a long-lived seabird, Leach's Storm-Petrel. *Biological Conservation*, 2003.

[4] Nirupama Bulusu, Vladimir Bychkovskiy, Deborah Estrin, and John Heidemann. Scalable, ad hoc deployable, RF-based localization. In *Pro-*

ceedings of the Grace Hopper Conference on Celebration of Women in Computing, Vancouver, Canada, October 2002.

[5] Vladimir Bychkovskiy, Seapahn Megerian, Deborah Estrin, and Miodrag Potkonjak. Colibration: A collaborative approach to in-place sensor calibration. In *Proceedings of the 2nd International Workshop on Information Processing in Sensor Networks (IPSN'03)*, Palo Alto, CA, USA, April 2003.

[6] Karen Carney and William Sydeman. A review of human disturbance effects on nesting colonial waterbirds. *Waterbirds*, 22:68–79, 1999.

[7] Alberto Cerpa, Jeremy Elson, Deborah Estrin, Lewis Girod, Michael Hamilton, and Jerry Zhao. Habitat monitoring: Application driver for wireless communications technology. In *2001 ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean*, San Jose, Costa Rica, April 2001.

[8] Dan Chalmers and Morris Sloman. A survey of Quality of Service in mobile computing environments. *IEEE Communications Surveys*, 2(2), 1992.

[9] Benjie Chen, Kyle Jamieson, Hari Balakrishnan, and Robert Morris. Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. In *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking*, pages 85–96, Rome, Italy, July 2001. ACM Press.

[10] Eveready Battery Company. Energizer no. x91 datasheet. `http://data.energizer.com/datasheets/library/primary/alkaline/energizer_e2/x91.pdf`.

[11] Michael Hamilton, Michael Allen, Deborah Estrin, John Rottenberry, Phil Rundel, Mani Srivastava, and Stefan Soatto. Extensible sensing system: An advanced network design for microclimate sensing. `http://www.cens.ucla.edu`, June 2003.

[12] David C.D. Happold. The subalpine climate at smiggin holes, Kosciusko National Park, Australia, and its influence on the biology of small mammals. *Arctic & Alpine Research*, 30:241–251, 1998.

[13] Tian He, Brian Blum, John Stankovic, and Tarek Abdelzaher. AIDA: Adaptive Application Independant Data Aggregation in Wireless Sensor Networks. *ACM Transactions in Embedded Computing Systems (TECS), Special issue on Dynamically Adaptable Embedded Systems*, 2003.

[14] Jason Hill and David Culler. Mica: a wireless platform for deeply embedded networks. *IEEE Micro*, 22(6):12–24, November/December 2002.

[15] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David Culler, and Kristofer Pister. System architecture directions for networked sensors.

In *Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-IX)*, pages 93–104, Cambridge, MA, USA, November 2000. ACM Press.

[16] Chuck Huntington, Ron Butler, and Robert Mauck. *Leach's Storm Petrel (Oceanodroma leucorhoa)*, volume 233 of *Birds of North America*. The Academy of Natural Sciences, Philadelphia and the American Orinthologist's Union, Washington D.C., 1996.

[17] Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin. Directed diffusion: a scalable and robust communication paradigm for sensor networks. In *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking*, pages 56–67, Boston, MA, USA, August 2000. ACM Press.

[18] Philo Juang, Hidekazu Oki, Yong Wang, Margaret Martonosi, Li-Shiuan Peh, and Daniel Rubenstein. Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with ZebraNet. In *Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)*, pages 96–107, San Jose, CA, USA, October 2002. ACM Press.

[19] Bhaskar Krishanamachari, Deborah Estrin, and Stephen Wicker. The impact of data aggregation in wireless sensor networks. In *Proceedings of International Workshop of Distributed Event Based Systems (DEBS)*, Vienna, Austria, July 2002.

[20] Jie Liu, Patrick Cheung, Leonidas Guibas, and Feng Zhao. A dual-space approach to tracking and sensor management in wireless sensor networks. In *Proceedings of the 1st ACM International Workshop on Wireless Sensor Networks and Applications*, pages 131–139, Atlanta, GA, USA, September 2002. ACM Press.

[21] Samuel Madden, Michael Franklin, Joseph Hellerstein, and Wei Hong. TAG: a Tiny AGgregation service for ad-hoc sensor networks. In *Proceedings of the 5th USENIX Symposium on Operating Systems Design and Implementation (OSDI '02)*, Boston, MA, USA, December 2002.

[22] Alan Mainwaring, Joseph Polastre, Robert Szewczyk, David Culler, and John Anderson. Wireless sensor networks for habitat monitoring. In *Proceedings of the 1st ACM International Workshop on Wireless Sensor Networks and Applications*, pages 88–97, Atlanta, GA, USA, September 2002. ACM Press.

[23] Ian Nisbet. Disturbance, habituation, and management of waterbird colonies. *Waterbirds*, 23:312–332, 2000.

[24] Onset Computer Corporation. HOBO weather station. `http://www.onsetcomp.com`.

[25] Kris Pister, Barbara Hohlt, Jaein Jeong, Lance Doherty, and J.P. Vainio. Ivy: A sensor network infrastructure for the University of California, Berkeley College of Engineering. `http://www-bsac.eecs.berkeley.edu/projects/ivy/`, March 2003.

[26] Joseph Polastre. Design and implementation of wireless sensor networks for habitat monitoring. Master's thesis, University of California at Berkeley, Berkeley, CA, USA, 2003.

[27] Bruno Sinopoli, Cory Sharp, Luca Schenato, Shawn Schaffert, and Shankar Sastry. Distributed control applications within sensor networks. *Proceedings of the IEEE*, 91(8):1235–1246, August 2003.

[28] Marco Toapanta, Joe Funderburk, and Dan Chellemi. Development of Frankliniella species (Thysanoptera: Thripidae) in relation to microclimatic temperatures in vetch. *Journal of Entomological Science*, 36:426–437, 2001.

[29] Hanbiao Wang, Jeremy Elson, Lewis Girod, Deborah Estrin, and Kung Yao. Target classification and localization in habitat monitoring. In *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2003)*, Hong Kong, China, April 2003.

[30] Hanbiao Wang, Deborah Estrin, and Lewis Girod. Preprocessing in a tiered sensor network for habitat monitoring. *EURASIP JASP Special Issue on Sensor Networks*, 2003(4):392–401, March 2003.

[31] Kamin Whitehouse and David Culler. Calibration as parameter estimation in sensor networks. In *Proceedings of the 1st ACM International Workshop on Wireless Sensor Networks and Applications*, pages 59–67, Atlanta, GA, USA, September 2002. ACM Press.

[32] Ya Xu, John Heidemann, and Deborah Estrin. Geography-informed energy conservation for ad hoc routing. In *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking*, pages 70–84, Rome, Italy, July 2001. ACM Press.

[33] Jerry Zhao, Ramesh Govindan, and Deborah Estrin. Computing aggregates for monitoring wireless sensor networks. In *Proceedings of the 1st IEEE International Workshop on Sensor Network Protocols and Applications*, Anchorage, AK, USA, May 2003.