An Analysis of a Large Scale Habitat Monitoring Application

Robert Szewczyk, Alan Mainwaring, Joseph Polastre and David Culler

April 9, 2004

Abstract

Habitat and environmental monitoring is a driving application for wireless sensor networks. We present second generation sensor networks deployed during the summer and autumn of 2003. These networks produced unique datasets for both systems and biological analysis. This paper focuses on nodal and network performance, with an emphasis on lifetime, reliability, and network performance. Of particular interest are the the static and dynamic aspects of single and multi-hop networks over the course of several months as these networks run completely unattended. This analysis sheds light on a number of network design issues from network deployment, through selection of power sources to optimizations of routing decisions.

1 Introduction

Sensor networks offer enormous potential benefits for habitat and environmental monitoring. They enable precision sensing at the spatial and temporal scale of the phenomena of interest rather than of the human observer or sparsely deployed instrument. Automated observation and measurement not only exposes the long-term unattended operation potential of low-power microelectronics, but it minimize observer effects and study site intrusions as well. The efficiency of data collection as well as the quality of the data should increase, while the costs as compared with traditional human-centric methods should decrease. Moreover, sensor networks offer new capabilities, such as near real-time data access via the Internet to recent and archival data. This further minimizes observer effects and site intrusions, and establishes a new paradigm of data publication and dissemination.

A broad class of important habitat and environmental monitoring applications are within reach of contemporary technology. They are both scientific and commercial in nature, representing, for example, applications in ecology, plant physiology, and precision agriculture. They application share a common structure where fields of mostly stationary sensors are *queried* periodically for their sensor readings and report results to a central data center. The two principal characteristics that distinguish points within the space are the sensor sampling rate and the network bandwidth requirements. The application discussed in this paper resides in the low sample rate, low bandwidth quadrant. The architecture and its implementation as described herein is quite general and can support applications in other quadrants with varying sampling and bandwidth demands.

To demonstrate the practical benefits of sensor networks and to develop first hand experience with sensor network technologies in the field, we have several deployed sensor networks of increasing scale in a rugged outdoor environment. The networks have produced a unique datasets for the life sciences and biological analyses are underway. From a systems perspective, they contain a great deal of information. For example, some sensor nodes ran unattended for 3 to 4 months and there are packet logs for them. There are packet traces showing emergent network behavior after weeks of operation. And although the application on each sensor node was simple, it exhibited ample amounts of interesting behaviors after days, weeks and months of runtime. Bit by bit we're gaining practical experience from real world deployments.

This paper presents an analysis of that data collected during the summer and autumn of 2003. It is organized as follows: Section 2: system architecture and realization. Section 3: data Analysis Section 4: discussion Section 5: Related Works and Section 6: concludes.

2 System

As preparation for the performance analyses in Section 3, this section presents the system architecture and its implementation. In particular, we describe the tiered network architecture typical of habitat monitoring applications and the design and implementation of its major components.

2.1 Architecture

The system has the tiered architecture depicted in Figure 1. The lowest level consists of *sensor nodes* that perform communication, computation and sensing. They are typically deployed in dense patches. Depending on the application, their spatial distribution within a *sensor patch* can vary from transects to patches of nodes to three-dimensional monitoring. Each sensor patch has a *gateway* that sends data from the patch through a *transit network* to a remote *base station* via the *base station gateway*. We expect mobile *field tools* will allow on-site users to interact with the base station and sensor nodes to aid the deployment, debugging and management of the installation. The base station provides Internet connectivity and database services. It needs to provide high availability and be able to deal with disconnections. Remote management facilities are a crucial feature of a base station. Typically the sensor data is replicated via Internet. These replicas are located wherever it is convenient for the users of the system. In this formulation, a *sensor network* consists of one or more sensor patches spanned by a common transit network and base station.

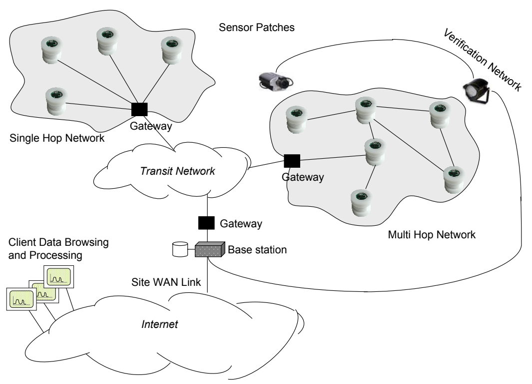

Figure 1: Architecture of the habitat monitoring system

Sensor nodes are small, battery-powered devices capable of general purpose computation, bi-directional wireless communication, and application-specific sensing. The sizes of nodes are commensurate with the scale of the phenomenon to be measured. Their lifetime varies with duty cycle and sensor power consumption; it can be months or years. They use analog and digital sensors to sample their environment, and perform basic signal processing, e.g., thresholding and filtering. Nodes communicate with other nodes either directly or indirectly by routing through other nodes. Nodes communicate directly to their gateway in a *single hop* network or through intermediates in a *multihop* network.

Independent *verification networks* collect baseline data from reference instruments that are used for sensor calibration, data validation, and establishing ground truth. Typically verification networks utilize conventional technologies and are limited in extent due to cost, power consumption and deployment constraints.

2.2 Concrete Deployment

Our habitat monitoring deployment is a concrete realization of the general architecture from Figure 1. The sensor node platform was a Mica2Dot, a repackaged Mica2 mote produced by Crossbow, Inc. with a 1 inch diameter form factor. The mote used an Atmel ATmega128 microcontroller running at 4 MHz, a 433 MHz radio from Chipcon operating at 40Kbps, and 512KB of flash memory. The mote interfaced to sensors digitally using I2C and SPI serial protocols and to analog sensors using the onboard ADC. The small diameter circuit boards allowed a cylindrical assembly where where sensor boards were endcaps with mote and battery internal. Sensors could be exposed on the endcaps; internal components could be protected by O-rings and conformal coatings. This established a mechanical design where sensor board and battery diameters were in the 1 to 1.5 inch range but the height of the assembly could vary.

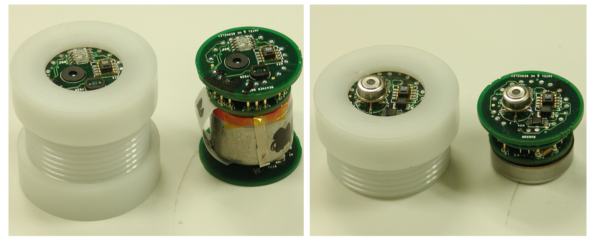

Figure 2: Mote configurations used in the deployment: climate mote (left) and presence mote (right)

We designed two different motes for the application, *presence motes* for detecting occupancy using non-contact infrared thermopile sensors and *climate motes* for monitoring microclimates. The presence motes, being the more electromechanically constrained of the two forms, pushed miniaturization more than the climate mote. The presence mote had to be extremely small to be deployed unobtrusively in nests typically only a few centimeters wide. Batteries with lithium chemistries were chosen with discharge curves where voltages remained in tolerance for mote, radio and sensors almost to the end of the battery lifetime. We elected not to use a DC boost regulator to reduce noise, standby power consumption, and power consumption as the battery discharges. Its operation and the quality of its sensor readings depends upon the voltage remaining within tolerance. As the voltage falls outside the operating range of the mote, radio, or sensors, the results are unpredictable.

2.2.1 Presence and Climate Motes

Presence motes monitor temperature, humidity and occupancy of nesting burrows using non-contact passive infrared temperature sensors. They have two sensors: a Melexis MLX90601 non-contact temperature module and a Sensirion SHT11 temperature and humidity sensor. The Melexis part measures both ambient temperature ($\pm$ 1$\circ$C) and object temperature ($\pm$ 2$\circ$C). The Sensirion part measures relative humidity ($\pm$ 3.5% but typically much less) and ambient temperature ($\pm$ 0.5$\circ$C), which is used internally for temperature compensation. The motes used 3.6V Electrochem batteries rated at 1Ahr, with a 1mA rated discharge current and a maximum discharge of 10mA. The 25.4mm diameter by 7.54mm tall dimensions of the cell were well suited for the severely size constrained presence enclosures.

Climate motes monitor temperature, humidity, and barometric pressure. (They also measure ambient and incident light, both broad spectrum as well as photosynthetically active radiation, but these were not used in this application.) They have the following sensors: Sensirion SHT11, Intersema MS5534A barometer, 2 TAOS TSL2550 light sensors, and 2 Hamamatsu S1087 photodiodes. The Intersema measures barometric pressure ($\pm$ 1.5 mbar) and ambient temperature ($\pm$ 0.8$\circ$C) used for compensation. The motes used 2.8V SAFT LO34SX batteries rated at 860mAhr, with a 28mA rated discharge current and a maximum discharge exceeding 0.5A. A 25.6mm diameter by 20.3mm height were similar to the Electrochem, permitting a similar packaging technique. The battery exhibits a flat voltage profile for nearly its entire lifetime.

2.2.2 Single and Multihop Networks

The first network deployed was an elliptical single hop network The total length of the ellipse was 57 meters. The network gateway was at the western edge. Nodes in this network perform no routing, they sampled their sensors every 5 minutes and sent results to their gateway. The gateway mote used a TESSCO 0.85 dBi omnidirectional antenna for the sensor patch and a Hyperlink 14dBi yagi for a long distance point-to-point link to the base station over a distance of about 120 meters.

The second deployed network was a multihop network with a kite-shape to the southwest and a tail to the northeast. Its total length is 221 meters with a maximum width of 71m at the southwest but it narrows to 8m at the northeast. Nodes in this network sampled every 20 minutes and routed packets destined for its gateway. An identical antenna configuration to the single hop network was used, with a few important distinctions. The gateway mote in the single hop network operated at 433 MHz, and its packets were retransmitted on the transit network at 915 MHz. The gateway mote in the multihop network operated at 435 MHz, and its packets were retransmitted on the transit network at 916 MHz as well. The 128kHz frequency bandwidth of the radio allowed the two sensor network deployments to operate independently but share a common transit network and upstream infrastructure.

2.2.3 Transit and Verification Networks

There were two redundant transit networks operating on identical channels that connected the single and multihop networks to the base station gateways. The gateway nodes were co-located in the sensor patches for convenience.

To understand the correlation between infrared sensor readings from presence motes and true occupancy, the verification network collected 15 second movies using in-burrow cameras equipment with IR illuminators. Using a combination of off-the-shelf equipment–cables, 802.11b, power over Ethernet midspans, and Axis 2401 camera servers–eight sites were instrumented and five operated successfully. All verification network equipment was physically distinct from the transit and sensor networks with the exception of the laptops at the base station. Scoring the movies by hand in preparation for analysis with sensor data is underway.

2.2.4 WAN and Base Station

A DirecWay 2-way satellite system with Optistreams service provided WAN connectivity with 5 globally routable IP addresses for the base stations and other equipment at the study site. This provided access to the PostgreSQL relational databases on the laptops, the verification image database, administrative access to network equipment, multiple pan-tilt-zoom webcams and network enabled powerstrips. A remote server computes the set of database insertions and uploads the differences every 20 minutes via the satellite link. Should the link be unavailable, updates are queued for delivery. The remote databases may be queried by replicas for missing data upon link reconnection. Although the upstream bandwidth is small (128Kbps), we have not witnessed an overrun situation. The base station, satellite link and supporting equipment were powered off a standalone photovoltaic system with an average daily generating capacity of 6.5kWh/day in the summer.

2.3 Media Access and Routing

The network software was designed to be simple and predictable. The radio was duty cycled in our deployments with a technique called low power listening. Low power listening periodically wakes up the node, samples the radio channel for activity, and then returns to sleep if the channel is idle. Packets sent to a low power listening node must be long enough that the packet is detected when the node samples the channel for activity. Once activity is found, the node stays awake and receives the packet, otherwise it goes back to sleep.

The single hop network utilized low power listening but with normal sized packets to its transit gateway. Packets with short preambles can be used because the gateway does not duty cycle the radio–instead it is capable of always receiving packets. The sensor nodes periodically transmitted their sensor readings. They used low power listening to allow the base station to issues commands to change their sample rate, to read calibration data them, and to *ping* the node for health and status information.

The multihop network integrated low power listening with adaptive multihop routing developed by Woo [11]. Each node selected its parents by monitoring the channel and using the path it expected to be most reliable. Nodes periodically broadcasted their link quality estimates to their neighbors. This data was used to find bidirectional reliable links. The nodes communicated with each other using low power listening and long packets. We estimated a network neighborhood size of 10 nodes and calculated that a 2.2% radio duty cycle would maximize the node's lifetime. We deployed the nodes with these settings and allowed them to self-organize and form the network routing tree.

The software running on the presence and climate motes implements a data stream architecture. Each mote samples it sensors and sends its readings to the base station once per sampling period. Single hop motes sample every five minutes and multihop motes every twenty

Table 1: Estimated power consumption of the application

Task	Energy (mAs)	Singlehop current (μA)	Multihop current (μA)
Baseline sleep	-	20	20
Timer	0.0012	22	22
Incoming packet detection (low power listening)	0.166	166	332
Packet transmission (short preamble)	1.4	5	-
Packet transmission (long preamble)	14.0	-	23
Climate sensing	13.0	43	11
Occupancy sensing	12.6 (est)	42	10.5
Total climate mote w/o routing & overhearing)	-	256	408
Total presence mote w/o routing & overhearing)	-	255	407.5

minutes. Each mote also listens for incoming packets and takes appropriate action upon receipt. If the packet is destined for the mote, it parses the contents and takes appropriate action, e.g., changing its sampling rate or responding to a diagnostic ping message. Otherwise, the packet is forwarded by the routing subsystem towards its destination.

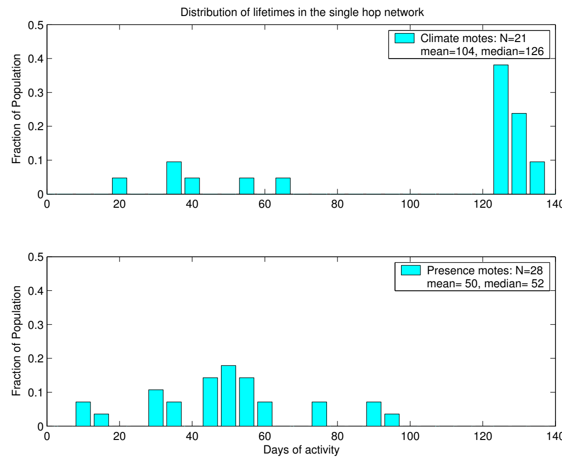

Figure 3: Node lifetimes in the single hop network

3 Analysis

This section analyzes the performance of the sensor networks from a systems perspective, analyzing packet power, and network properties. During their combined 115 days of operation, the networks produced in excess of 650000 sensor observations. The first deployment started June 8th with the incremental installation of a single hop network. At its peak starting June 16th, the network had 48 motes (28 presence and 21 climate). A second deployment began July 8th with the incremental installation of a multihop network. At its peak starting August 5th, the network had a total of 98 motes (62 presence and 36 climate).

3.1 Lifetime

The design goal for the networks was to provide observations throughout the field season. We first examine the lifetime of single- and multihop motes, and compare the achieved performance with estimates. The simplicity and regularity of the single-hop network makes the power analysis straightforward. Table 1 summarizes the main sources of power drain in the system. Excluding packet reception costs, our original estimates called for a 256 μA and 255 μA current drain in the single hop climate and presence motes respectively – or 138 days of operation assuming a full battery discharge with the SAFT lithium cells.

Figure 3 shows the distribution of mote lifetimes in days in the single hop networks for both climate and presence motes. The key distinctions between the climate and presence lifetimes are the significantly different mean and median days of operating. Whereas over 75% of the climate mote population lives for 120 days or longer, only

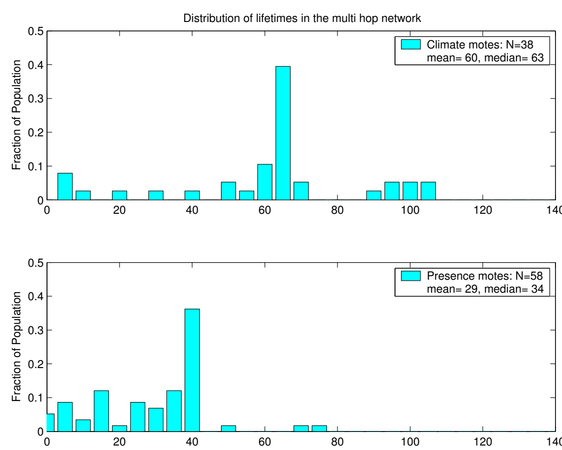

Figure 4: Node lifetimes in the multihop network

50% of the presence mote population lives for 52 days or longer. We have met the lifetime goal for the climate motes – motes were operational at the end of the deployment. We believe the different characteristics of batteries powering the two different kinds of devices explain this; a more detailed power usage discussion follows below.

Figure 4 shows the distribution of mote lifetimes in days in the multihop network for climate and presence motes. Just as in the single-hop case, the presence motes have considerably shorter lifetimes than the climate motes: whereas over 62% of the climate mote population lives for 63 days or longer, only 42% of the presence mote population lives for 34 days or longer. The multihop network operates for a much shorter period. This is to be expected: the average current consumption, excluding listening and packet forwarding, is estimated at 400 μA, or a mere 90 days of lifetime (Table 1). However, these lifetimes still fall short of the predictions. To locate possible causes of the shortened operation, we turn to analyzing the battery voltages during the final hours of operation.

Figure 5 shows the average voltages reported by motes during their last three hours of operation. Mote that may have continued to operate but that were unable to communicate successfully with the base station were considered to have died. A threshold is highlighted on each graph that indicates the lowest voltage at which a mote will reliably operate.

The passing remnant of hurricane Isabel required shutting down the base station between September 15th to October 9, which produced a 23-day-long outage. The base stations were revived until October 20th when they were shutdown for the season because of persistent early winter storms. Motes that ceased operation during each shutdown have their lifetimes truncated; this leads to formation of indicated clusters in the scatter plot.

Figure 5 shows that of the original 21 single hop climate motes, 15 are still operating at the end of the season on October 20th. Improper sealing may play a part in shortening the lifespan of the 6 remaining motes. Only 3 of the multihop climate motes were prematurely terminated because of clear battery problems. We note that the all multihop motes experience increased power consumption when they actively forward packets from the network or when they overhear packets from a busy neighborhood; we attempt to quantify these issues below.

Both kinds of multihop motes show significant clustering of points above the threshold voltage. If the base station was operating properly, these motes would have operated for another few days. This leads us to conclusion that the shutdown had a significant impact on mean and median recorded lifetimes.

The presence motes also seem to run their battery to exhaustion. In the single hop case, motes drain their battery thoroughly: 8 motes fall below the conservative 3V threshold. We observe a very sharp voltage drop at the end of battery capacity – it is possible that the remaining experienced a drop so rapid that it was not recordable. We discuss a number of improvements in collection of diagnostic information in Section 4. The multihop presence motes exhaust their supply very rapidly. We conjecture that it is caused by two factors: inability of battery to source current for the long preambles and by excessive overhearing of other multihop traffic. We attempt to verify the second hypothesis below.

3.2 Packet Delivery Effectiveness

We examine whether the packet delivery mechanisms was effective. Recall that both single- and multihop networks used the streaming data architecture and relied on oversampling (2x) to deliver the required sampling rates. Previous studies using the same multihop routing algorithm have indicated a 90% packet yield across a 6-hop network [11]. We study the packet delivery over both a short term and a long term and determine whether the streaming data approach, with no acknowledgments or retransmissions, was sufficient. First, we note that the networks were operating at such a low duty cycle that collisions or network congestion should be insignificant.

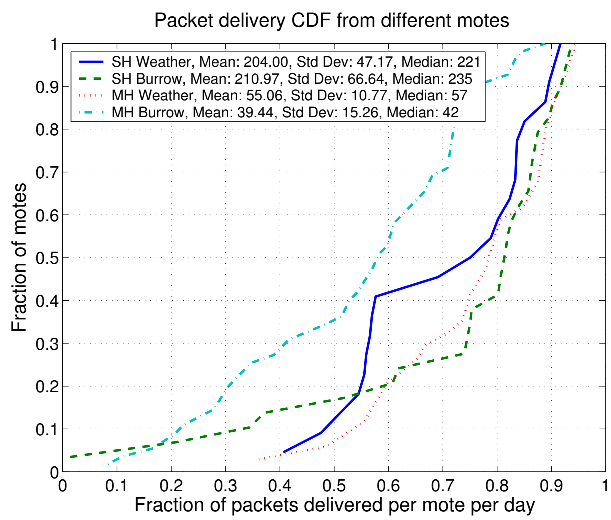

Figure 6: Packet delivery CDF on the first day of complete deployment of the single-hop (June 18, 2003) and the multihop network (August 6, 2003). multihop weather motes had a median packet delivery of 42 packets (58%). All other motes achieved a median packet yield of over 70%.

Figures 6 and 7 show the packet yields from the 4 different kinds of motes. Figure 6 shows the packets delivered to the base station during the first full day of deployment of each network. The results match the expectations built up in the indoor lab environment: the climate motes, on average, deliver well over 70% of the packets. The single-hop presence motes deliver similar yield. The multihop presence motes perform worse (with a median yield of 58%), but within expectations: the application was designed to oversample the biological signal, and consultations with biologists indicated that the delivered sampling rate would be sufficient over the long term. Figure 7 tells a different story. It plots the CDF of the packet yields for

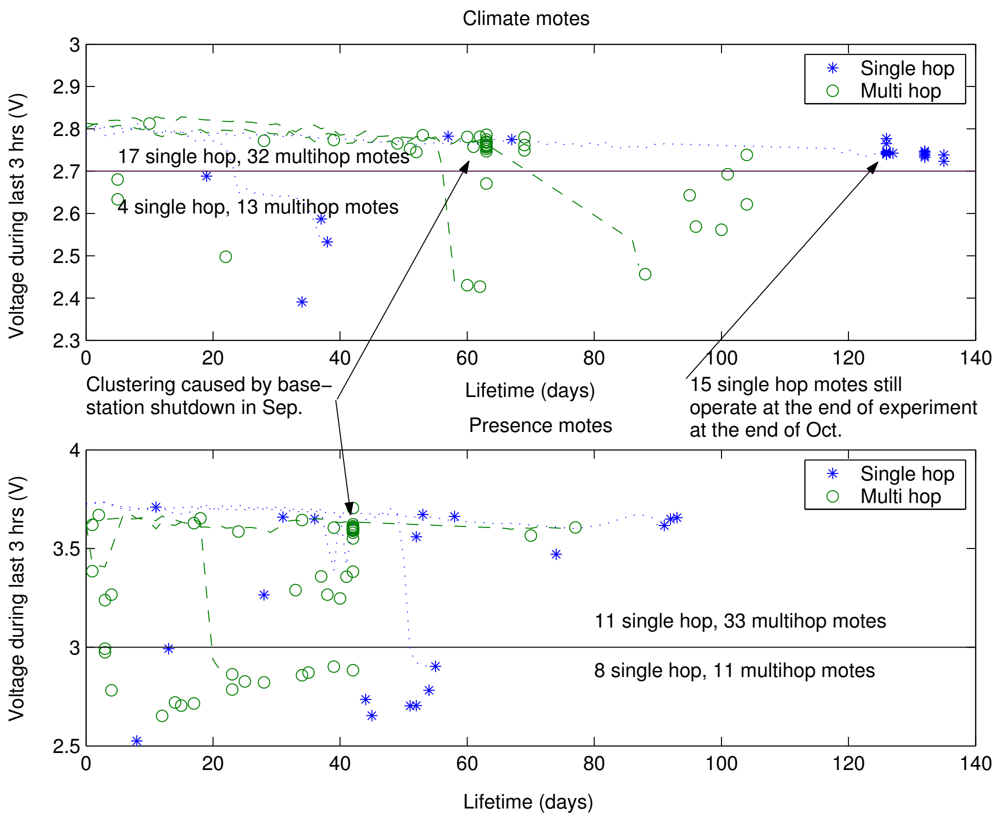

Figure 5: Mote voltage at the end of operation. The cutoff voltages are conservative and have been selected from battery datasheets. The dotted and dashed trails illustrate average daily voltages from representative motes; these are in line with the discharge curves in the datasheets.

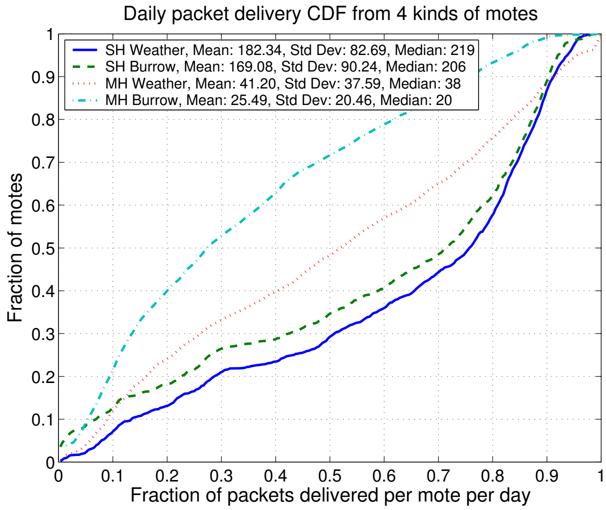

Figure 7: Daily packet delivery CDF over the entire length of the deployment. Motes in the single-hop network deliver a median yield of 70% of packets. The multihop network fares much worse, with multihop burrow nodes delivering a median yield of just 28%.

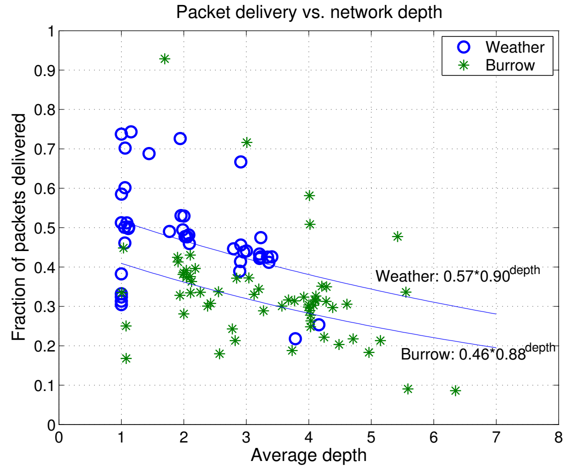

Figure 8: Packets delivered v. average depth. We model packet delivery as a function of the Al^d, where l is link quality, d is the average depth in the routing tree, and A represents packets lost for reasons unrelated to multihop routing, like base station or transit network outage

each mote on every day that mote was active, i.e. delivered a packet. The results are satisfactory for the single hop: the median yield still remains over 70%. In contrast to Figure 6, the distribution contains many mote-days with substandard performance. On the first day, no mote delivered fewer than 35% of its packets; over the course of the deployment nearly a quarter of the motes performed that badly. The situation get even worse when we look at the multihop network: the climate motes deliver a median yield of a mere 28%, a performance that begins to jeopardize the required delivery rates. Some portion of this is due to partial day outages at the base station, but other effects are also at play. This unexpected behavior illustrates how results from lab experiment may be optimistic in the long term outdoor deployments.

We conclude that the simple streaming approach was sufficient in the single hop network, but fell short in the multihop case.

Next, we evaluate the effectiveness of routing. Figure 8 depicts the packet yield for each multihop mote as a function of mote's average depth in the routing tree, weighted by number of packets. Recall that the multihop routing software performs no retransmissions. Given that we have observed non-negligible packet loss in the single hop network, we expect to observe the same behavior per hop in the multihop case. In turn, this behavior would produce a packet yield that is a decaying exponential function of depth in the tree. Indeed, the data in Figure 8 shows such behavior. If we model the packet yield P as l^d where l is link quality and d is a packet depth, then the best fit link quality l is 0.72 and the mean squared error is 0.03, for both climate and presence motes. This matches closely the mean packet delivery in the single hop network in Figure 6 as well as the link quality data reported in [11] (0.7 and 0.73 for climate and burrow, respectively), and is considerably better than the mean packet yield over the lifetime of the single hop network.

We note that a fixed fraction of data was never recorded in the database – either because of losses in the transit network (with potentially different link characteristics) or because of the base station outage. It is therefore appropriate to model the packet yield as Al^d, where A corresponds to that deterministic loss for all motes. The best fit parameters curves from this model are shown in Figure 8. The MSE using this model is 0.015 for presence motes and 0.025 for climate motes. The average link quality estimate of nearly 0.9 shows that the routing layer is capable of picking high quality links. The deterministic loss A is very high: 0.57 and 0.46 for climate and presence motes respectively.

The observed packet yields make a strong case for more reliable routing subsystems. To reduce the exponential drop as a function of depth, a simple hop-by-hop retransmission strategy could be used. In addition, custody transfer mechanisms could be used to improve or eliminate the deterministic loss.

3.3 Multihop network structure

Low-power listening, as used in the multihop network, lowers the cost of listening, while increasing the cost of both transmitting and receiving. Overhearing packets is also costly, since there is no early packet rejection. Consequently, the connectivity of the motes in the multihop network has a large impact on the power consumption of individual nodes. We did not log the information about neighborhood sizes or packet overhearing, at best we can make conservative estimates from the dataset. We briefly analyze the structure of the multihop routing graph and point out opportunities for improvement.

Given their reduced power supply, the presence motes should never be used to route packets. In the deployment, we did not specifically enforce this constraint, consequently 48 of the presence motes were used as parents at some point during their operation; on average these motes routed 75 packets, with maximum of 570 packets. While the packets routed through the presence motes were a small fraction of the overall traffic (3600 out of 120000 packets), this was an optimization that would have been easy to implement.

Larger savings may be obtained from reducing overhearing. We calculate the number of packets routed through each mote. We observe that the communication between parent and child is bidirectional – the algorithm for parent selection ensures that. Therefore we can approximate the number of packets overheard by any mote by summing all packets routed through *any* of the parents this mote selected. On average, each of the presence motes overheard nearly 17000 packets, a significant burden on the limited power supply. Overhearing reduction is more complex to implement than restricted route selection.

Finally, we examine the distribution of children of every node. Figure 9 shows that a large portion of the network – 32% – consisted of leaf nodes that never routed any packets. In a stable network topology, the leaf nodes (most of which are climate motes) can dramatically reduce the discovery rate.

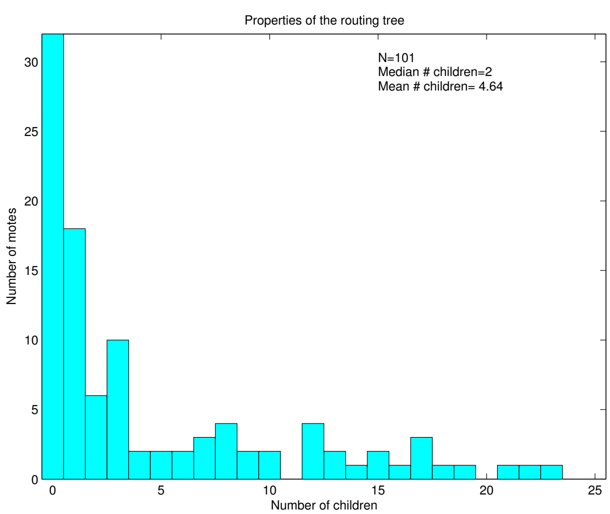

Figure 9: Distribution of children in the routing graph.

3.4 Routing Stability

Lab measurements have documented routing stability over periods of hours. We evaluate the stability over weeks and moths in a real world deployment. Previously published results, show the random fluctuations in link quality in the real world. In addition, the logistic of the application – incremental installation, and node attrition – contribute significantly to parent switching.

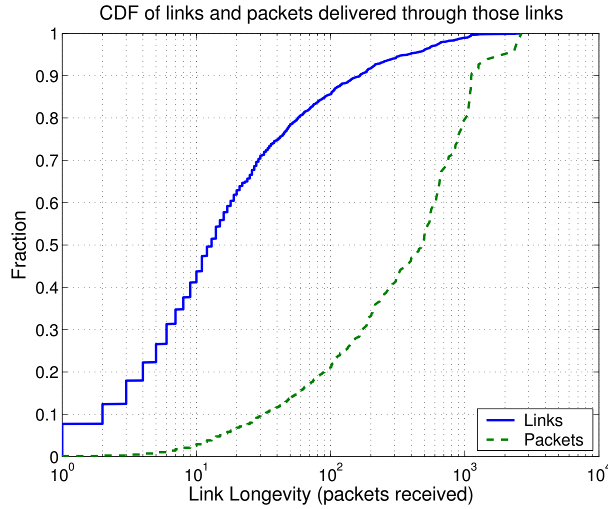

Figure 10: CDFs of parent-child relationship lengths and packets delivered through those links. Long-lived, stable links (ones that delivered more than 100 packets) constitute 15% of all links, yet they are used for more than 80% of the packets.

We begin the analysis by looking at the lengths of parent-child relationships within the routing tree. Figure 10 shows a CDF of both the links and the packets delivered over links. Link longevity is measured as a number of consecutive packets successfully delivered through a particular parent. We only count packets sourced rather than routed through the link since packets are sourced at a fixed rate. Because of great range in the link longevity, it is more appropriate to plot it on a logarithmic scale. Most of the links are short lived: the median link is used to deliver only 13 packets. However, mote of the packets are transmitted over stable stable links. We observe an 80-20 rule: 80% of the packets are delivered over less than 20% of all links. These stable links last for more than 100 packets – or more than a day and a half. While the distribution may be skewed by nodes that communicate directly with the root of the tree, it still can be used to choose the beaconing rates for updating routes.

Another way to look at the stability of the tree is to look at the number of parent changes per time window. Because of the fluctuating network size, we normalize the number of parent changes by the number of motes active over that particular window. Window size has a significant impact on the analysis; we chose a window of 6 hours, which is longer than the median longevity of the link. Figure 11 offers a time series view of the multihop network. In order to understand the stability over time, we must look at two related variables: network size and quality of data logging. Recall that the multihop network was installed in 3 stages, concluding with a deployment of presence motes on August 5. Prior to presence mote installation, we see a stable mote population. The parent change rate spikes rapidly after the installation of each mote group installation, but settles quickly thereafter. After the initial parent turnover, the network is stable; this matches the stability results in [11]. After the deployment of burrow motes the parent change rate remains high at 0.5 for a week. This behavior is likely caused by a set of nodes faced with a choice between a few equally good links. The behavior is not directly caused by changes in population size – a period at the end of August corresponds to a similar mix of motes, is accompanied my motes disappearing, and yet the parent change rate is nearly 0. We also note a high turnover in the last part of the network lifetime – this is related to the network periodically disappearing, and reappearing.

Figure 12 attempts to look at this behavior as a distribution. The takeaway is that the over 55% of time intervals corresponded to times with no changes in the links; 75% had experienced less that 0.1 parent changes per 6 hour interval.

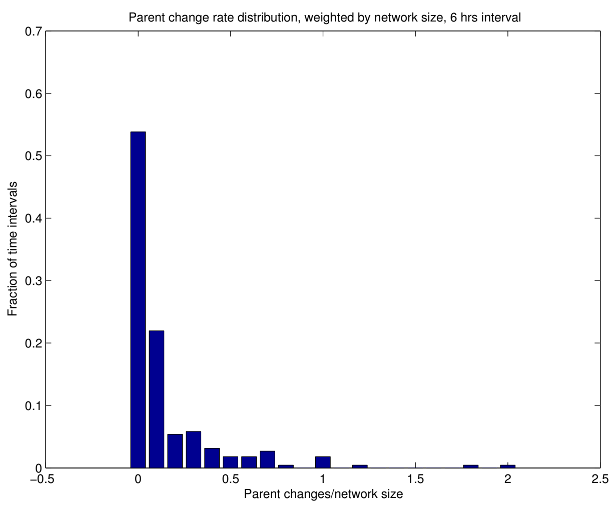

Figure 12: Distribution of parent change rates

4 Discussion

This section discusses insights we have gained from the data analysis as well as deploying and working with sensor networks in the field. In some cases, as in the case of integrated data logging, we recommend functionality for future systems. In others, as in the case of field tools, we talk about our struggles with very primitive tools for the on-site installation and remote monitoring of sensor networks.

4.1 Mote Design

The shorter lifespan of presence motes in the multihop network - nearly 50% - shown in Figure 3 was surprising. The presence motes were a challenge to design and to engineer for damp and dirty environment. We were unable to determine the root cause of presence mote failures although were able to identify a few potential factors. In this section we identify solutions that may assist in future deployments with root cause analysis.

When using batteries with a constant operating voltage, such as the lithium batteries used in our deployment, battery voltage does not indicate how much capacity is remaining. A more adequate measure of how much work has been performed by the node is needed to calculate each node's expected lifetime. Since our deployment, we have implemented energy counters at the MAC layer and in our sensing application. Each counter keeps track of the number of times each operation has occurred (*e.g.*sensing, receiving or transmitting bytes, total amount of time the cpu is active). By keeping track of this data, nodes can report on the number of packets forwarded or overhead. We can improve our lifetime estimate through additional health information from each mote. Metrics, such as these energy counts, are crucial to predicting the performance of the deployed network.

Since the presence motes were so closely integrated with a damp volatile environment, their packaging must protect the internal components of the mote. With only external sensors on our mote, we were unable to detect water breaches in the packaging. We have decided to include a humidity and temperature sensor inside the mote for future platforms. The internal sensor can cause the mote to notify end users of packaging failures.

4.2 Integrated Datalogging

In the current system, packets are sent without any acknowledgments or retransmissions. Sensor readings can be lost at any level of the network from the node to the base station gateway. However, nodes at each level of the network could log data into local non-volatile memory as it heads towards the base station.

A protocol can attempt to incorporate data logging with reliable message delivery or custody transfer protocols to reduce data loss. For example, given 512KB of local storage, 64 byte sensor readings, and a sampling interval of 20 minutes, each node could store 113 days of its own readings. This buffering could mitigate downtime at the added expense of buffering packets along the path. The buffering is free unless the packet is writing to flash memory; writing flash memory is approximately four times more costly than sending the message.

Integrated logging allows nodes to retain data when disconnections occur. A node may be disconnected from another mote (such as a parent), the gateway, the base station gateway, one of the base stations, or the data base service. Since large scale deployments are relatively costly, it may be worth taking measures to retain the data if the reduction in longevity can be tolerated.

4.3 Field Tools

Sensor networks are challenging to install and monitor. Currently a great deal of expertise is required to install these networks; rather we would prefer to enable the spe-

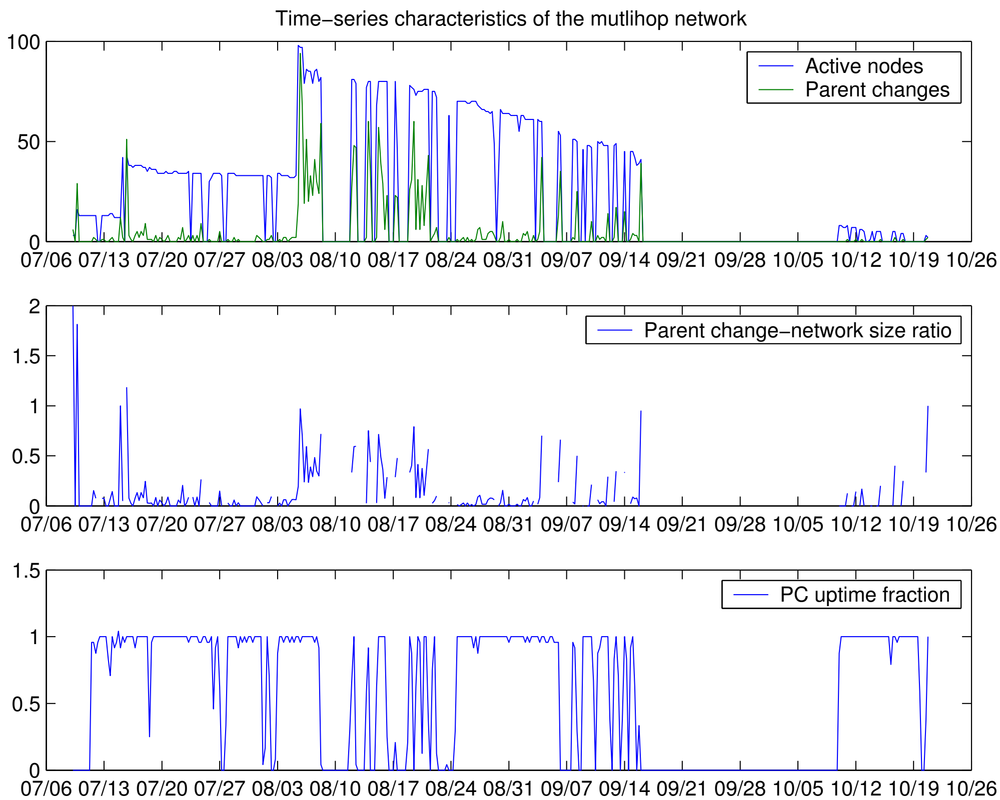

Figure 11: multihop network stability over time. The top graph shows the size of the multihop network. Parent change rate is shown in the middle figure. The bottom graph shows the state of the base station logging the data. The base station exhibits a number of outages that impact both the observed network size and the calculation of parent change rates.

cialist and the non-specialist alike to accomplish that easily. We identify two levels of tools that provide assistance when deploying sensor networks: field tools that run on small, PDA-class devices and client tools that run on larger laptops class machines situated either at the local field station or more distantly at businesses or universities many miles away. For each class, we note the functionality that would be useful from our own experiences in the field.

4.3.1 Field Tools Functionality

1. Run self-check. Before placing a mote in the field, where is may not be touched again for months or years, it should be possible to run a final self diagnostic to verify the mote's health. If the device is healthy, it can be left alone, but otherwise repairs must be taken.

2. Show network neighborhood. While incrementally deploying a network in the field, oftentimes one needs to see whether a mote has been placed within range of the rest of the network. Placement is often guided by non-networking factors, e.g., factors of biological interest.

3. Show network statistics. While within range of the mote, it can be useful to query it for basic packet level statistics. Once a mote has joined a network, it can be useful to monitor how well its networking subsystem is operating before moving onto new locations.

4.3.2 Client Tools Functionality

1. Retask nodes (reprogram the sensor network). From time to time, application software upgrades and patches will become available and it will become necessary to upgrade application software. It should be possible to do this from a field station or remotely on the Internet.

2. Show who-can-hear-whom relationships. Show the graph estimating the radio neighborhoods around each node as well as the routing tree currently in use for the network. This is typical of the diagnostic information that technicians would need to monitor network operation.

3. Show when mote lasted reported in. Report when each mote was last heard from. This is a very common and very useful statistic for the layperson looking for signs of a malfunctioning mote or network.

The usefulness of these tools is proven by a large number of scenarios that arose in the study area. Although the spatial distribution of nodes was driven by the interests of our biologist, these tools can show the density of each level of the multihop network. A simple GUI could display when each node was last heard from. An optional alarming mechanism to notify on-site staff when a node failed to report is needed functionality for non-technical on-site staff.

4.4 The One Touch Mote

For sensor networks to scale to hundreds and thousands of nodes, motes can be touched just once. Assembling a mote, programming it in a test fixture, enclosing it in a package (with no external on/off switch), positioning it the field, and acquiring its survey GPS location is impractical for large numbers of motes. Even the end user who receives a complete, pre-assembled mote ready for deployment faces usability problems of scale. Ideally, the mote should be a completely encapsulated, or packaged, with only the sensors exposed, a non-contact on/off switch, and robust network booting and reprogramming. Issues of programming and network construction should be handled with tools that operate on aggregates of nodes rather than individuals wherever possible.

4.5 Metadata Management

When motes are deployed in the environment, they assume a geospatial context that is critical when interpreting their sensor data. Their locations potentially change; this is one form of metadata that must remain bound to the mote for the duration of its lifetime. When motes are calibrated, coefficients of calibration become another form of metadata that must remain bound to the device for its lifetime. Whatever the form of metadata, it is critical that it remains bound to its mote in the presence of potentially disruptive events, such as being moved or being cleaned. Ambiguity will cast doubt on the data analysis.

4.6 Lessons Learned

This remainder of this section discusses how we would approach and what we would do differently in future deployment.

4.6.1 Deployment Tools

The back-end infrastructure such as the transit network, base stations and relational databases were deployed before the motes so that packet logs were ready to be captured as soon as sensors nodes began to report in. When deploying new motes, it was possible to see their records being inserted into the database and thus know they were alive. This was a primitive but sufficient means of creating a network when combined with the ability to issue SQL queries against the growing database. The queries allowed one to retrieve statistics such as the network size and when motes were last heard from.

A variety of tools have since been developed for data analysis as well as network monitoring. Ideally, these GUIs, visualizations, and statistical tools should be available at deployment time as well to enrich the suite to client tools that are available. The statistics one performs on a corpus of data, such as lifetime analysis or voltage profiling, may have great practical utility during phases of network construction and early days or weeks of operation as well. Many of the graphs in Section 3 would be of interest to practioners deploying their own large scale sensor networks out in the field.

4.6.2 Node Packaging

There were several interdependent constraints on node packaging and power design. We sought a single, waterproof enclosure design with integrated internal battery that was small enough for our occupancy cavity yet with enough interior space to accommodate a larger battery with more capacity for the climate motes. From these constraints, the remaining packaging and power issues and solutions followed. In retrospect, we can more fully appreciate their implications.

Rather than building semi-custom enclosure from off-the-shelf stock plastic, custom machined enclosures may have been advantageous and more cost effective because the current enclosures (1) require adhesive sealant for truly watertight seal in harsh environments but this makes disassembling them difficult, (2) have end caps that screw onto the enclosure but this limits how tightly they attach compared to an assembly that uses screws, (3) have holes to accommodate the external antenna but a new design could properly integrate an antenna into the package, (4) limit the form factors of available batteries to esoteric cells but a new design could make broader range of less expensive, high capacity cells available.

The choice of batteries reflected the requirement to find a cell small enough for presence nodes whereas extra room was available inside the climate nodes. Electrochems were rated at 1mA constant discharge. When motes were operating, they used nearly 25mA. Two large capacitors were added to prevent the in-rush of current when coming out of sleep and the resulting voltage drop from resetting the mote. The climate mote could be larger, however the desire for a common enclosure strategy limited the battery selection to a slightly taller and wider cells.

4.6.3 Node Reclamation

Reclamation is an important issue facing contemporary deployments. Because presence and climate motes are small, inconspicuous devices, they were easy to misplace and to lose in the field. Even with their GPS locations, motes deployed in early spring can be difficult to find in late fall after an entire summer's of vegetation growth. (For this reason, it is worthwhile augmenting locations experimental deployments with other markers, such as survey flags.) Whether localization is an intrinsic service of the sensor network or part of the installation process, geographic location is critical for the reclamation of nodes at the conclusion of a deployment.

Reclamation represents another application for field tools where for example, it could integrate GPS, directional mote antenna, and access to mote metadata like locations. With such a tool, a person one could zero in on a mote, pinging and listening for responses from different angles. The prospect of littering pristine and often protected habitats with nodes and lithium batteries may be unattractive, and in some cases, may be unacceptable, too. Strong economic incentives for reclaiming nodes will remain for the foreseeable future as long as climate motes continue to cost on the order of $400.00.

5 Related Work

There are two fundamental differences between traditional data loggers and presence and climate motes: loggers lack both networking and open source programming environments for application-specific customizations. One such data logger is the Hobo Data Logger [6]. Traditional data loggers can be larger and more expensive. They require that intrusive probes and corresponding equipment immediately adjacent. They are typically used since they are commercially available, supported, and provide a variety of sensors. 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" from the area of interest. With this method, biologists cannot gauge whether the weather station actually monitors a different micro-climate due to its distance from the organism being studied. Using the readings from the weather station, biologists make generalizations through coarse measurements and sparsely deployed weather stations. Instead, we strive to provide biologists the ability to monitor the environment on the scale of the organism, not on the scale of the biologist [2, 8].

Habitat monitoring for WSNs has been studied by a variety of other research groups. Cerpa et. al. [1] 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 [10], but no long term results or reliability data has been published. Wang et. al. [9] implement a method to acoustically identify animals using a hybrid iPaq and mote network.

ZebraNet [5] is a WSNfor monitoring and tracking wildlife. ZebraNet nodes are 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.

At UC James Reserve in the San Jacinto Mountains, the Extensible Sensing System (ESS) continuously monitors ambient micro-climate below and above ground, avian nest box interior micro-climate, and animal presence in 100+ locations within a 25 hectare study area. Individual nodes with up to 8 sensors are deployed along a transect, and in dense patches, crossing all the major ecosystems and environments found on the Reserve. The sensor data includes temperature, humidity, photosynthetically active radiation (PAR), and infrared (IR) thermopile for detecting animal proximity.

ESS is built on TinyDiffusionn [3, 7] routing substrate, running across the hierarchy of nodes. Micro nodes collect low bandwidth data, and perform simple processing. Macro sensors organize the patches, initiate tasking and process the sensor patch data further. They often perform functions of both cluster heads and patch gateways. In case of a macro sensor failure, the routing layer automatically associates macro sensors with the nearest available cluster-head. The entire system is time-synchronized, and uses SMAC for low power operation. Data and timestamps are normalized and forwarded to an Internet publish-and-subscribe middleware subsystem called Subject Server Bus (SSB), whereby data are multicast to a heterogeneous set of clients (e.g., Oracle, MatLab, and LabVIEW) for processing and analysis of both historical and live data streams. ESS makes an aggressive use of hierarchy within a patch; the diversity of sensors can also be used for verification of data. The SSB is a noteworthy departure from the architecture in Figure 1 – it allows for natural integration of triggered features into the system in addition to data analysis.

California redwoods are such large organisms that their life cycle can be measured through microclimate observations. Having developed models for their metabolism, biologists are now using sensor networks to verify and refine these models. The sensor network measures direct and incident photosynthetically active radiation (PAR), temperature, and relative humidity. In the fall of 2003, 70 nodes were deployed on a representative tree in the middle of the forest, reporting data every minute. Biologists intend to grow the network to both interior and edge trees in a grove.

The network collecting this information is an instantiation of the Tiny Application Sensor Kit (TASK) [4]. The macro sensors in the patch run a version of TinyDB query processing engine that propagates queries and collects results from a multi-hop network. There is no separate transit network – the patch bridges directly to the base station. The base station runs a TASK server that logs data, queries, network health statistics and keeps a journal of the experiment. TASK server is capable of running on a macro sensor. Deployment and in the field debugging are aided by a PDA-class device running a field tool, that allows for connectivity assessment and direct querying of individual sensors. To achieve low power operation, the entire network is time-synchronized and duty-cycled. TASK has a health query as one of the options, which could obtain voltage, routing, neighborhood and other networking state that could facilitate the analyses in Section 3

6 Conclusions

We have presented the system design and deployment of two wireless sensor networks. We deployed 150 nodes in both single- and multi-hop configurations. During four months of operation, these networks produced a rich dataset, valuable both to life and computer scientists.

Among the life scientists, Graphical Information Systems (GIS) have become the lingua franca for the visual presentation, analysis and exchange of geospatial data. Figure 13 shows a visualization of our study site of underground cavities on the left and corresponding surfaces on the right. Data for this visualization was collected from presence and climate motes, respectively, at midnight on a typical summer evening. Darker colors correspond to warmer temperatures, lighter colors correspond to cooler temperatures, and and shaded regions of similar colors correspond to isoplethic temperature regions. From this visualization, the effect of radiational cooling on the surface is apparent due to cooler temperatures, whereas the buffering and insulating properties of the cavities cause them to maintain their a constant temperature. Hot-spots in underground cavities are of special interest to life scientists, because the most likely source of heat would be a nesting sea bird.

The data allows computer scientists to compare the lab benchmarks with the phenomena observed in the real world over a period of months. In a number of cases (*e.g.*lifetime prediction for the single-hop network), the lab micro-benchmarks transfer well into a deployed application. In other cases, we see an interesting departure from the expected behavior, *e.g.*the lower packet yields from the multihop network. Often, these new behaviors emerge as the network is reconfigured, expanded, or redeployed; others are an artifact of aging. Most of those behaviors cannot be observed without actually constructing the complete system and deploying it in the field. Our analysis begins to extract some important characteristics, and compares them against the previously published results. Ultimately, this dataset should be used to inform models of real-world behavior; algorithms may be designed and tested in a lab with the characteristics of a real world outdoor deployment.

References

[1] CERPA, A., ELSON, J., ESTRIN, D., GIROD, L., HAMILTON, M., AND ZHAO, J. 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, Apr. 2001).

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

[3] HEIDEMANN, J., SILVA, F., AND ESTRIN, D. Matching data dissemination algorithms to application requirements. In *ACM SenSys 2003* (Los Angeles, CA, USA, Nov. 2003), ACM Press.

[4] HONG, W., MADDEN, S., FRANKLIN, M., AND HELLERSTEIN, J. TASK: Tiny Application Sensor Kit. `http://berkeley.intel-research.net/task/`, 2004.

[5] JUANG, P., OKI, H., WANG, Y., MARTONOSI, M., PEH, L.-S., AND RUBENSTEIN, D. 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)* (San Jose, CA, USA, Oct. 2002), ACM Press, pp. 96–107.

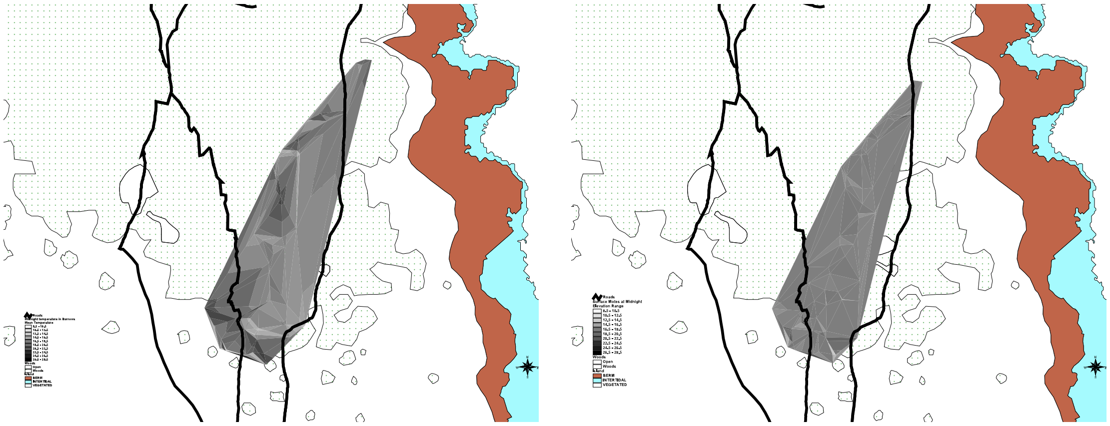

Figure 13: GIS map of midnight cavity (left) and surface(right) temperatures.

[6] ONSET COMPUTER CORPORATION. HOBO weather station. `http://www.onsetcomp.com`.

[7] OSTERWEIL, E., AND ESTRIN, D. TinyDiffusion in the Extensible Sensing System at the James Reserve. `http://www.cens.ucla.edu/~eoster/tinydiff/`, May 2003.

[8] TOAPANTA, M., FUNDERBURK, J., AND CHELLEMI, D. Development of Frankliniella species (Thysanoptera: Thripidae) in relation to microclimatic temperatures in vetch. *Journal of Entomological Science 36* (2001), 426–437.

[9] WANG, H., ELSON, J., GIROD, L., ESTRIN, D., AND YAO, K. Target classification and localization in habitat monitoring. In *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2003)* (Hong Kong, China, Apr. 2003).

[10] WANG, H., ESTRIN, D., AND GIROD, L. Preprocessing in a tiered sensor network for habitat monitoring. *EURASIP JASP Special Issue on Sensor Networks 2003*, 4 (Mar. 2003), 392–401.

[11] WOO, A., TONG, T., AND CULLER, D. Taming the underlying challenges of reliable multihop routing in sensor networks. In *ACM SenSys 2003* (Los Angeles, CA, USA, Nov. 2003), ACM Press, pp. 14–27.