ENVIRONMENTAL MONITORING THROUGH WIRELESS SENSOR NETWORKS

by

Tyler W. Davis

B.S., University of Pittsburgh, 2007

M.S., University of Pittsburgh, 2011

Submitted to the Graduate Faculty of the

Swanson School of Engineering in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2012

UNIVERSITY OF PITTSBURGH

SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

Tyler W. Davis

It was defended on

November 20, 2012

and approved by

Jorge Abad, Ph.D., Assistant Professor, Civil and Environmental Engineering

Piervincenzo Rizzo, Ph.D., Assistant Professor, Civil and Environmental Engineering

Stephen Tonsor, Ph.D., Associate Professor, Biological Sciences

Dissertation Director: Xu Liang, Ph.D., Associate Professor, Civil and Environmental Engineering

ENVIRONMENTAL MONITORING THROUGH WIRELESS SENSOR NETWORKS

Tyler W. Davis, Ph.D.

University of Pittsburgh, 2012

Wireless sensor networks (WSN) deployed for the purpose of environmental monitoring are becoming increasingly popular. This research examines the use of WSNs for the application of studying the hydrologic cycle with respect to the soil-plant-atmosphere continuum. The goal was to pair WSN technology with inexpensive hydrological sensors for the purpose of affordable and reliable environmental monitoring. This work encompasses the design, construction, and calibration of sap flow sensors; an examination into the power characteristics of the environmental sensors used to study hydrology and the wireless motes used to communicate data in the WSN; the results from deploying a pilot WSN test bed; the deployment, maintenance, and findings of a main WSN test bed; and the primarily results from the processed environmental sensor data. Lab experiments were used to determine an optimal sampling rate of environmental sensors, which was a compromise between the battery life of the motes and the temporal resolution of the environmental measurements. Based on the sampling rate, it was determined that the thermal dissipation sap flow method was the most practical for WSN applications. For the purposes of large-scale deployments, lab-made sap flow sensors provide a comparable cost-effective alternative to their commercial counterparts. Environmental monitoring of soil moisture, soil water potential, and sap flow was successful and exhibited spatial distributions and temporal changes within the main test bed. Analysis of the WSN communication revealed numerous factors impact network stability. While environmental monitoring through WSNs is feasible, its practicality is dependent on numerous variables.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

PREFACE

This research was supported by the National Science Foundation under the grants of CNS-0758372 and CNS-0721474 to Indiana University-Purdue University Indianapolis and the University of Pittsburgh, respectively.

This work was made possible by the collaborative efforts and support by the following people. First, I would like to acknowledge my dissertation advisor and research mentor, Dr. Xu Liang. Her support and encouragement throughout the past five years have made me a better researcher. Second, I would like to acknowledge my research supervisor and friend, Chen-Min Kuo. Working under his tutelage during the first two years of my graduate study were integral in laying the foundation for this dissertation. Third, I would like to thank the faculty who provided feedback and support for this dissertation including: Dr. Ravi Sharma, Dr. Rafael Quimpo, Dr. Jorge Abad, Dr. Piervincenzo Rizzo, and Dr. Stephen Tonsor. Lastly, I would like to thank my collaborators at IUPUI, namely Mr. Miguel Navarro, Mr. Newlyn Erratt, and Dr. Yao Liang. Their guidance with the wireless sensor network and web-based management system have contributed greatly to this work.

1.0 INTRODUCTION

1.1 MOTIVATION

The sustainable condition of freshwater resources partially depends on understanding the natural system in which it is cycled. Exploring the status and trends of soil moisture and transpiration at the local (i.e., plot) scale can help improve estimates of the water budget. Modeling the water budget of a region in an effort to reproduce the natural cycling can help forecast, and therefore help mitigate, natural hardships (e.g., floods and droughts). Currently, there is a partiality towards physically-based distributed hydrological models for reproducing regional-scale hydrology. However, these models require an understanding of the spatiotemporal variability of the variables studied which is difficult without an abundance of ground-based measurements for calibration and validation.

Ground-based measurements and monitoring of environmental variables has been impacted over the past decade by wireless sensor network (WSN) technology. Due to the expense of traditional data logging methods, researchers are forced to compromise between having high temporal, low spatial resolution (i.e., data logger measurements) or high spatial, low temporal resolution measurements (i.e., measurements taken manually). Recent technological advances have brought small, inexpensive, wireless monitoring devices to the researching field.

These devices are providing higher resolution data at a declining cost per unit area (Akyildiz et al., 2002).

As described by De Couto (2004), a WSN is a network of computers and devices (i.e., nodes) which are connected by wireless communication links. The links between nodes are supported by mote radios. Each radio link has a limited communication range and therefore many node pairs cannot communicate directly. As a result they must forward messages between one another by taking advantage of one or more intermediate nodes. The direction of messages through the network is based on the location of a specific sink node (i.e., the base station). Each link that an origin node's packet is transmitted over (i.e., parent node) is referred to as a hop. The set of links between an origin node and the base station is called a path. Transmission paths throughout a multi-hop network are referred to as a mesh. The study examined in this dissertation uses a multi-hop mesh network.

This dissertation presents a comprehensive examination of WSNs through the conception, design, and deployment of a WSN test bed being used for environmental monitoring. The novelty of this work can be seen in the size and longevity of the WSN deployment and the incorporation of sap flow monitoring, which is not a commonly studied environmental variable.

1.2 LITERATURE REVIEW

WSNs have found applications in a variety of environmental monitoring fields, including: habitat monitoring (Juang et al., 2002; Mainwaring et al., 2002; Szewczyk et al., 2004b), microclimate monitoring (Szewczyk et al., 2004a, Tolle et al., 2005; Hamilton et al., 2007), seismology (Werner-Allen et al., 2005, 2006), understory sunlight studies (Selavo et al., 2007),

agricultural studies (Panchard, 2008; Li et al., 2008; Martinelli et al., 2009), ecological studies (Musăloiu-E. et al., 2006; Rundel et al., 2009; Burgess et al., 2010), glacial studies (Martinez et al., 2004; Hart et al., 2006), and hydrological studies (Trubilowicz et al., 2009; Ingelrest et al., 2010; Kerkez et al., 2012). Each of these studies present deployments under unique environmental conditions. Because of the novelty of these applications, researchers appreciate the importance of understanding how this technology performs. In the majority of environmental studies using WSNs, researchers will spend as much time, if not more, studying the performance of the network collecting the data as they do studying the data being collected.

This dissertation looks to extend on the current research using WSNs to study the environment, specifically the spatial and temporal patterns of soil moisture, soil water potential, and tree transpiration in a forested watershed. As a part of this research, two WSN test beds were deployed: a pilot study and a main test bed. The main test bed is located at the Audubon Society of Western Pennsylvania (ASWP) and throughout this dissertation it will be referred to as the ASWP test bed.

WSNs are becoming a popular medium to monitor soil moisture because of the calibration needs for the increasing amount of satellite-based remote sensing data. Transpiration measurements, including sap flow, are not as well explored. As a part of this study, it was necessary to integrate environmental sensors into a wireless network and test the performance of the sensors in the network; both of which were successfully studied at the ASWP test bed. It was also necessary to determine if the WSN met the research needs of this project. In order to determine if this platform would be successful for environmental monitoring, WSN performance was explored. There will be a discussion in regards to whether using a WSN for real-time

environmental monitoring is a practical option considering the time, expense, and complications of maintaining the network.

The integration of sap flow monitoring into wireless applications has been attempted by both researchers and manufacturing companies. ICT International utilize a wireless modem connection, via an RS232 communication port on their SL5 Smart Logger, to provide remote access to the ICT data logger. Despite having remote access to network data, this system still relies on an expensive data logger for data collection.

In 2011, Dynamax Inc. introduced a wireless sap flow system called SapIP. They partnered with Digi International Inc. to create a multi-hop mesh platform for Dynamax sensors using the XBee-PRO DigiMesh 2.4 wireless module. Due to its recent release, no published research is currently available for studies utilizing this technology. It is also important to note that this technology, while it provides wireless mesh communications, is costly.

Starting in August 2006, Dr. Stephen Burgess, from the School of Plant Biology at the University of Western Australia, and Dr. Todd Dawson, from the Department of Integrative Biology at the University of California, Berkeley, conducted sap flow studies using wireless sensor networks at the University of California Berkeley Whitaker Forest property, located southwest of the Sequoia and Kings Canyon National Park, California. The TmoteSky wireless module and the ICT International HRM-30 sap flow sensors were used in this study. They successfully integrated the heat ratio method into a wireless networking system; however, it required extensive programming and software engineers to be realized (S.S.O. Burgess, personal communication, 30 April 2012). In addition to the sap flow sensors, their study also used meteorological sensors. Results from this study have been published (Burgess and Dawson, 2008; Burgess et al., 2010).

The field study examined in this dissertation is unique from Burgess and Dawson's research, in that the ASWP test bed studied uses the MICAz motes and Granier's thermal dissipation sap flow method (Granier, 1985, 1987b). The ASWP test bed study focuses on both the networking and the environmental data collected, which are two components that have yet to be extensively studied simultaneously. Also, sap flow is specifically studied in the context of the water cycle, as opposed to Burgess and Dawson, who looked at the biology behind sap flow. Due to budgetary restrictions, the ASWP test bed research also focuses on creating/using affordable technologies and techniques. For example, lab-made sap flow sensors, which are proven to be an acceptable alternative to the commercial grade sensors, are used at the ASWP test bed (Davis et al., 2012a).

In 2007, Phillip Rundel and Eric Graham, from the University of California, Los Angeles (UCLA) Center for Embedded Networked Sensing (CENS) together with Fabio Silva and Wei Ye, from the University of Southern California (USC) Information Sciences Institute (ISI) started conducting a wireless sap flow study at the Stunt Ranch Santa Monica Mountains Reserve. The study, called the Sensor Processing and Acquisition Network (SPAN), was conducted to examine the long-term influence of the 2006-2007 Southern California drought on local shrub and tree species (Ye et al., 2008; Rundel et al., 2009). No known sap flow sensor data has been published from this study.

Other studies of particular interest are the researchers looking at agricultural, ecological and hydrological monitoring with WSN systems. In 2005, Răzvan Musăloiu-E., Andreas Terzis, Katalin Szlavecz, Alex Szalay and Josh Cogan from the Johns Hopkins University and Jim Gray from Microsoft Research performed a proof-of-concept WSN deployment in an urban forest. In their study, Watermark soil moisture sensors and thermistors were attached to MICAz motes

using MTS101 data acquisition boards. The sampling rate was set to a 1-min interval. Temperature and soil moisture measurements for a six-week period within the 147 day test duration are published. In their results, they address the issues faced using the motes and discuss the power characteristics of the motes deployed (Musăloiu-E. et al., 2006).

This team used the MICAz motes for their field deployment, which is the same mote platform used in the ASWP test bed. This group studied the battery life of the MICAz motes during their deployment and found that the radio and processor operate below the recommended operational voltage. Similar results were found during the pilot test bed study. The transmission success of status messages they received from the field is similar to that of the ASWP test bed.

In 2005, Ecole Polytechnique Federale De Lausanne (EPFL), which is one of two federal institutes of technology located in Lausanne, Switzerland, in partnership with the Center for Electronic Design and Technologies (CEDT) of the Indian Institute of Science, explored the potential of WSNs for agriculture in developing countries. The title of the project was COMMON-Sense Net (CSN). The research for CSN was conducted in farming regions of India and the group looked at practical applications for WSNs to be operated by nontechnical users. Their study also compared power consumption and transmission distance for the MICA2 mote by Crossbow and the Tinynode mote by Shockfish. They ultimately decided to use the Tinynode mote. They have published results for the WSN performance (Panchard, 2008).

The ASWP test bed research is similar in that postgreSQL is used for database management, which CSN also utilize in their study. CSN had a web management tool, which allowed them to send commands to the WSN and view data and network statistics. The ASWP test bed also has a web management tool that provides similar capabilities.

Joel Trubilowicz, from the Department of Forest Resource Management, and Kan Cai, from the Department of Computer Science, both at the University of British Columbia and Markus Weiler, from the Institut fur Hydrologie, Universitat Freiburg, conducted studies from July 2006 to April 2007. Their field study was comprised of a 41-node deployment in the forested catchment in the Malcolm Knapp Research Forest, located in Maple Ridge, BC, Canada. They measured the air temperature, humidity, soil temperature, rainfall intensity, soil moisture, groundwater head, and overland flow, to see if the technology worked for the application of hydrological study. Technology used includes: MICA2 wireless motes and MDA300 data acquisition boards. Published results are available for this study (Trubilowicz et al., 2009).

While this study found that the network was not consistently reliable and mote configuration was timely and posed difficulties, it was still a successful short-term deployment in a forested area. While the duration of the Trubilowicz WSN was short-term, similar issues were found in the long-term ASWP test bed. Similar issues faced include consistency and reliability of the network as well as mote configuration challenges. In the case of the ASWP test bed, mote configuration issues were resolved with the assistance of researchers from the Institutes of Energy and the Environment at Penn State University. Trubilowicz also found that purchasing new motes, as the technology continues to advance, is not realistic for most researchers due to cost considerations. Investing heavily for long-term studies does not allow for purchasing mote platforms when new products come on the market, as was the case with the ASWP test bed.

In 2006, two research groups from Switzerland (EPFL: LCAV and EFLUM) collaborated on a WSN. EPFL: LCAV performed signal processing and networking and EFLUM contributed hydrology and environmental fluid mechanics research. The project was called SensorScope and the field work was conducted throughout the Alpine region of Switzerland. The aim of the

project was to help environmental engineers address long-term monitoring questions in challenging environments (Ingelrest et al., 2010). The technology used in this study includes the Shockfish Tinynode wireless module and sensing stations that contain a suite of environmental and hydrologic sensors. Published results are available for this study (Ingelrest et al., 2010).

EPFL: LCAV and EFLUM note important considerations for outdoor deployments, which were useful for the ASWP test bed deployment.

Brando Kerkez and Steven Glaser, from the Department of Civil and Environmental Engineering at the University of California, Berkeley together with Roger Bales and Matthew Meadows, from the Sierra Nevada Research Institute, University of California, analyzed the performance of a 57-node WSN, in which they monitored water balance across a 1 km^2 forested catchment in the southern Sierra Nevada of California. They used wireless devices developed by Dust Networks. Published results are available for WSN performance but not for the environmental data (Kerkez et al., 2012).

This group proposes a deployment strategy that ensures reliable performance and reduces network setup time. This team uses a 15-min sampling interval for their wireless measurements, which is the same sampling interval used at the ASWP test bed. They propose the use of two metrics, the packet delivery ratio (PDR) and received signal strength indicator (RSSI), for evaluating the operational path behavior of the WSN. In Section 5.2.3.2 of this dissertation, the RSSI is used for evaluating the operational WSN paths of data through the ASWP network.

1.3 SCOPE

The decision to create a test bed was initiated in 2007 when the applications of WSNs were becoming popular. In hydrology, there is an ever growing need for high quality ground-based measurements for validating and calibrating hydrological models. An investigation on using this new technology to meet this need seemed prudent. Technological understanding of WSNs (in terms of its functionality, performance, and usability) was limited outside its main field of research. Therefore, a broad analysis of using WSNs for environmental monitoring was necessary.

To begin with, a suite of environmental sensors were decided upon. Soil moisture is one of the most important factors in studying hydrology; therefore soil moisture sensors were used in the ASWP test bed. For a better understanding of soil moisture, accompanying soil water potential sensors were used. Finally, to investigate water's interaction between the soil and the atmosphere (primarily taking place within vegetation in a forested region), the last sensor type was chosen. For monitoring water movement through vegetation (i.e., sap flow), there are two thermometric sensor styles. Both of these sensors are investigated to determine which is more appropriate for WSN applications.

This leads to the first topic investigated in this dissertation. The designs of two types of sensors for monitoring sap flow are discussed. The sensors were built in the lab for the purpose of making sap flow monitoring affordable for large-scale deployments. The current expense of sap flow sensors limits this type of monitoring. To accompany the two designs, methods for quality control are presented. Included in this analysis is a comparison of the performance of the two lab-made sap flow sensor designs to a commercial sap flow sensor.

Sap flow sensors were not initially compatible with WSN hardware. The sensor probes did not produce a strong enough signal for the data acquisition boards to detect. To integrate the sap flow sensors into the WSN, a signal amplifier was necessary. To accommodate this need a circuit board was designed to allow for sap flow monitoring to be used in the WSN. The two sap flow sensor designs and their construction are investigated, in addition to a study of quality control, within Chapter 2.0

When using WSNs for environmental monitoring there is a tradeoff between the temporal resolution of measurements (i.e., the sampling rate) and the battery life of the motes in the network. Note that the motes refer to the individual units responsible for the data collection and transmission throughout the WSN. Since motes sample sensors and transmit data a power source is required, which is why battery life is a factor in understanding WSN performance. The WSN platform used at the ASWP test bed does not discern between sampling and transmission rates. Therefore, sampling rates are a concern because each sample requires a transmission and the number of transmissions is the main consumer of a mote's battery power. An investigation of various sampling rates, for the sap flow sensors used in this study, was conducted to determine a compromise between capturing natural temporal patterns and maximizing battery life. Included in this investigation is the impact of a mote's battery power on the three environmental sensor measurements. Results of these analyses are presented in Chapter 3.0

Once the environmental sensors were chosen and tested the next phase was to take the sensors and the WSN technology and investigate their performance in a pilot test bed. Approximately ten nodes were deployed on a residential parcel of land in the northern suburbs of Pittsburgh. Soil moisture and soil water potential sensors were added to the network. Sap flow was not successfully integrated into the pilot study. The main focus of this study was on network

performance, rather than environmental sensor performance. Specifically, power efficiency and battery savings were examined. Performance of the network in the field was compared to laboratory experiments. Included in this investigation is a description of a web-based WSN management system. The findings from the pilot test bed, including a discussion on the characteristics of the transmission of messages through the network, are presented in Chapter 4.0

Following the success of the pilot test bed the next phase was to deploy the main test bed (i.e., the ASWP test bed) within the Pittsburgh region. A relatively small catchment consisting of mostly undeveloped land was selected for the study area. Within the catchment, the test bed was located at the Audubon Society of Western Pennsylvania's (ASWP) Beechwood Farms Nature Reserve (BFNR) in Fox Chapel, Allegheny County, Pennsylvania. Utilizing the lessons learned from the pilot study, the deployment was broken down into three phases starting in the spring of 2010. Within three months, a 40-node network was in place. The network consisted of 25 sensor nodes and 15 relays. Within those 25 sensor nodes there were 55 soil moisture sensors, 25 soil water potential sensors, and 22 sap flow sensors.

The network was analyzed over a two-year study period. Over this two-year time period the number of soil moisture and soil water potential sensors remained the same; however the amount of sap flow sensors deployed at the ASWP test bed varied season to season. A maximum number of 22 sap flow sensors were installed in any given season.

Chapter 5.0 presents an evaluation of the WSN in terms of cost and complications. First, a thorough description of the ASWP test bed is given. Included in this description is a discussion of the base station and gateway remote connectivity, node hardware and enclosures, routing protocol of the wireless motes, and timeline for the network deployment. The network is then evaluated in terms of financial considerations and maintenance time. This is followed by an

analysis of the communication of the motes within the network using health statistics data and transmission path simulations. Complications of the network are also examined in the context of gateway connectivity, mote communication, physical damage to equipment, and battery life.

The analysis of the environmental measurements is presented in Chapter 6.0 It is important to note that the scope of this dissertation does not include an in-depth analysis of the hydrological data collected. The majority of the research conducted was done in an effort to establish a data collection system that could be used for environmental monitoring in the future. The realization of this goal was touched on in this chapter, but it was not the main focus of this work.

Within Chapter 6.0 the analysis of environmental measurements includes an examination of the seasonal spatial distribution of soil moisture and soil water potential, a comparison of WSN soil sensor measurements to data logger measurements, and an examination of the impact of duplicate packets on WSN soil sensor measurements. Additional analysis includes the temporal variations in sap flow data and an examination of noise in sap flow sensor measurements. Within the discussion, methods for compensating sap flow measurement noise, the monitoring cost per unit area of the network, and natural phenomenon (i.e., a lightning strike that impacted equipment within part of the ASWP test bed) are also examined.

2.0 SAP FLOW SENSORS: CONSTRUCTION, QUALITY CONTROL AND COMPARISON

The majority of the content contained in this chapter was previously published in the open access journal, *Sensors*, under the title, "Sap Flow Sensors: Construction, Quality Control and Comparison." References to the published paper are cited as Davis et al (2012a).

The heat ratio (Burgess et al., 2001), heat balance (Sakurtani, 1981), and thermal dissipation (Granier, 1985) sap flow methods make up the three main sensor types that are used today for estimating plant transpiration. With respect to certain tree sizes (>25 cm diameter), only the heat ratio method (HRM) and the thermal dissipation method (TDM) are applicable (Swanson, 1994). Both of these methods use cylindrical thermocouple and heater probes inserted into the tree xylem. While the installation and methodology of using these sensors are different, the fundamental mechanics and operation of the probes are similar. Swanson (1994) suggests that a combination of both pulsed and thermal dissipation sap flow measurement methods would be the most cost effective solution for long-term studies. In addition, the high cost associated with commercial grade sensors of the same types (i.e., TDM and HRM) diminishes the ability for researchers to make large-scale deployments. Therefore, this chapter describes the design of both a HRM and TDM sap flow sensor utilizing inexpensive parts for the purpose of estimating tree sap flow.

To accompany the design of these two sensors, a quality control checking method is also presented. Due to the inconsistencies that may arise when producing lab-made sap flow sensors, it is important to calibrate each probe's output to a standard measure. This chapter presents a simple method to do so and it also presents the results from calibrating 24 temperature and heater probes' output (measured in voltage) to a standardized temperature (measured in degrees). This accounts for the variations in individual sensors by applying a linear conversion equation to the sensor probes' output.

To demonstrate the sensor's performance, results from a field study using both HRM and TDM handmade sensors are also presented. For further comparison, a commercial sap flow sensor, Dynamax Inc.'s TDP30, was field tested alongside the handmade sensors. Five consecutive days' results from all three sensors are presented. Simultaneous measurements from three different sensor types measuring at three different locations in the same tree cannot present the same quantitative results due to heterogeneities present in the thermal properties of the tree's sap-conducting wood (Burgess and Dawson, 2008). Therefore, only a qualitative comparison is made.

2.1 BACKGROUND AND MOTIVATION

2.1.1 Overview of sap flow methodologies

Sap flow is an important environmental variable in hydrology for estimating the amount of water transpired by vegetation. Sap flow is the water that travels inside the sapwood of a tree (or other woody plants), typically taken up from the soil by the roots, travels through the tree trunk and is

released to the atmosphere by the leaves. The quantification of sap flow is made by measuring the velocity of the water being transported through the active sapwood. There are two thermometric methodologies which aim to measure sap flow.

The TDM was first introduced by Granier (1985). This method is based on the assumption that the heat input by the sensor under steady sap flow conditions is equal to the heat dissipation (via convection and conduction) along the interface between the sensor and the tree when the sensor and the tree are in thermal equilibrium. Daily fluctuations in the heat dissipated from the sensor probe are compared to the unheated temperature of the tree sap and wood. To measure the heated and reference temperatures, two probes, vertically aligned, are inserted radially into the sapwood of a tree. The downstream probe consists of a coiled metal wire, which supplies the heat via the Joule effect, and a thermo-junction, which measures the temperature via the Seebeck effect. A constant voltage is added to the coiled metal wire to power the heater while the thermo-junction measures the temperature of the heat source. The upstream probe consists of only a thermo-junction for measuring the unheated reference temperature at a fixed distance away from the heater probe. Daily temperature differences are used to calculate the sap flux density from an empirical equation developed by Granier (1987a) and terminology based on Edwards et al. (1996):

Equation 2.1

$$Q_s = 0.000119 \cdot \left(\frac{\Delta T_0 - \Delta T}{\Delta T} \right)^{1.231}$$

where: ΔT = Temperature difference between the upper and lower probes (°C)

ΔT_o = Maximum daily value of ΔT, i.e. zero sap flow (°C)

Q_s = Sap flux density ($m^3 \cdot m^{-2} \cdot s^{-1}$)

Lu et al. (2004) outlines the various research papers that have reported calibrating the Granier-style sensor on a variety of tree species and porous media.

The HRM is an improved heat-pulse method developed by Burgess et al. (2001) in response to the limitations of the compensation heat-pulse method (CHPM) first introduced by Huber and Schmidt (1937) and later developed by Marshall (1958) and Swanson (1962). This method uses the relative temperature increases following a heat pulse measured at equal distances upstream and downstream from the source to determine the convective velocity of the heat pulse which can then be compensated by the tree's physical characteristics (i.e., a weighted average of the stationary wood and flowing sap) to the sap velocity. In this method, three probes are inserted radially into the tree sapwood. The middle probe generates a heat pulse and the temperature changes are measured by the other two probes at locations equidistant upstream and downstream from the heater. The heat pulse velocity is calculated according to Marshall (1958) and terminology according to Edwards et al. (1996):

Equation 2.2

$$v_h = \frac{k}{x} \cdot \ln\left(\frac{v_1}{v_2}\right)$$

where: v_h = Heat pulse velocity ($mm \cdot s^{-1}$)

k = Thermal diffusivity of sapwood ($mm^2 \cdot s^{-1}$)

x = Distance between the heat pulse and temperature probes (mm)

v_1/v_2 = Ratio of time-dependent up/down stream temperature differences where $v(t) = T(t) - T(0)$

The times, t, during which temperature measurements are made should be between 60–100 s following the generation of the heat pulse when v_1/v_2 is effectively linear with respect to time

(Burgess et al., 2001). The sap velocity is then calculated based on the methodology of Barrett et al. (1995) and terminology according to Edwards et al. (1996):

Equation 2.3

$$v_s = \frac{\rho_{sm} \cdot c_{sm}}{\rho_s \cdot c_s} \cdot v_h$$

where: v_s = Sap velocity ($mm \cdot s^{-1}$)

ρ_{sm} = Density of sap in a woody matrix ($kg \cdot m^{-3}$)

c_{sm} = Specific heat of sap in a woody matrix ($J \cdot kg^{-1} \cdot °C^{-1}$)

ρ_s = Density of sap ($kg \cdot m^{-3}$)

c_s = Specific heat of sap ($J \cdot kg^{-1} \cdot °C^{-1}$)

The heat pulse velocity in Equation 2.2 can be corrected for sensor misalignment and tree wounding effects which can affect the sap flow measurements (Burgess et al., 2001). Results of either Equation 2.1 or Equation 2.3 can be multiplied by the cross-sectional area of conducting sapwood to convert sap flux density (Q_s) or sap velocity (v_s) to sap flow (Q).

2.1.2 Existing measurement technologies

The means for measuring sap flow have been designed and sensors for the HRM and TDM are commercially available. Lu et al. (2004) lists three major companies as distributers of the Granier-style sap flow sensors, including UP Gmbh (Germany), PlantSensors (Australia), and Dynamax (USA). Others include Ecomatik (Germany) and ICT International (Australia). ICT International is the only known supplier of a sensor specifically for the HRM. These sensors,

however, are expensive and often times packaged with proprietary data collection equipment and software which can also be expensive and limited in functionality.

The fundamental principles of these sap flow sensors are based on the thermo-electric and Joule heating laws. Using these principles, temperature and heater probes can be created for the HRM and TDM similar to those described in their respective works or those sold by commercial sap flow sensor retailers. In their design of a variable length TDM sensor, James et al. (2002) explains that using lab-made sap flow sensors is inexpensive and they are simple to manufacture. It is the intention of this work to exemplify this with straightforward sensor designs, quality control measures, and an example application comparing the lab-made sensors to their commercial counterparts.

Lab-made sap flow sensors have been presented in numerous other works, e.g., James et al. (2002); Braun and Schmid (1999); Clearwater et al. (1999); Pearsall et al. (2011), but most fail in describing at least one of the following: the process in which the sensors were made, the materials used, or the quality assurance of each sensor. This chapter details each of these key points.

2.1.3 Benefits of using lab-made sap flow sensors

There are many benefits to making sap flow sensors rather than purchasing them commercially. These benefits are mainly due to the selection and pricing of materials used for the sensor building.

The parts for making the probes comprised in the sensor design for both the HRM and TDM sap flow sensors are the same, only in different quantities. Therefore, researchers can take advantage of using both sensor designs in their study which has been suggested for improving

the quality of measurements (Swanson, 1994). Purchasing the sensor parts also allows for custom fitting the sensors to the specific application. This can include the probe length (to account for various sapwood depths), probe spacing (especially for the TDM where Lu et al. (2004) show that the distance between the heater and reference temperature probes may affect results), and wire connection type (to work with various data logger or collection systems).

The price of lab-made sensors, when considering both time and materials, is much less than their commercial counterparts for large deployments. For small deployments, however, the trade-off between the quality and price between lab-made and commercial sensors may not be as justifiable. However, large deployments allow researchers to install a greater number of sensors over larger areas and species types. This can help reduce errors in transpiration estimation made by single point measurements. It is well known that the TDM is prone to errors due to variations of sap flow along the sensor length, improper probe placement, disruption to conducting cells during installation, improper probe spacing, and temperature gradients within the sap wood (Wullschleger et al., 2011). To help accommodate for this, it is recommended that multiple sensor types be deployed together. Because both the TDM and HRM sap flow sensors share the same parts, they can be built and deployed together for a dual perspective of sap flow.

Time and effort are necessary to assemble sap flow sensors. The experience of the worker will affect the quality of the sap flow sensors produced. The time spent on sensor production may delay deployment, however, the experience gained from making sensors is valuable.

2.2 METHODS AND MATERIALS

This section presents two sensor designs for the TDM and HRM sap flow calculation methods. The designs for building the sap flow sensors are based on Chen and Miller (2012a, 2012b). The principal components of these two sensors have been reduced to three probe designs. From these three probes, both the TDM and HRM sap flow sensors can be constructed. The TDM and HRM sap flow sensor designs both utilize a temperature probe (for the reference temperature in the TDM and the equidistant temperature measurements upstream and downstream in the HRM). The difference between these two sap flow sensors are their heater probes. The TDM utilizes a constant heater probe, consisting of both a heater and a temperature measurement, while the HRM utilizes a heat pulse probe, consisting of only a heater.

To supply the constant heat required by the TDM, a design for a voltage regulator is necessary and is included in this chapter. Unlike the HRM, which uses a short but high power heat pulse, the TDM delivers a steady and constant power output which in this design is controlled via a voltage regulator.

To control the quality of these probes, their measurements (measured in mV) are calibrated to a standardized thermocouple (measured in °C). The testing procedure and results from two tests are presented below.

2.2.1 Construction and quality control

2.2.1.1 Temperature probe

The temperature probe used by both the HRM and TDM consists of a thermocouple junction, or thermo-junction, inside of a metal needle. Traditionally, a copper-constantan (Type

T) thermocouple is used. In this design, the thermocouple type was changed to a chromium-constantan (Type E) to increase the voltage output. Figure 1 shows the comparison between Type T and Type E thermocouples. The thermocouples' voltage response to a given temperature can be represented by a polynomial (of order seven and nine for Type T and E thermocouples, respectively) and have a nearly linear response curve in the temperature range expected for sap flow monitoring conditions. However, the Type E thermocouples produce over 50% higher voltage response than Type T for a given temperature (e.g., 1.4 mV compared to 0.9 mV at 23 °C for Type E and T thermocouples, respectively). Lu (1997) shows that the TDM sap flux density can be directly calculated from the voltage measurements as opposed to the temperature, due to the cancellation of the conversion factor (Seebeck coefficient) in the numerator and denominator. Based on the ratio of time-dependent temperature differences in Equation 2.1, the same is also true for the HRM calculation. Therefore using a thermocouple with a higher Seebeck coefficient will not change the results in the calculation. Instead, for the same temperature measurement the Type E thermocouple will produce a higher voltage that can be more easily resolved by instrumentation.

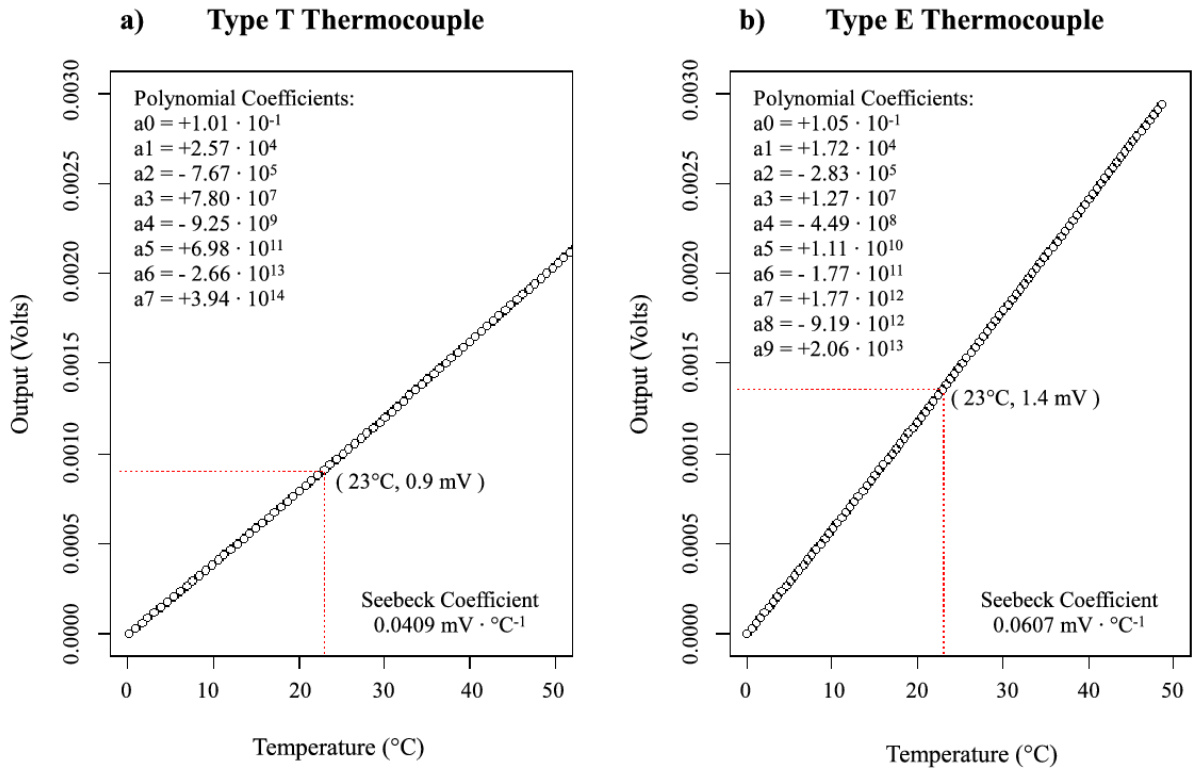

Figure 1. (a) Type T thermocouple voltage to temperature conversion plot. (b) Type E thermocouple voltage to temperature conversion plot. The polynomial voltage (V) to temperature (°C) conversion equation coefficients, Seebeck coefficients for the near linear temperature range shown and voltage responses at 23 °C are given for both thermocouple types. (Davis et al., 2012a).

The construction of the temperature probe begins with 36 gauge (0.127 mm diameter) chromel and constantan thermocouple wire. The wires are soldered to form a thermo-junction and inserted into a 41 mm long, 0.556 mm inner-diameter glass micropipette. The chromel wire is insulated with a PFA copolymer resin to prohibit short-circuits. The thermo-junction is glued into place inside the micropipette using a cyanoacrylate adhesive. The location of the thermo-junction in the micropipette determines the depth temperature measurements will be taken inside the tree. For this design, the thermo-junction is located at half the length of the micropipette. At this point, the thermocouple inside the micropipette can be tested using a multimeter to measure the resistance between the exposed leads to ensure the soldered joint was not damaged during installation. Depending on the quality of the soldered joint, resistance measurements typically are between 1–4 Ω. If resistance measurements are found to vary greatly from the expected, the thermocouple is discarded. Type E extension wires, 20 gauge (0.8128 mm diameter), are wrapped and soldered to their respective thin leads of thermocouple wire (Figure 2a). The connections are insulated with polytetrafluoroethylene (PTFE) tape. The micropipette is carefully inserted into an 18 gauge (0.965 mm inner-diameter) 38.1 mm long stainless steel dispensing needle (Figure 2b). It should be noted that the needle size was chosen to closely match that of the Dynamax TDP30 commercial sap flow sensor. The micropipette and needle size may be changed to match other designs. The tip of the steel needle is sealed with solder to restrict water from reaching the wiring. To hold the assembly together, heat shrink tubing is tightened around the needle hub and thermocouple extension wire. For additional water protection, the seam around the heat shrink tubing is glued (Figure 2c).

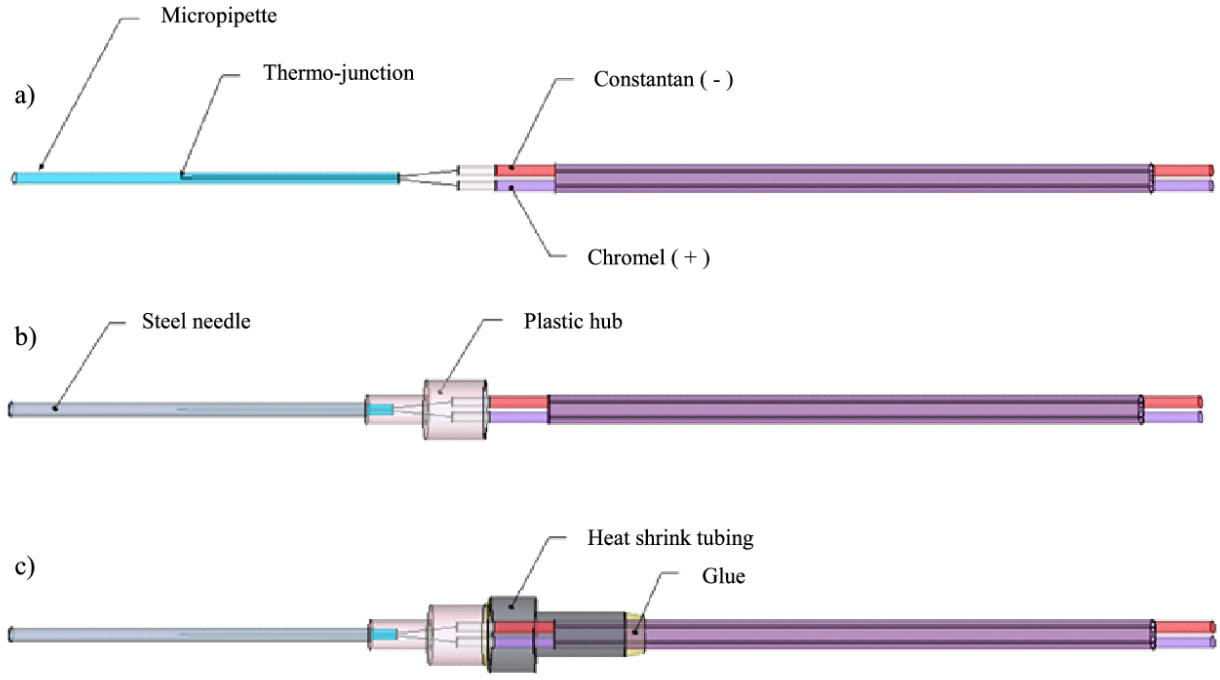

Figure 2. (a) Temperature probe schematic of thermo-junction location inside the micropipette and connection to Type E thermocouple extension wire (constantan/chromel). (b) Temperature probe schematic showing micropipette located inside the steel dispensing needle. (c) Temperature probe schematic of the heat shrink tubing and glue used to secure the thermocouple extension wire to the plastic hub of the needle. (Davis et al., 2012a)

2.2.1.2 Constant heat probe

The constant heat probe used by the TDM is a modification of the temperature probe described in Section 2.2.1.1. To provide the constant heat source required by the TDM, a heater wire is wrapped around the outside of the temperature probe's needle. To reduce the potential for short circuiting the heater wire, the metal needle is first insulated with PTFE tape (see Figure 3a). Nickel-chromium wire (also known as Nichrome 60 and Chromel C) is then wrapped around the insulated needle (see Figure 3b). Note that the wire coil is not drawn to scale. At this point, the heater wire can be tested using a multimeter to measure the resistance between the two leads. If the resistance measured varies greatly from the expected, the wire is discarded. Positive and negative leads are attached to either end of the heater wire and are connected to a voltage regulator. For additional security against wire breakage, the heater leads can be tied to the thermocouple extension wire with electrical tape or by other means.

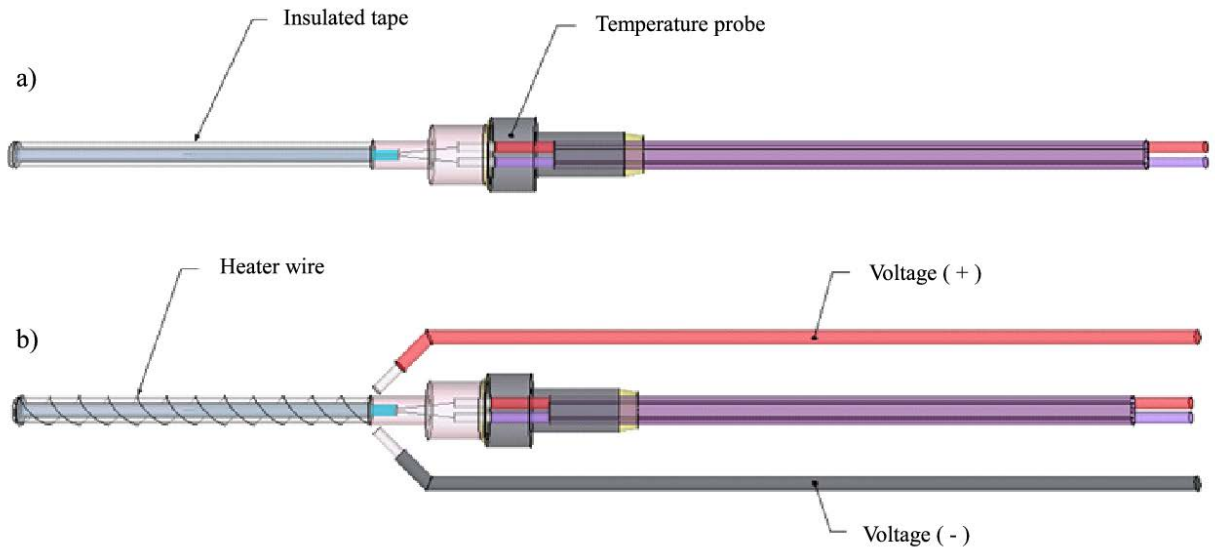

Figure 3. (a) TDM heater probe schematic of insulated tape covering the temperature probe's steel dispensing needle. (b) TDM heater probe schematic of heater wire coiled around the insulated tape and connection to the positive and negative voltage extension wire. (Davis et al., 2012a)

The 36 gauge (0.127 mm diameter) nickel-chromium wire has an approximate resistance of 88.6 $\Omega \cdot m^{-1}$. Granier (1985) calibrated the heater at 10 Ω and used a current of 0.141 A resulting in a 0.2 W powered probe. It is important to note that the calibration of this sensor depends on the heat field created by the probe (characterized by its size and shape) and the heating power used (Lu et al., 2004). For this design, the heater resistance is set to approximately 45 Ω; therefore to maintain the 0.2 W power requirement of this sensor design, a 3.0 V power supply is required. It should be noted that the heater resistance and power voltage were decided based on the comparison with the Dynamax TDP30 sap flow sensor, which is provided later in this chapter. It is assumed that the heater resistance and power voltage can be adjusted to various combinations so long as the output power of 0.2 W is maintained.

The PTFE tape provides some support in holding the heater wire in place during installation. In Granier's original design, the heater was placed into an aluminum sheath. The extent in which the aluminum sheath improves the thermal contact between the heater and the wood is unknown. Most installations today use wax or grease to improve thermal contact between the heater and wood in lieu of the aluminum sheath.

2.2.1.3 Heat pulse probe

The HRM heat pulse probe, unlike that used by the TDM, does not measure temperature. Therefore the wire heater can be placed inside the probe needle where the thermo-junction is located, as seen in the temperature probe from Section 2.2.1.1. To reduce the potential for short circuiting the heater, the wire coil is placed inside a micropipette which is then inserted into the steel needle of the probe. The positive end of the heater wire is attached to heater extension wire, while the negative end of the heater wire is either also attached to the extension wire or soldered directly to the steel needle. In the latter case, negative extension wire is also soldered to the metal needle. The steel needle tip is sealed with solder to hold the negative end of the heater wire in place and prevent water intrusion. For quality assurance, the resistance of the heater probe between the positive and negative leads can be measured using a multimeter. If the resistance varies greatly from the expected, the heater probe is discarded.

The heater (Figure 4) is made of 36 gauge nickel-chromium (Nichrome 60/Chromel C) wire. To increase the temperature enough for detection in the upper and lower temperature probes during the short heat pulse, the resistance in this design was set at 20.1 Ω. Using a 12 V battery source, the probe delivers over 7.0 W of power.

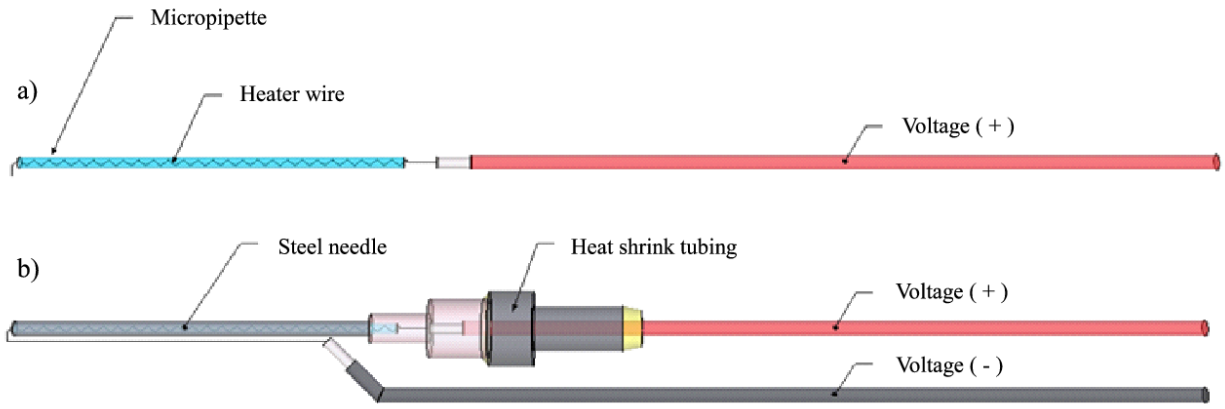

Figure 4. (a) HRM heater probe schematic of heater wire coil inside the micropipette connected to the positive voltage extension wire. (b) HRM heater probe schematic of micropipette located inside the steel dispensing needle, connection of the negative voltage extension wire to the heater wire, and heat shrink tubing and glue securing the positive voltage extension wire to the needle's plastic hub. (Davis et al., 2012a)

2.2.1.4 Voltage regulator

The TDM requires a steady constant power supply. Based on the TDM sensor design, the voltage that is required to maintain 0.2 W of power to the heater probe is approximately 3 V. To achieve this, a voltage regulator was designed which can reduce and split a single 12 V power source into two 1.25–11 V variable outputs. Each regulator requires only five components: an LM317 voltage regulator, a 240 Ω resistor, a 0–2.5 kΩ potentiometer, and two 0.1 µF capacitors. The dual design, given here, splits the 12 V input power in parallel between two regulators that are located on the same circuit board. This allows regulated voltage output to two separate constant heat probes. This design may be reduced to a single voltage regulator if desired. Figure 5 shows the layout of the components and a general wiring scheme.

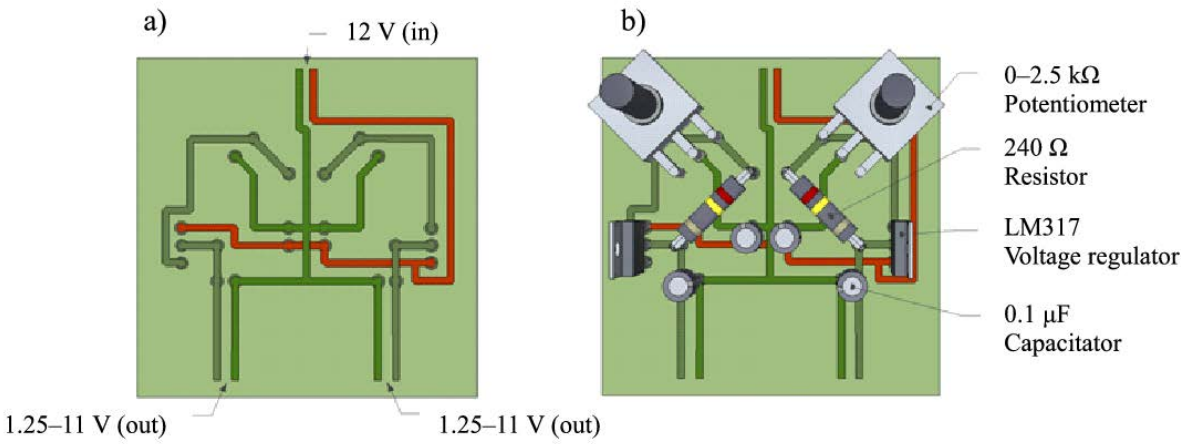

Figure 5. (a) Voltage dual regulator general wiring schematic. (b) Voltage dual regulator component locations and wiring schematic. (Davis et al., 2012a)

2.2.1.5 Sensor output calibration

Both the temperature probes and the TDM heater probes were tested for their response and accuracy to temperature measurements. This was to make certain that errors were not introduced into the probe temperature measurements due to variances in the craftsmanship. The testing consisted of comparing the temperature readings made by the sensor probes to a standardized thermo-junction. Each pair of measurements was recorded simultaneously in a series of three different water temperatures. The three water temperatures were: hot (~40 °C), warm (~25 °C), and cold (~10 °C). A testing pattern of *cold* → *warm* → *hot* → *cold* → *hot* → *warm* was used such that each sensor probe was tested at each temperature twice.

Individual tests took 3 min (30 s for each temperature) and captured a gradual heating response, a rapid cooling response, a rapid heating response, and a gradual cooling response. The water temperature was measured and recorded every 1 s using an Omega thermometer and temperature monitoring program (180 data points).

The TDM heater probes were also tested against the standardized thermocouple. Due to the insulation around the probe needle for the heater wire, each of the temperatures were tested for twice as long, resulting in a 6 min total test time for each heater probe (360 data points).

2.2.2 Parts and pricing

Table 1 shows the quantities for the various parts comprising the sap flow sensor probes and voltage regulator. The prices in Table 2 were taken from five commercial wholesale vendors for the year 2011. Table 2 shows the vendor and part numbers for all the materials listed in Table 1.

Table 1. Material quantities for a temperature probe, constant heat probe, heat pulse probe and voltage regulator. (Davis et al., 2012a)

Item	Temp. Probe	Const. Heat Probe	Heat Pulse Probe	Voltage Regulator
Pipette (10 μL)	1 ea.	1 ea.	1 ea.	—
Heat shrink tubing	2 cm	2 cm	2 cm	—
Stainless steel needle	1 ea.	1 ea.	1 ea.	—
Cyanoacrylate	0.015 mL	0.015 mL	0.005 mL	—
PTFE tape	5 cm	15 cm	5 cm	—
Chromel TC wire	5.5 cm	5.5 cm	—	—
Constantan TC wire	5.5 cm	5.5 cm	—	—
Type E TC extend wire	10 cm	10 cm	—	—
NiCr heater wire	—	50.8 cm	22.7 cm	—
Heater extend wire	—	12 cm	12 cm	—
Electrical tape	—	5 cm	5 cm	—
Resistor (240 Ω)	—	—	—	2 ea.
Capacitor (0.1 μF)	—	—	—	4 ea.
Voltage reg. (LM317)	—	—	—	2 ea.
Potentiometer	—	—	—	2 ea.
Dual pc board	—	—	—	0.5 ea.
Extend wire	—	—	—	20 cm

Table 2. Commercial vendor 2011 part numbers and pricing for sap flow probe and voltage regulator materials. (Davis et al., 2012a)

Item	Vendor	Part No.	Price/unit
Pipette (10 μL)	Cole-Parmer	EW-07950-20	$13.50/100 pc.
Heat shrink tubing	DigiKey	A014B-4-ND	$1.69/1.22 m
Capacitor (0.1 μF)	DigiKey	P984-ND	$0.26 ea.
Resistor (240 Ω)	DigiKey	240QBK-ND	$0.07 ea.
Voltage reg. (LM317)	DigiKey	LM317TFS-ND	$0.68 ea.
Potentiometer	DigiKey	CT2263-ND	$1.55 ea.
Stainless steel needle	McMaster-Carr	75165A75A	$14.41/50 pc.
NiCr heater wire	McMaster-Carr	8880K85	$25.01/487.68 m
Heater extend wire	McMaster-Carr	9697T1	$35.00/76.20 m
Extend wire	McMaster-Carr	7587K911	$6.56/30.48 m
Electrical tape	McMaster-Carr	76455A22	$4.33/20.12 m
Cyanoacrylate	McMaster-Carr	74555A44	$4.92/0.1 oz.
PTFE tape	McMaster-Carr	4591K12	$1.74/15.24 m
Chromel TC wire	Omega	TFCH-005-50	$16.00/15.24 m
Constantan TC wire	Omega	TFCC-005-50	$16.00/15.24 m
Type E TC extend wire	Omega	EXTT-E-20-50	$72.00/15.24 m
Dual pc board	RadioShack	276-148	$1.99 ea.

2.3 RESULTS

Based on the three probe designs, the HRM and TDM sap flow sensors can be constructed in a lab environment. From the quantities given in Table 1 and the bulk pricings shown in Table 2, the temperature probe costs approximately $1.07, the constant heat probe costs approximately $1.17, and the heat pulse probe costs approximately $0.54. The HRM sap flow sensor, which is comprised of two temperature probes and one heat pulse probe, costs approximately $6.71 including three meters of extension wire. The TDM sap flow sensor, which is comprised of one reference temperature probe and one constant heater probe, costs approximately $6.27, which includes three meters of extension wire. The dual voltage regulator for the TDM costs approximately $6.68. The 2010 price guide from Dynamax lists the TDP30 sap velocity probes with 10 feet of extension wire at $330 each and the Dynamax Dual-Adjustable Voltage Regulator is priced at $340 each.

The TDM sensor designed for the field experiment was closely matched to the Dynamax TDP30 sensor. However, it should be noted that there are some minor differences between the TDM and Dynamax sensors. These include the thermo-junction type and the probe needle dimensions. The Dynamax TDP30 sap flow sensor uses Type T thermo-junctions and has a metal probe with a 30 mm length and a 1.27 mm outer diameter. The thermo-junction type was changed to Type E in the lab-made design. As for the needle probe, the closest length available with a 1.27 mm outer diameter was 38.1 mm. Even though a longer needle is used for the lab-made probes, the measurement length was found not to be a problem since the location at which

the thermo-junction is positioned inside the probe can be placed to match that of the commercial sensor.

Figure 6a and Figure 7a show the calibration results from testing the sensor response of typical temperature and constant heat probes to a standard thermocouple temperature, respectively. Figure 6b and Figure 7b show the residuals based on the regression between the probe output (measured in mV) and the standard temperature (measured in °C). The mean of the residuals of 24 temperature probes tested is 0.0276 °C, and the mean of the standard deviations of the residuals is 2.15 °C. For the constant heat probe, the insulation around the needle tends to cause underestimation of the actual temperatures (see Figure 7a). Lag times for the rapid heating and cooling periods caused more outliers in the datasets, hence the longer test time. The mean of the residuals of 24 heater probes is 0.0315 °C, and the mean of the standard deviations of the residuals is 2.50 °C. In both cases, the mean of the residuals is close to zero and the mean of the standard deviations is small, indicating the stability and quality of the lab-made temperature probes. However, the values are slightly larger for the constant heat probe compared to the lab-made temperature probes.

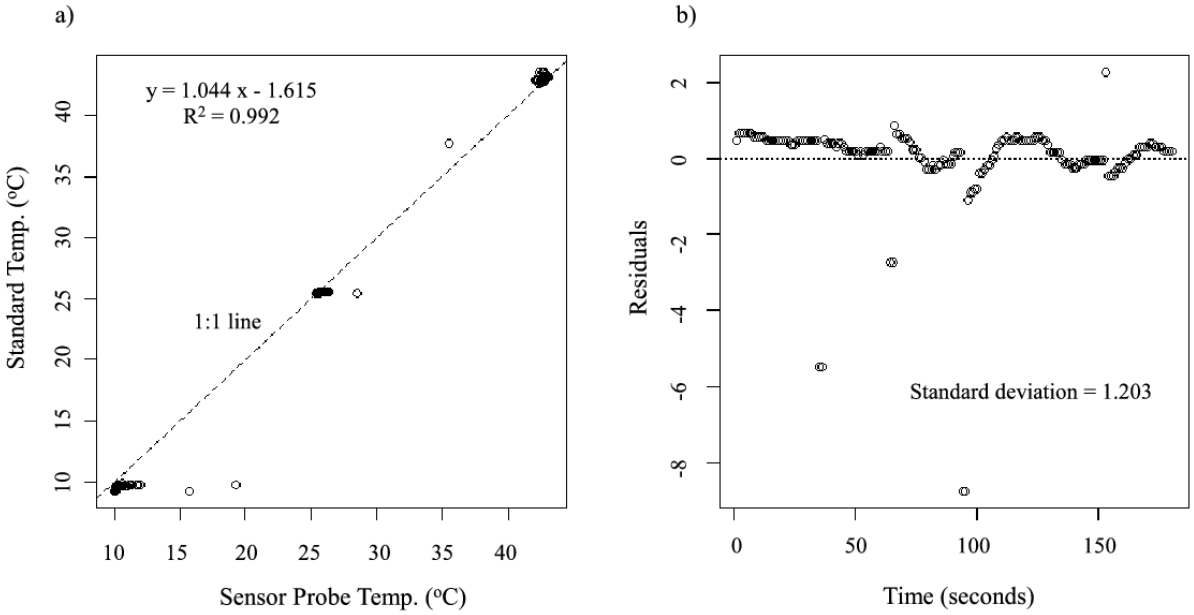

Figure 6. (a) Linear regression of a typical temperature probe's output to the standard temperature. (b) Regression residuals from the temperature conversion equation and their standard deviation. (Davis et al., 2012a)

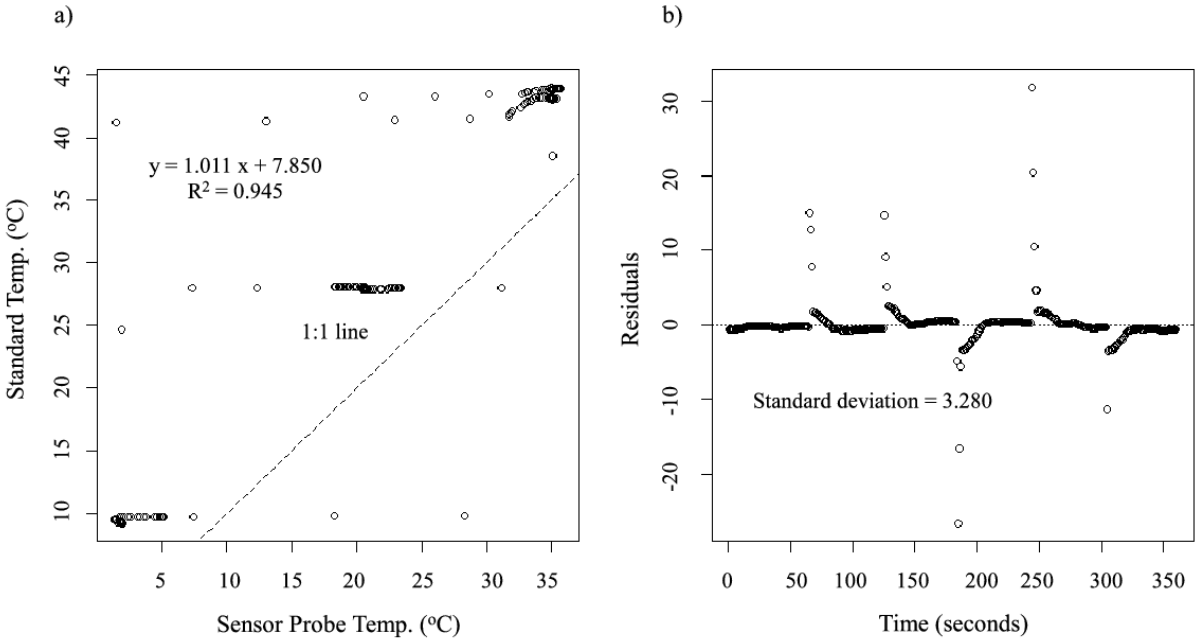

Figure 7. (a) Linear regression of a typical heater probe's output to the standard temperature. (b) Regression residuals from the temperature conversion equation and their standard deviation. (Davis et al., 2012a)

The field measurements were made using both the lab-made TDM and HRM sap flow sensors alongside the commercially available Dynamax TDP30 transpiration sensor. All three sensors were tested in a maple tree, which is approximately 40 cm in diameter, over a period of five days near the end of August 2008. Data was collected on a Campbell Scientific CR-1000 data logger at a 1-s sampling interval. Sap flow rates and pulse velocities were calculated every 15 min. Based on the observed minimum sap flow rates, the 24-hr period for determining the maximum temperature difference for the TDM calculations was set to start at 06:00:00. Corrections for the temperature and constant heat probes, as described in Section 2.2.1.5, were applied to both sets of lab-made sap flow sensors. The correction equations for the TDM constant heat probe (HP) and reference temperature probe (TP) are:

Equation 2.4

$$HP_{temp}(corrected) = 1.1065 \cdot HP_{temp}(raw) - 1.4742$$

Equation 2.5

$$TP_{temp}(corrected) = 1.0539 \cdot TP_{temp}(raw) - 2.8209$$

The correction equations for the HRM downstream (DP) and upstream (UP) temperature probes, respectively, are:

Equation 2.6

$$DP_{temp}(corrected) = 1.0532 \cdot DP_{temp}(raw) - 2.6328$$

Equation 2.7

$$UP_{temp}(corrected) = 1.0451 \cdot UP_{temp}(raw) - 2.5195$$

The results of these measurements are shown in Figure 8. Due to the complexities in determining the dynamic in situ thermal properties of the tree's woody matrix, only the heat pulse velocity for the HRM sensor is plotted.

Based on the general shape of the diurnal cycle and magnitude of the measurements (for the TDM), Figure 8 suggests that the lab-made sensors are capable of adequately capturing daily sap flows. It can be seen that the lab-made TDM sensor results in higher sap flow values compared to the Dynamax TDP30 sensor for all five days of measurements. The magnitude differences are not substantially different given that the measurements were made in different locations on the tree. Notice that the peak sap flow occurs near the same time of day on each of the five days for the lab-made TDM sensor. The spike in sap flow for the lab-made TDM sensor occurs between 15:00 and 15:30 each day. This mid-day spike in the lab-made TDM sap flow measurements, which is not present in the commercial sensor's measurements, may be attributed to a natural thermal gradient caused by solar radiation over the measurement area as described in Lu et al. (2004). The commercial TDP30 sensor peaks sometime between 13:00 and 15:00 except for the first day where it peaks closer to 12:45. The morning and evening sap flow trends are similar in slope and time of occurrence for both the lab-made TDM and commercial TDP30 sensors.

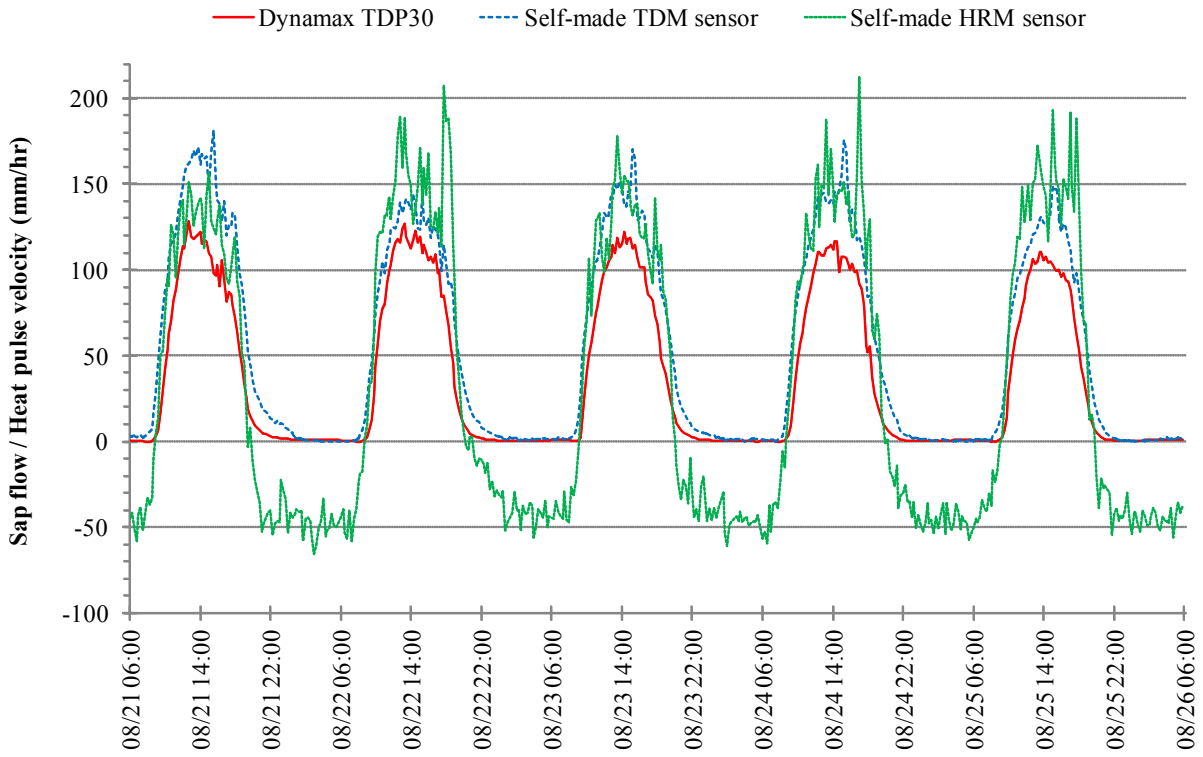

Figure 8. August 21–26, 2008 comparison of Dynamax TDP30 sap flow sensor (solid red line) to lab-made TDP (dashed blue line) and HRM (dotted green line) sensors in a maple tree. (Davis et al., 2012a)

Figure 8 shows that the HRM sensor results are in good agreement with the lab-made TDM sensor. The onset of positive sap flow in the HRM sensor occurs at the same time for both the lab-made TDM and commercial TDP30 sensors. The high variability during times of peak sap flow rates makes it difficult to compare the time of peak sap flow occurrence. Note that the HRM calculations are for the heat pulse velocity, not the sap velocity, so the magnitudes of the HRM results cannot be directly compared to the results of the lab-made TDM or commercial TDP30 sensors. While the lab-made TDM and commercial TDP30 sensors show a lagging tail to the day's sap flow measurements, the HRM sensor is quick to measure the end of the day's positive sap flow and immediately begins to measure reverse sap flow rates. The negative sap flow rates in the HRM curve correspond to reverse sap flow (downward movement) that the HRM sensor design measures. The magnitude of the reverse rates are close to others reported using the HRM method with measurements showing reverse sap flow around 25 $mm \cdot hr^{-1}$ for tap roots in *Eucalyptus camaldulensis* (Burgess et al., 1998) and approximately 100 $mm \cdot hr^{-1}$ in the main stem of *Eucalyptus salmonophloia* (Burgess and Bleby, 2006). It is interesting to note that the magnitude of the nighttime HRM sap velocities is approximately one-third of the daytime rates.

TDM sap flow calculations were also made based on the raw temperature measurements from the lab-made sap flow probes. Due to the calculation method for the HRM, the pulse velocity calculations were not significantly affected by the uncorrected temperature data (maximum deviation was 1.16 $mm \cdot hr^{-1}$). The TDM, however, showed a significant difference in sap flow velocities between the corrected and uncorrected temperature calculations. Deviations between the corrected and uncorrected temperature calculations were in excess of 90 $mm \cdot hr^{-1}$ near peak hours, as shown in Figure 9.

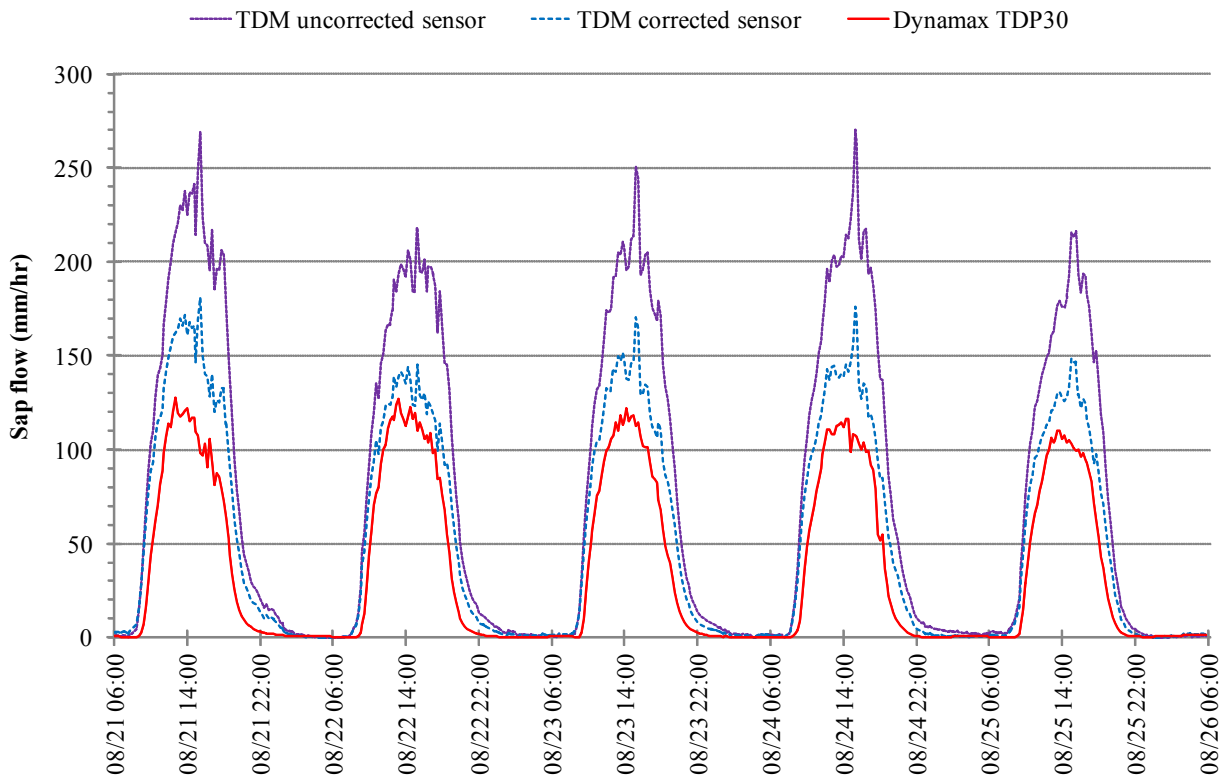

Figure 9. Comparison of sap flow calculations between Dynamax TDP30 sensor (solid red line), temperature-corrected lab-made TDM sensor (dashed blue line) and uncorrected lab-made TDM sensor (dotted purple line). (Davis et al., 2012)

2.4 DISCUSSIONS AND CONCLUSIONS

There are some important considerations that need to be noted when comparing the results of the lab-made and commercial TDM sensors. The first consideration is with the method used for setting the constant heat input of both sensors. The heat input level of the lab-made and commercial sensor was based in this research on matching the maximum temperature difference, ΔT_o in Equation 2.1 similar to that of James et al. (2002) as described in Lu et al. (2004). Adjusting the heat input in this way may alter the heat field around the sensor probe such that Granier's empirical equation used to estimate the sap flow may no longer be completely valid. There is yet to be a study that examines the effects of adjusting the heat field on TDM sensor measurements. The second concern is with the circumferential variations in sap flow. The measurements shown in Figure 8 were made simultaneously on the same tree. The two sap flow sensors were located approximately 90^o apart from one another and at slightly different elevations (approximately 12 cm separation). Numerous studies have shown and reported circumferential sap flow variations with fluctuations as high as 50% (Burgess and Dawson, 2008; James et al., 2002; Jiménez et al., 2000; Čermák et al., 2004). Lastly, there have been suspicions that the Dynamax TDP sap flow sensor systematically underestimates transpiration (Lu et al., 2004; Wilson et al., 2001). Currently there is no information regarding the calibration of the Dynamax TDP sap flow sensor with regards to Granier's original empirical equation. These concerns may account for the apparent over-estimation in sap flow results given by the lab-made sensors compared to their commercial counterparts. Considering these factors and the general shapes of the diurnal patterns (see Figure 8) of the lab-made sap flow sensors, it is

encouraging to see that the lab-made TDM sensors can provide acceptable sap flow measurements compared to the commercial ones (Dynamax TDP30) at roughly 2% of the cost of the commercial sensors.

It should be noted that calibrating Equation 2.1 for the TDM to specific tree species is important. The sensor design given in this chapter does not correct nor address the validation of Equation 2.1 for individual trees. Therefore, the researcher must exercise caution when analyzing sap flow measurements made on tree species that have not been properly calibrated for the TDM.

The large fluctuations present in the peak sap flow values shown in Figure 8 for the lab-made TDM and HRM and in the nighttime sap flow rates for the lab-made HRM can be attributed to the method used to calculate these sap flow values. The 1-s dataset collected for the TDM measurements was sampled at every 15-min value. Given the variance in the measurement values, it is possible that the measurements made at the 15-min time stamps were not the best representative values for the calculations. If the sap flow was instead calculated at the 1-s time interval and averaged over each 15-min period, this would smooth the fluctuations observed in the lab-made data. This suggests that the variance in the commercial sensor is much lower than that with the lab-made sensors. The HRM calculation averages over 40 data points for each 15-min sap flow value, i.e., measurements from 60–100 s. If either v_1 or v_2 are found to be negative or zero, then their ratio, v_1/v_2, is assumed zero. This was chosen to avoid holes in the dataset. Therefore, sensor fluctuation could cause zero values within the 40 data points which would affect the average for that period.

The costs of the lab-made sensors are considerably small, although this may be compensated by the time and effort required to build and test them. One set of sap flow probes

can be built on average in approximately 1 hr. This is based on moderate to large scale production where many sensors are made at the same time. This method of construction increases efficiency and improves the craftsmanship by repeating the same procedure over and over again. While the quality checks and output calibration add additional time to the procedure, it is important to identify and eliminate as many errors as early in the deployment as possible. Figure 9 shows the potential problems that improper temperature correction can cause when using the TDM. It may be assumed that all sap flow sensors made by the same person under the same conditions will behave similarly. This may reduce the need for testing every individual sensor probe's measurement output. Therefore, for lab-made sensor applications, it may be better to use the HRM which is less susceptible to individual probe variations if one does not have time to calibrate individual probes.

In summary, 1) the most cost-effective method for sap flow monitoring is a combination of pulsed and thermal dissipation methods; 2) HRM and TDM sap flow sensors can be constructed based on three probe designs; 3) utilizing common parts from wholesale vendors saves money for large productions; 4) field tests show that careful quality control of sensor output is more important for the TDM sensor than the HRM sensor; 5) compared to the Dynamax TDP30 sap flow sensor, both the HRM and TDM lab-made sensors provide an acceptable alternative to commercial sensors at a fraction of the price.

3.0 ANALYSIS OF POWER CHARACTERISTICS FOR SAP FLOW, SOIL MOISTURE AND SOIL WATER POTENTIAL SENSORS

The majority of the content contained in this chapter was previously published in *IEEE Sensors Journal* under the title, "Analysis of Power Characteristics for Sap Flow, Soil Moisture and Soil Water Potential Sensors in Wireless Sensor Networking Systems." References to the published paper are cited as Davis et al. (2012b).

This chapter examines the power characteristics of three environmental sensors used for monitoring aspects of hydrology (i.e., sap flow, soil moisture, and soil water potential). The power characteristics are examined for two reasons. First, two different methods of sap flow monitoring are examined to determine which one is practical for use in a WSN. Second, the mote power is examined (for all three environmental sensors) to determine how power levels impact sensor measurements.

In an early report discussing WSNs for environmental monitoring, Mainwaring et al. (2002) outline the general application requirements of wireless monitoring systems. Included in these requirements are the reliability and longevity of the network nodes and their measurements. The longevity of the system is highly dependent on the battery life of the individual nodes (or motes) that make up the network. One of the prevailing concerns of WSNs is the characteristics of battery power of the individual wireless motes. Previous studies focus on the power characteristics mainly from a WSN communication perspective (e.g., Anastasia et al., 2009;

Beutel et al., 2010; Ye et al., 2004). The effect on battery life for a given variety of adopted sensors has not yet been investigated. This chapter first examines, from a WSN sensing perspective, how different sensor characteristics, due to their underlying design (such as different types of sap flow sensors), are affected by the sampling/transmission rate and thus affect the mote's battery. This chapter then looks at various transmission intervals to determine its effect on the mote's battery life.

In addition, an individual mote's battery power level can significantly influence the accuracy of measurements made by its sensors. Data acquisition boards are generally used to connect wireless motes to environmental sensors. These acquisition boards utilize the power from the wireless mote to take their measurements. When mote battery power is low, it may not adequately power the sensors and data acquisition board for making and collecting the measurements. A wide variety of environmental sensors are used in WSNs including, for example, microclimatological and soil water status sensors (SensorScope project[1] seismoacoustic sensors (Werner-Allen et al., 2006); solar radiation sensors (Tolle et al., 2005); Medusa light sensors (Selavo et al., 2007); and passive infrared detectors (Szewczyk et al., 2004b) to name a few. While most of the battery power considerations of previous works deal with implementing energy efficient protocols or network health monitoring, none delve deeply into the relationship between mote battery power and its affect on sensor measurements. Ruiz-Garcia et al. (2008) showed in their work that ZigBee-based wireless motes (Xbow) had large errors in their temperature and relative humidity measurements at low battery voltages (around 2160 mV) during their monitoring of refrigeration chamber temperature fluctuations. Their findings of the large errors in those measurements, together with the findings of this study, was the motivation

[1] http://sensorscope.epfl.ch

to examine the impacts of low battery power on measurement accuracy. Therefore this work investigates how to effectively eliminate the negative impacts of low battery power on the measurement accuracy, thus equivalently extending motes' battery life for the adequacy and reliability of data sampling.

3.1 MEASUREMENT QUALITY OF SAP FLOW SENSORS AND THEIR APPLICABILITY TO WIRELESS SENSOR NETWORKS

The most applicable sap flow sensing method is determined based on the quantity of measurements required to achieve quality calculation results. The two typical methods of calculating sap flow, based on active thermometric measurements, are the thermal dissipation method (TDM) and the heat ratio method (HRM). While heat is used as a tracer for determining sap velocity in both methods, the principle behind each is quite different. The HRM determines the sap velocity by measuring the travel time of a heat pulse between measurement points while the TDM uses a mass balance approach to empirically determine the sap velocity based on the amount of heat diffusion within the porous wood between two measurement points (Swanson, 1994). The main differences between these two methods are their power and sampling requirements. Because both methods require external power (not supported by the wireless mote) for the heat tracer, the only effect on mote battery power examined in this study is from the measurement sampling scheme. In particular, the influence of the sampling scheme on the two different sap flow monitoring methods (i.e., TDM and HRM) is examined based on data performance and power considerations of a WSN application. The tradeoff between the sampling interval and mote battery life is used to determine an appropriate method for sap flow

monitoring for WSN applications. These findings can be important for future applications of WSNs in which sap flow sensors are used.

3.1.1 Sap flow methods

The purpose of this section is to examine how the accuracy of sap flow calculations based on the two widely used methods (i.e., TDM and HRM) changes under various sampling schemes. This is important for determining the effect each sampling method has on mote battery power. The sampling interval will be incremented for each method to determine a range which produces results with reasonable accuracy. The TDM and HRM and their respective sensor designs are presented in Chapter 2.0 .

3.1.2 Sap flow sampling measurements

Sap flow estimates are compared at various sampling intervals to investigate the impact that different sampling schemes have on the accuracy of the sap flow calculations. In order to ensure the highest quality of readings at a fine resolution for the sap flow estimates, measurements for both methods were made on a CR1000 data logger, manufactured by Campbell Scientific. The measurements were taken over four days, from 22–25 August 2008, on a sugar maple tree, approximately 40 cm in diameter and 15 m tall, located near 40.61°N, 80.05°W. The sap flow sensors used were lab-made based on designs by the University of California at Berkeley and were thoroughly tested and validated for their satisfactory accuracy compared to commercial sap flow sensors (Davis et al., 2012a, also Chapter 2.0). Measurements were collected each day at a 1-s time interval. For the HRM, the heater probe is powered in 2 to 4 s pulses and remains

unpowered during a 15-min interim between pulses. This restrains the finest resolution for sap flow calculations for the HRM to be every 15 min.

To create the other sampling interval datasets, the 1-s data series was down-sampled to specified larger time intervals. No data aggregation took place during this process. From the 1-s TDM measurement time series, five time series sub-datasets were created with sampling intervals of 10, 60, 180, 300, and 900 s. From the HRM measurement time series, eight time series sub-datasets were created with sampling time intervals of 2, 3, 4, 5, 10, 20, 30, and 60 s. These fifteen datasets (two original and thirteen sub-datasets) were then used to calculate, accordingly, the sap flux density, Q_s, for the TDM (Equation 2.1) and pulse velocity, v_h, for the HRM (Equation 2.2). The datasets calculated from the 1-s time interval measurements serve as the basis for comparison in the sampling interval analysis.

3.1.3 Sap flow calculation and results

Figure 10a–h shows the results for each of the four days' measurements. For the sake of simplicity, only the pulse velocity is given for the HRM in Figure 10e–h. The sub-datasets are also plotted with the 1-s time series. The greatest amount of variability was found near the time-of-peak sap flow rates.

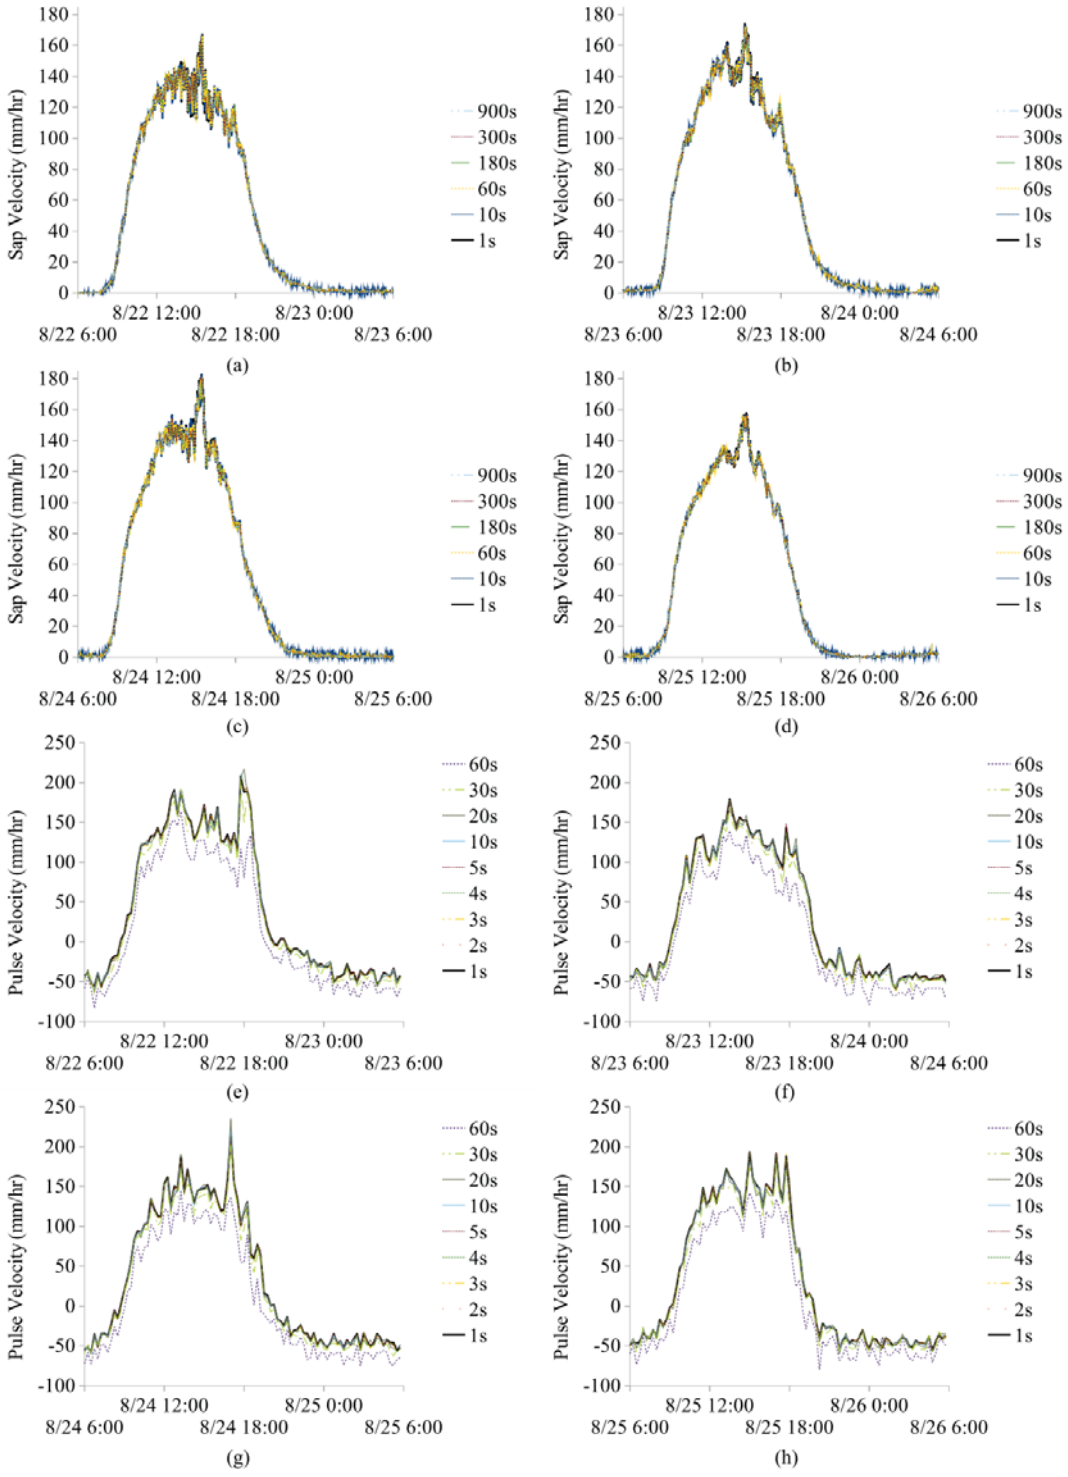

Figure 10. Sap flow measurements plotted with the sub-datasets for the TDM (a)-(d) and HRM (e)-(h). (Davis et al., 2012b)

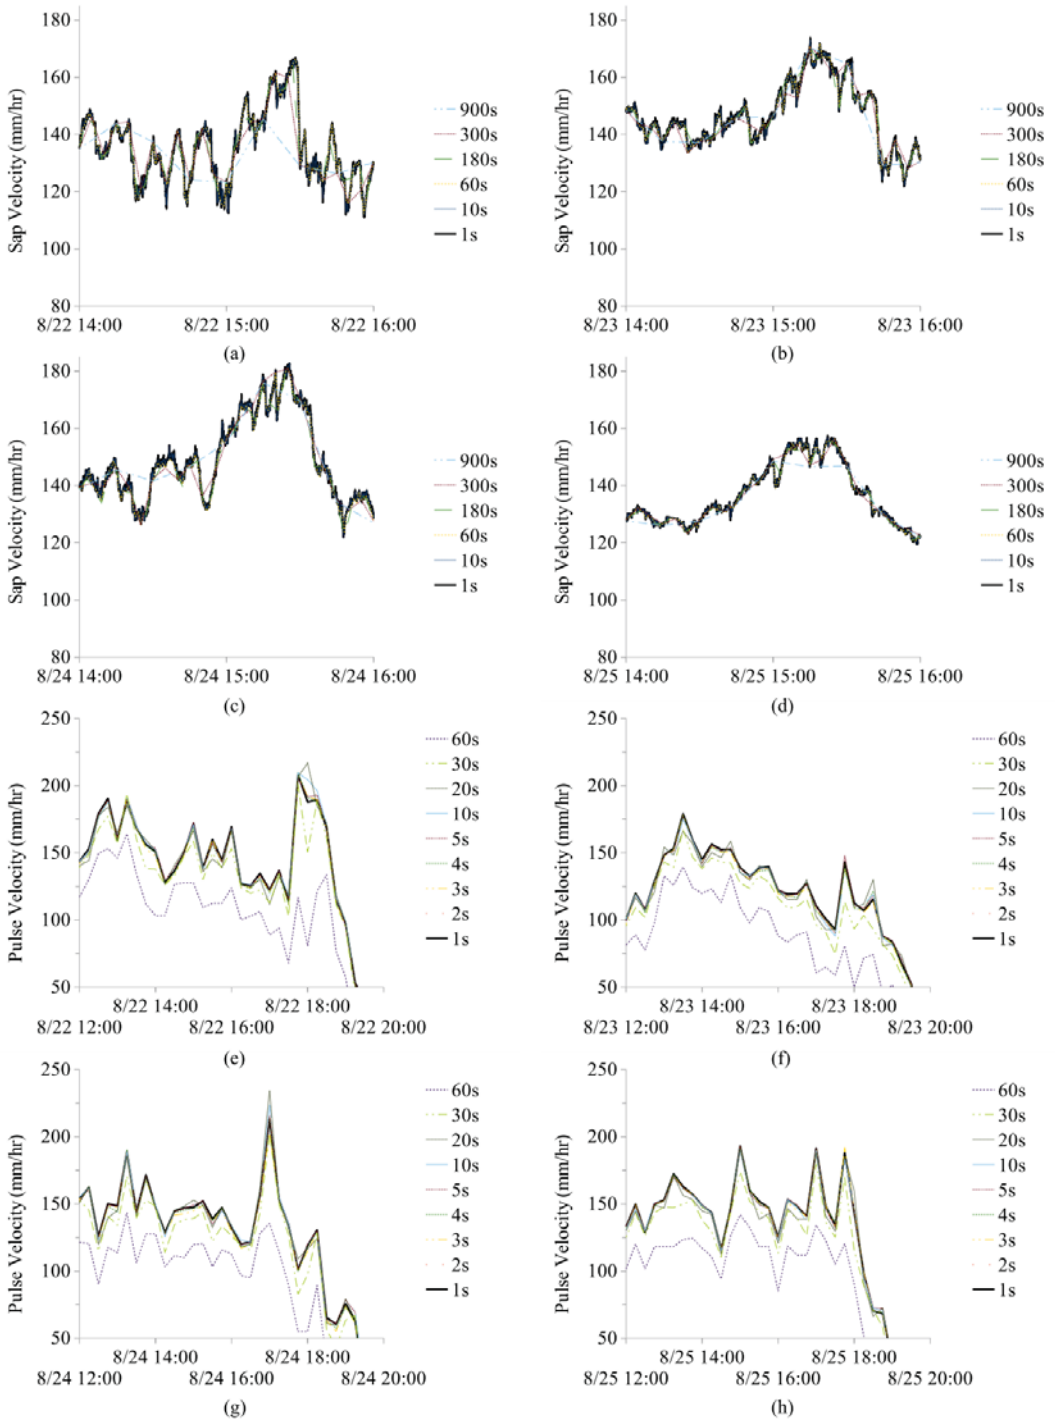

Figure 11. Time-of-peak sap flow calculations for the measurements and sub-datasets of the TDM, 14:00-16:00 hours (a)-(d), and HRM, 12:00-20:00 hours (e)-(h). Note different scales are used for the vertical axis between plots (a)-(d) and plots (e)-(h). (Davis et al., 2012b)

To examine the error due to the different sampling intervals, a 2-hr window over the peak sap flow occurrence, 14:00–16:00 hours, is shown in Figure 11a–d. Due to the higher variability in the daily peak HRM measurements, an 8-hr window over the maximum sap flow occurrence, 12:00–20:00 hours, is shown in Figure 11e–h.

From Figure 11a–d, there is agreement between the 1-s data and sampling intervals up to 300 s. The 900-s sampling interval begins to have large separations from the major peaks and valleys shown in the original dataset. The largest separation can be seen on 22 August, Figure 11a, where the 900-s peak sap flow underestimates the actual peak sap flow rate by approximately 20 $mm \cdot hr^{-1}$ (12.6% of the peak sap flow). On average, for the four days, the 900-s sampling interval underestimated the 1-s data by approximately 5 $mm \cdot hr^{-1}$ (3.2% of the maximum sap flow).

From Figure 11e–h, there is agreement up to the 20-s time series dataset. The 30-s series begins to show some slight under-estimation while the 60-s series is consistently under-estimating the pulse velocities. Similar to the TDM results, the largest under-estimation occurs again on 22 August, as shown in Figure 11e. The average 60-s velocity during peak hours under-estimated the average 1-s maximum velocity by approximately 25%.

A full error analysis was done on each of the four days' data. The root mean squared error (RMSE) was calculated for each of the sub-datasets of both sap flow methods. First, each of the TDM sub-datasets (10, 60, 180, 300 and 900-s) were refined to a 1-s resolution by means of averaging between pairs of given data. The averaging was done exclusively to maintain all of the original sampled data points. All of the data between two consecutive sampled points are equivalent to the average of the two sampled data points (a constant value, not linearized). All of the new 1-s TDM datasets were aggregated to a 15-min resolution. This includes the original 1-s

TDM dataset, which was believed to be the "best achievable" 15-min aggregated dataset. For the HRM, the sap flow calculations based on the 1-s dataset were used as the basis for comparison. The RMSE is calculated as follows:

Equation 3.1

$$RMSE = \sqrt{\frac{\sum (u - U)^2}{N}}$$

where: u = Original 1-s dataset values

U = Values from a sub-dataset

N = Total number of data points (e.g., 96 points/day)

Table 3 shows the RMSE, in units of $mm \cdot hr^{-1}$, for each of the sub-datasets for the four days of measurements. Also included is the average RMSE of each sub-dataset for the four days.

Table 3. RMSE calculation results for the TDM and HRM ($mm \cdot hr^{-1}$). (Davis et al., 2012b)

	TDM				
	10 s	**60 s**	**180 s**	**300 s**	**900 s**
22 Aug	0.028	0.135	0.502	0.953	2.93
23 Aug	0.142	0.178	0.504	0.651	1.45
24 Aug	0.026	0.482	0.954	0.838	1.88
25 Aug	0.035	0.148	0.231	0.346	1.17
Average	0.058	0.236	0.548	0.697	1.86

	HRM							
	2 s	**3 s**	**4 s**	**5 s**	**10 s**	**20 s**	**30 s**	**60 s**
22 Aug	1.42	2.01	1.77	1.39	2.79	4.93	8.68	30.4
23 Aug	1.25	1.99	1.77	1.44	2.50	4.50	8.99	27.4
24 Aug	1.47	2.21	1.61	1.45	2.64	4.21	8.69	28.6
25 Aug	0.73	1.03	1.21	1.30	2.30	4.55	8.94	27.7
Average	1.22	1.81	1.59	1.40	2.56	4.55	8.82	28.5

The RMSE results in Table 3 show that the HRM is much more sensitive to increases in the sampling interval than the TDM, with RMSE values 50-100 times greater than the TDM for similar sampling intervals. Figure 12 shows a graphical comparison of the RMSE results for the TDM and HRM and includes the one-standard deviation error bounds (based on the four days' results). From Table 3, it can be clearly seen that the TDM calculated sap flow results at the 900-s sampling interval are comparable to the sap flow estimations based on the 10-s sampling interval using the HRM.

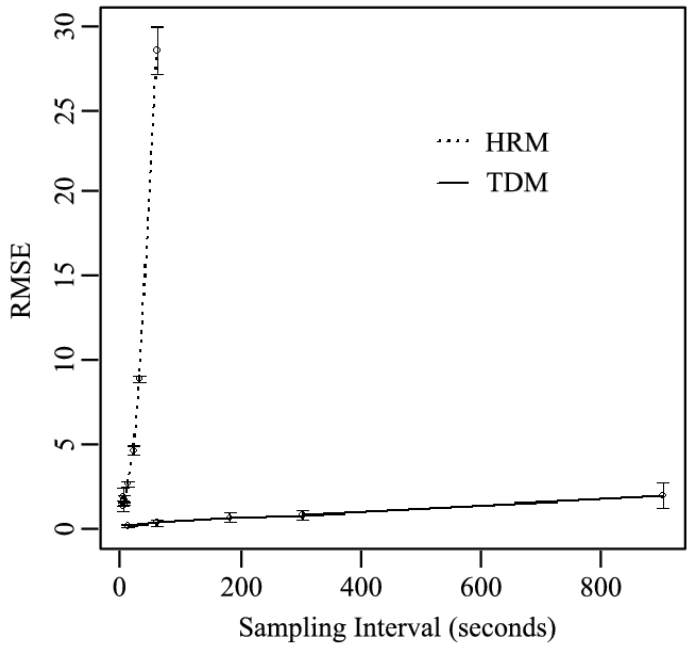

Figure 12. RMSE comparison for the TDM and HRM sampling frequencies with the one-standard deviation bounds. (Davis et al., 2012b)

This analysis shows that one can achieve similar levels of accuracy with the TDM sampled every 900 s to those with the HRM sampled every 10 s. This difference in the sampling interval indicates that the TDM can provide accurate sap flow estimates using significantly less data than the HRM. Most of the power consumed by a wireless mote happens during the receiving and transmission of data packets as shown, for example, in Mainwaring et al. (2002, Table 2), Beutel et al. (2010, Table 3), and Crossbow (2007b, Table 6-2). Since less data needs to be sampled and transmitted through the WSN by using the TDM, this implies that the TDM promotes battery power savings. While the HRM remains a popular method for sap flow monitoring due to its minimal heating requirement and theoretical basis, the variable sampling rates (e.g. high sampling rate following a heat pulse and low sampling rate between pulses) required to get accurate measurements makes it less favorable for WSN applications. Therefore the use of the TDM for sap flow monitoring is more appropriate for WSN applications. This finding is applicable to any WSN that uses sap flow sensors.

3.2 EXTENDING WIRELESS MOTE SENSOR MEASUREMENTS INTO LOW BATTERY POWER CONDITIONS

Additional battery life can be obtained in WSNs by extending the time of measurements during motes' low battery power levels. It has been shown that wireless communications continue at low battery power levels. Understanding the uncertainties in sensor measurement accuracy taken at low battery power levels is the focus of this section. The natural variability inherent in the data collection device, namely the wireless mote, is first examined.

Sap flow, soil moisture, and soil water potential sensors are then tested to determine the effect of mote battery power on their measurements. Each sensor represents a different power requirement by the mote. The sap flow sensors require no input from the mote and return an analog voltage which is first processed by an amplifier by an appropriate gain. The soil moisture sensors require a short excitation of 2.5 V. Measurement results are calibrated for this power level only. The soil water potential sensor also requires a short excitation voltage from the mote, but does not require calibration for input voltages between 2–5 V.

Based on the individual sensor measurement results, it will be determined if a correction algorithm is necessary for each individual sensor or whether a general correction can be applied to the wireless mote's results collectively.

3.2.1 Mote testing and variability

Motes were first tested to determine if any uncertainty present in the battery life is caused by variability in individual motes or sets of batteries. These tests were conducted using the MICAz wireless mote, equipped with an MDA300 data acquisition board, both manufactured by Crossbow Technology (now a part of Memsic Inc.). The purpose of using the older style MICAz motes is twofold. The first is because of the large investment that has already been made by so many researchers, agriculturalists, environmentalists, etc. using this platform. While new wireless modules and architectures are being developed (e.g., Iris, Lotus, Wasp, etc.), there still exists a large populace using the MICAz platform. The second is that MICAz wireless modules are still being produced and distributed today, making them a viable option for continued WSN applications. The wireless motes were programmed using their multi-hop, ad-hoc, low power mesh networking protocol (i.e., XMesh) using the NesC programming language supported on the

mote's TinyOS operating system. The mote's CC2420 radio transmits data in the 2.4 GHz band and was set in the range of 2.405–2.425 GHz in these experiments. The radio's transmission power was left at its default maximum power (0 dBm) setting. The sampling and transmission interval, which is designed to be the same value, was set to 1 s. Eight motes were selected at random. Four were powered using two 2500 mAh AA batteries and another four motes were powered using two 2450 mAh AA batteries. Each mote was tested four times. Table 4 summarizes the results of these tests for the eight individual motes tested each four times. The averages and standard deviations are given based on the four tests. The starting and ending battery voltages (based on the two AA batteries) for the experiment represent the first and last recorded voltage value for the motes. The number of data packets, shown in thousands, represents a transmitted message that was received by the base station from a mote during its operation life. The battery life, shown in minutes, represents the length of time that the motes transmitted data packets.

Table 4. Summary of the mote and battery variability experiment of eight individual motes with the means and standard deviations based on four tests. (Davis et al., 2012b)

	Case 1: 2500 mAh			
Mote #	1	2	3	4
Start (volts)	2.73±0.13	2.63±0.12	2.82±0.14	2.63±0.13
End (volts)	1.97±0.00	1.82±0.11	1.99±0.00	1.87±0.01
Pkts (thou.)	137±23.6	119±29.0	138±25.6	132±22.9
Life (min.)	4750±856	4170±1078	4870±924	4700±850

	Case 2: 2450 mAh			
Mote #	5	6	7	8
Start (volts)	2.75±0.02	2.74±0.04	2.79±0.05	2.75±0.04
End (volts)	1.93±0.02	1.93±0.01	2.00±0.08	1.90±0.02
Pkts (thou.)	66.1±21.3	86.1±26.4	84.4±16.4	91.7±24.9
Life (min.)	2360±747	3100±952	3070±609	3380±927

Examining the results of the battery life (last row) in Table 4, the variability from both the mote and batteries is represented in the standard deviation, shown for each of the motes in both cases. The standard deviations of the battery life results are more representative of the variability of the batteries assuming that individual mote variability is minimal. To investigate the variability of the individual motes, the averages of the battery life results for the four motes in both cases were used to calculate the standard deviation, which can serve as an approximate measure of the variability of the motes. This is because the battery variability has already been accounted for in the standard deviations presented in Table 4. The means and standard deviations of the average battery life of the four motes in Case 1 (i.e., 4750, 4170, 4870, and 4700 min) and the four motes in Case 2 (i.e., 2360, 3100, 3070, and 3380 min) are 4620 (±310) and 2978 (±435) min, respectively. Comparing the standard deviations based on the four motes'

results (i.e., 310 min in Case 1 and 435 min in Case 2) to the individual standard deviation values presented in Table 4 shows that the mote variability is less than that of the batteries. Based on this, it can be said that variability is present in both the individual motes and their batteries. However, the variability of the batteries has a larger influence (approximately double) than that of the motes on the operational battery life. It is interesting to note that there is small variability in the ending voltage for all of the individual motes (second row in Table 4). This suggests that individual motes have a specific lower operating voltage. It should also be noted that all eight motes tested operated at battery voltage levels below 2.2 V.

3.2.2 Low mote battery power sensor measurements

The effect of low battery power on individual sensors is examined in this section for the possibility of retrieving measurement data when wireless motes are under low battery power conditions. The purpose for this is to exploit as much battery life from the mote as possible by extending the lower limit of the operating voltage. The recommended mote operational battery voltage range is 2.7–3.6 V (Crossbow, 2007b). Motes, however, can still function at battery voltages below 2.7 V. Data collected at battery voltages below 2.7 V are not guaranteed to be an accurate representation of their actual sampled values. This was shown by the attenuation in sensor data measured using the Crossbow motes at low battery voltages presented in Ruiz-Garcia et al. (2008) and in this study's measurements as well. In Ruiz-Garcia et al. (2008), they showed that data values rapidly increased in value as the battery voltage decreased below 2.16 V. In an effort to increase the mote operation life to include this lower voltage range where the motes are still functional, the impacts of such signal attenuations on the sensor data readings and how the

readings may be adjusted were investigated. The sensors used include those measuring soil moisture, soil water potential, and xylem sap flow.

3.2.2.1 Experiment design and results

Three motes were selected for this experiment, based on Table 4 results. Each mote was equipped with seven sensors. The sensors that were tested include the sap flow sensor (lab-made temperature probes), the EC-5 soil moisture sensor and the MPS-1 soil water potential sensor (both by Decagon Devices, Inc.). Four sap flow sensor temperature probes, two soil moisture sensors, and one soil water potential sensor were connected to the mote via the MDA300 data acquisition board (see Figure 13). Each of the three sets of sensors were placed into their own planter pot, approximately 7000 cm^3, of mixed soil and water and were covered in plastic to minimize evaporative losses. The soil moisture and water potential sensors were buried at a depth of 5–10 cm. The sap flow temperature probes were inserted vertically into the soil to needle depth (i.e., 4 cm). Although the temperature probes are not measuring sap flow, they are instead used to measure the variations of the soil temperature in the vicinity of the soil surface. Thus, the behavior of the thermocouple probes can be examined as a function of battery life in a similar way as they are used to measure xylem sap flow.

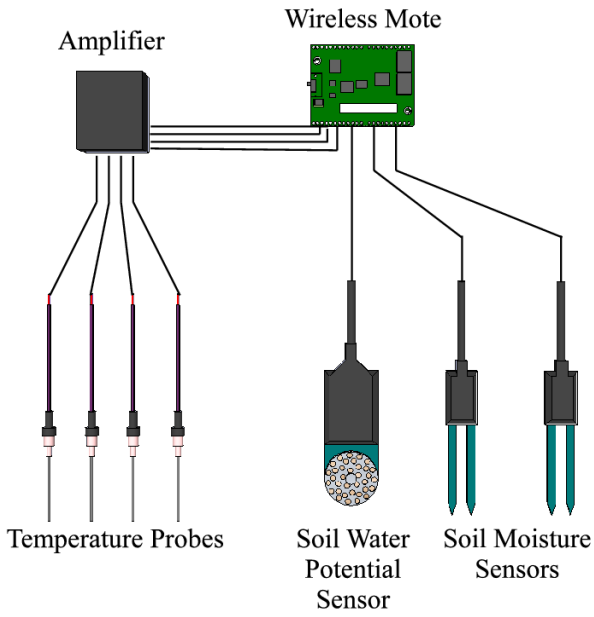

Figure 13. Schematic showing 4x sap flow temperature probes, 1x MPS-1 sensor, and 2x EC-5 sensors connected to a MICAz wireless mote. (Davis et al., 2012b)

Measurements were made every 30 s using two types of batteries: 1) rechargeable 1.2 V AA batteries and 2) industrial 1.5 V AA batteries. A comparison is made between the measured results from the rechargeable batteries and those from the industrial batteries to see if the battery type has any influence on the readings. Also, as the batteries run below the lower bound of the suggested operating voltage, a comparison of the results is made between the low powered readings and those made in the optimal operating voltage range. Figure 14 shows results of using the two different sets of batteries on the same mote.

The battery voltage curves in Figure 14 show the difference in how the rechargeable batteries and the industrial batteries discharge over time. The EC-5 soil moisture sensor shows little variation over the battery life until it loses stability for both the rechargeable and industrial batteries at a battery voltage of about 2.4 V. The MPS-1 and temperature probes (i.e., TC)

appear to be influenced by the low battery power. The shape of the MPS-1 measurements with the rechargeable batteries (Figure 14a) follows the voltage decay curve, in contrast with the industrial batteries (Figure 14b) which follows the inverse of the voltage decay curve. The temperature probes all show a nearly linear increase in their measurement values at battery voltages below 2.5 V (rechargeable batteries) and 2.6 V (industrial batteries). Since no additional water was added during these measurements, the small increase in soil moisture near the end of 23 November 2009 and from 27–29 November 2009 suggests that the battery voltage has influenced the EC-5 sensors to a lesser degree.

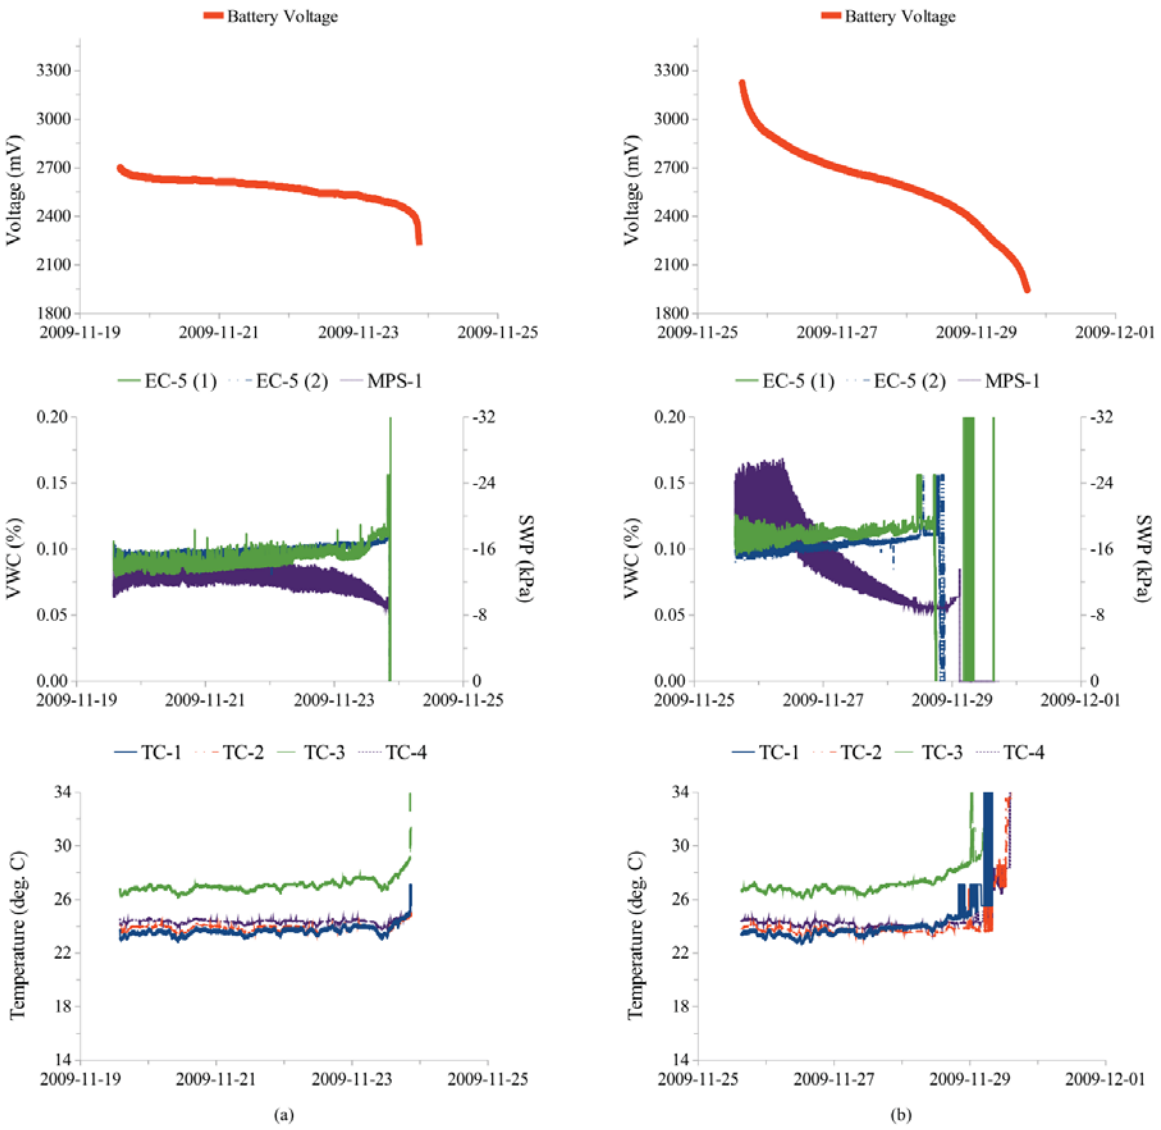

Figure 14. Battery voltage, soil moisture, soil water potential, and sap flow probe temperature for a mote powered by two (a) 1.2 V rechargeable AA batteries and (b) 1.5 V industrial AA batteries. The soil moisture and water potential sensors are buried at a depth of approximately 5–10 cm. The sap flow probes are inserted to needle depth, approximately 4 cm. (Davis et al., 2012b)

3.2.2.2 Sensor measurements under a controlled power supply

To investigate these characteristics further, the mote power supply was manually controlled over a specific voltage range. This was accomplished by utilizing standard 120 VAC regulated through an 800 mA, 3–12 VDC variable output converter[2]. The voltage converter was connected to a voltage regulator allowing the power supply to be decreased in increments ranging from 50–100 mV.

In this experiment, six EC-5 soil moisture sensors, three MPS-1 soil water potential sensors, and two sap flow temperature probes were tested. The EC-5 soil moisture sensors were divided into three sets of two sensors. The first two sets were placed in jars of water and the third set was placed in a pot of dry soil, providing high and low constant conditions for the measurements. The MPS-1 sensors were placed in pots of moist soil mixtures and were covered in plastic to minimize moisture loss. The sap flow temperature probes were placed in jars of room temperature water (~24 °C). An improved circuit board for amplifying the thermocouple response was used to boost the thermocouple probe's signal to the wireless mote (see Section 6.4.2.2 for further details). Figure 15 shows the changes to the connections between the sensors and the wireless mote.

The mote power supply was adjusted from 2.8–2.1 V. At each voltage, 10–12 measurements were taken from each sensor. Based on the sampling interval, each set of measurements took 5–6 min to collect. The average output voltages (mV) for each set of sensors' measurements are plotted in Figure 16a–c.

---

[2] RadioShack AC/DC Adapter (CAT. NO. 273-1667)

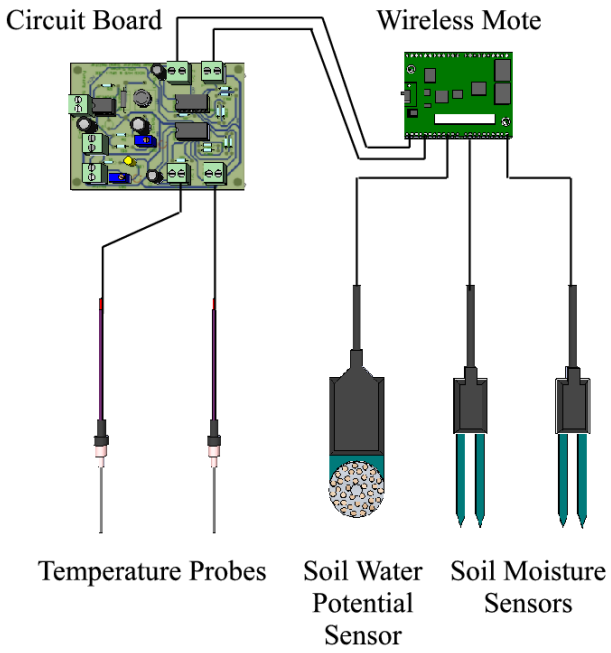

Figure 15. Schematic of improved circuit board to amplify the sap flow temperature probes to the wireless mote. (Davis et al., 2012b)

The EC-5 sensors (Figure 16a) in jars of water (i.e., 1-1, 1-2, 2-1, and 2-2) show a slight decrease in their measurement response between power supply voltages of 2.6–2.4 V. This corresponds to the slight increase in volumetric soil moisture seen in Figure 14a–b over a similar voltage range. Little to no change is seen in the EC-5 sensors placed in the dry soil. The EC-5 sensors began erratic measurement readings (or no readings) below a power supply of 2.3 V.

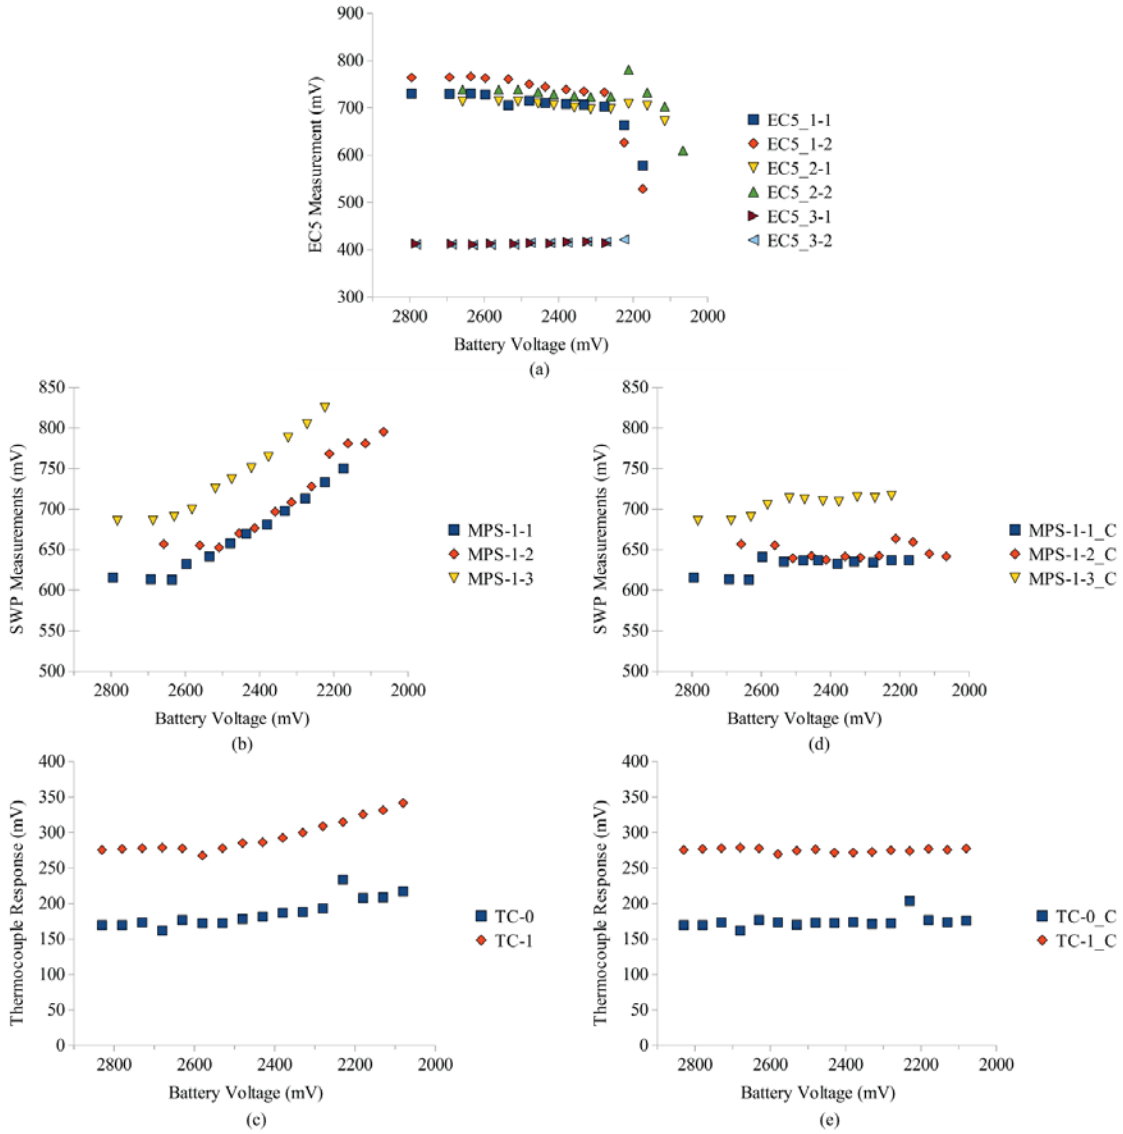

Figure 16. (a) Six EC-5 soil moisture sensor measurements' response to decreasing wireless mote battery voltage. Three motes were tested, each with two soil moisture sensors. The sensors 1-1, 1-2, 2-1, and 2-2 were placed inside jars of water and sensors 3-1 and 3-2 were placed in dry soil. (b) Three MPS-1 soil water potential sensor measurements' response to decreasing mote battery voltage. Three motes, each with one MPS-1 sensor, were used to measure the soil water potential of moist soil mixtures. (c) Two sap flow temperature probe (TC0 and TC1) measurements' response to decreasing mote battery voltage; measurements were taken inside a jar of room temperature water. (d) Three MPS-1 soil water potential sensor measurements' regression corrected response to decreasing mote battery voltage. (e) Two sap flow temperature probe measurements' regression corrected response to decreasing mote battery voltage.

The MPS-1 sensors (Figure 16b) show a constant measurement range for the first two or three power supply voltages tested until around 2.6 V. The measurements then begin a nearly linear increase as the power supply voltage decreases from 2.6–2.2 V. This increase in the measurement voltage (mV) corresponds to a decrease in the negativity of the water potential seen in Figure 14a–b.

The thermocouple measurements (Figure 16c) remain nearly constant until the power supply voltage drops below 2.6 V. After which, the measurements increase in a linear trend. This is similar to the gradual increase in temperature seen in Figure 14a–b.

This information suggests that a regression can be formulated to retrieve the original data from the attenuated results. It is important to note that the results from this second set of experiments (i.e., Figure 16a–c) show similar patterns as those from the first set of the experiments (i.e., Figure 14a–b). The major differences between the first and second experiments are that the battery voltage is carefully controlled and an improved circuit board for amplifying the thermocouple is implemented.

The findings show that the attenuation of the mote's power supply influences only specific types of sensors. Of the three sensor types tested, the EC-5 soil moisture sensor measurements (Figure 14a–b, and Figure 16a) did not show strong low battery power attenuation. One possible reason for this could be due to the specific excitation power requirements for the EC-5 sensor measurements. The EC-5 requires a 2.5 V excitation to power the sensor's measurements. As the mote's power supply decreases below the lower bound of the suggested operating range, the excitation voltage may also decrease which would under power the sensor. This under-powered measurement could be offset by the mote's acquisition board

signal attenuation. Further investigation into the actual reasons why the EC-5 sensor measurements are seemingly unaffected by the mote's low power attenuation is recommended.

3.2.3 Mote signal attenuation correction algorithm

As a final test to quantify and correct the influence of low power supply attenuation that the mote has on measurements, mote measurements were tested over a range of known signal inputs and power supply levels. Three motes were supplied with known input signals produced by a high precision source meter[3]. The input signals that were tested ranged from 100–950 mV, stepped down in increments of 50 mV. This range was chosen because it covers the expected voltage signals from all three of the sensor measurements (see Figure 16a–c).

The mote power supply was once again controlled by a voltage regulator, decremented from 2.6–2.2 V in increments of 50 mV. Measurements were taken on the same three motes from the previous section. For each mote power supply level and input signal, the average of ten output measurements was recorded. From all the data points collected (N=450), samples taken at input voltages of 100, 200, …, and 900 mV were separated from the samples taken at 150, 250, …, and 950 mV. The 150–950 mV half of the dataset was used to build the regression equation and the 100–900 mV remaining half was used for validation. A single outlier was identified in the validation dataset at battery voltage 2609 mV, input voltage 700 mV, and output voltage 395 mV. Figure 17 shows the regression and validation data that were collected. For each input signal tested, a family of outputs was collected based on the varying mote power supply.

[3] Keithley Model 2400 Series SourceMeter

The initial idea for the attenuation correction equation was to make it a linear combination of the two known measurements (i.e. the sensor measurement value and mote battery power) in the form $y = Ax_1 + Bx_2 + C$ where x_1 and x_2 are the two known measurements, y is the corrected sensor measurements and A, B and C are fitted coefficients. After examining the regression curve more closely, it was found that the slope of the regression curve was dependent on the mote's battery power. Therefore, it was decided that a better choice for the attenuation correction would be to use the battery power as a weight for the sensor measurement in the form $y = (Ax_2 + B)x_1 + C$. Looking for the best fit, variations of this form were tested including, $y = Ax_1x_2 + B$ and $y = Ax_1x_2$. The fit of the last form given here resulted in the largest coefficient of determination and smallest standard error.

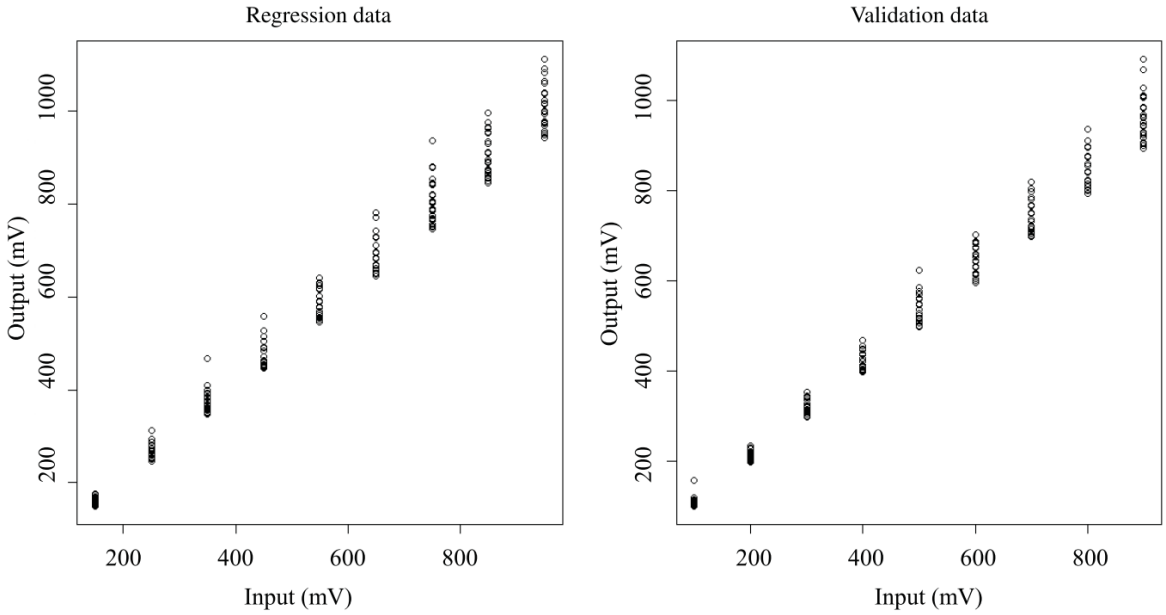

Figure 17. Regression and validation data used for the low power analysis of mote measurement attenuation. The measured (i.e., output) voltages for each of the input voltages are attenuated based on battery power levels ranging from 2.6–2.2 V. (Davis et al., 2012b)

A multivariate regression using the two known values, mote power supply and output signal, was developed such that the decreasing battery voltage would scale the output attenuation. An analysis of the fitted residuals from the linear model using the 225 regression data points (Figure 17) had a negative skew. To reduce the skew, 11 "outliers" were identified and removed. The linear model was re-run on the regression data (N=214) and the resulting equation is:

Equation 3.2

$$V_{actual} = 0.0003905 \cdot V_{battery} \cdot V_{output},$$

where: V_{actual} = Real measurement signal (mV)

$V_{battery}$ = Mote power supply (mV)

V_{output} = Output measurement value collected by the mote (mV)

Note that V_{output} is not sensor specific.

Figure 18a shows the variance of the error from the regression. The residuals now have a mean value near zero with a standard deviation of 7.6 mV. However, the variance is shown to increase in a fan pattern with the fitted value. The normal Q-Q plot (Figure 18b) is nearly linear with some dispersion in the residuals at the tail ends. Figure 19 shows that the majority of the standardized residuals fall within an acceptable range (±2 standard units) with a few out of the ordinary values (±3 standard units). This was not as great of a concern because some unaccounted variability in the mote's data collection was expected. Overall, the residuals appear normally distributed, suggesting that the model is a good fit ($R^2 = 0.9998$).

Figure 20 shows the comparison of the corrected output voltage signal (i.e., V_{actual} in Equation 3.2) and the uncorrected output voltage signal (i.e., V_{output} in Equation 3.2) against the input voltage signal for the validation data set (Figure 17). The variance in the corrected data is smaller and concentrated closer to the 1:1 line for each of the input voltage signals. Improvement can be seen in all levels of input tested. An RMSE analysis shows that the regression reduces the error from 54 in the original data (i.e., V_{output}) to 12 after the correction (i.e., V_{actual}).

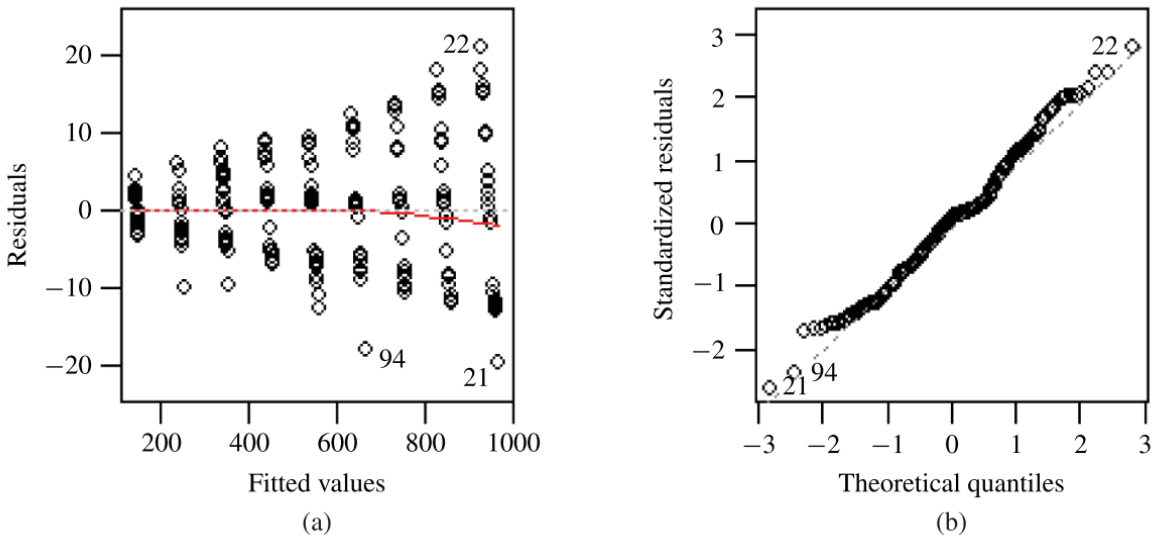

Figure 18. Results of the model fit to the regression data. (a) Residuals versus fitted values. (b) Normal Q-Q plot. (Davis et al., 2012b)

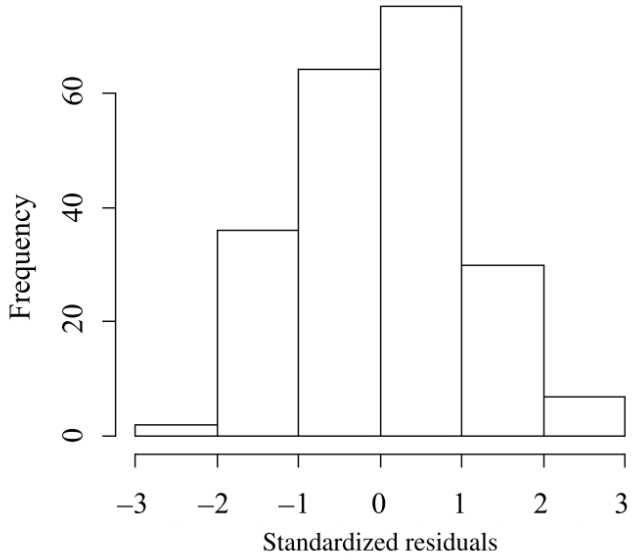

Figure 19. Histogram of the regression model's standardized residuals. (Davis et al., 2012b)

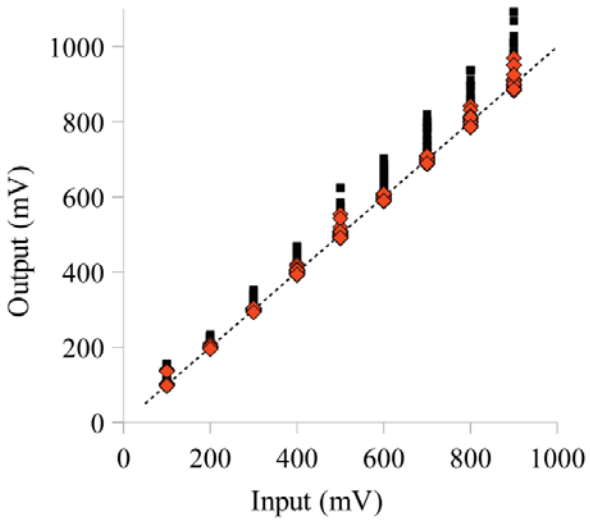

Figure 20. Corrected and uncorrected validation data output versus the known input from Figure 17. The diamonds and the squares are from V_{actual} and V_{output} in Equation 3.2, respectively. The dashed line shows the 1:1 relationship between the actual and input voltages. (Davis et al., 2012b)

After the validation, the regression (Equation 3.2) is used to correct the attenuated sensor data in Figure 16b–c. The EC-5 data were purposefully left out of this regression analysis because they did not exhibit much attenuation from the mote's low power state (see Figure 16a). The attenuation in the MPS-1 sensor data is corrected with the regression (see Figure 16d) for measurements made at mote power supply levels below 2.6 V. The corrected data are closer to the constant value measured at mote power supply levels greater than 2.6 V, overestimating the expected value by approximately 10–30 mV. The sap flow temperature probe measurements were also corrected with the same regression equation (see Figure 16e), overestimating the expected values by only 5–10 mV. The overestimation in both of these corrected datasets is within the range expected based on the variance of the residuals (Figure 19).

In environmental monitoring applications, WSNs are required to reliably deliver high quality measurements over long deployment cycles (Musăloiu-E. et al., 2006). Therefore, it is typical that sensor measurements are made continuously until the mote batteries are expired. When the validity of measurements is only guaranteed under the optimal battery voltage levels, measurements made below this are wasted. In research, it is important to make the best use of all the data available. Extending measurements below their optimal power level is one way of achieving this. In this study, battery life was extended from 2.7 to 2.3 V, adding approximately 1.5 times the battery life compared to the normal operation voltage range. From a practical perspective, none of the sensors tested show reasonable measurements with the battery voltage below 2.3 V. This lower voltage limit for sensor measurements is similar to that reported in Musăloiu-E. et al. (2006).

3.3 CONCLUSIONS

This chapter concludes that the TDM for sap flow monitoring is the better choice for WSN applications (see Table 3). This is due to the relatively low error introduced with the TDM in which the sampling interval can be increased up to 15 min (RMSE≈2.5) while it produces measurement accuracy (e.g., in terms of RMSE) greater than that by the HRM measuring at the 10-s sampling interval.

It is found that the mote's sensor measurements may be affected in the low-power states (2.7–2.3 V). Given the range of sensors tested (Figure 16a–c), the attenuation correction equation may be valid for any single-end analog sensor input. The 8 channel ADC on the ATMega128(L) microcontroller uses a reference voltage (either the digital supply voltage or an

internal 2.56 V reference voltage) (Atmel, 2011). The analog to digital conversion equation uses the inverse of the reference voltage to scale the analog measurement to the digital output. Therefore, when the reference voltage decreases due to low battery power, the conversion scale for the digital output would increase. The exception to this is when the sensor requires a specific power excitation voltage, which is supplied by the mote. Under low-power states, it is hypothesized that the mote's excitation voltage also decreases making the correction algorithm invalid. This was exemplified with the EC-5 sensor which was not significantly influenced by the mote's low battery power. In the future, a more thorough investigation is required to determine the validity of the EC-5 measurements collected at low mote battery voltage. A rigorous examination of the attenuation in both the MPS-1 and temperature probe measurements sampled while the mote's power supply is low (i.e., 2.7–2.3 V) shows that a simple regression model can be devised to correct the data to approximately their original values (see Figure 16d–e). With the newly developed regression, mote battery life can be significantly extended and wireless mote measurements can continue under motes' low power conditions (2.7–2.3 V) beyond the recommended operational voltage range (3.6–2.7 V). This suggests that an additional 150% of the mote's battery life can be utilized for these two types of sensor measurements. The regression equation for the sensors presented here needs further testing under various environmental conditions to verify that its applicability holds outside the laboratory setting. It is also recommended that this concept of the simple regression correction for low mote voltage sensor measurements be applied to other sensor types for the purpose of increasing the battery life for wireless monitoring.

4.0 AN EXPERIMENTAL STUDY OF WSN POWER EFFICIENCY: MICAZ NETWORKS WITH XMESH

The majority of the content contained in this chapter was previously published in *The International Journal of Distributed Sensor Networks*, which is an open access publication, under the title, "An Experimental Study of WSN Power Efficiency: MICAz Networks with XMesh." References to the published paper are cited as Davis et al (2012c).

WSNs have demonstrated their potential and promising application in different fields of science and engineering (Akyildiz et al., 2002; Liu et al., 2011), such as geophysical studies for volcanic activities (Werner-Allen et al., 2006), environmental monitoring in glacier regions (Ingelrest et al., 2010), structural monitoring (Xu et al., 2004), and healthcare applications (Chaczko et al., 2010). However, the severe resource constraints of wireless sensor nodes (e.g., battery power, memory size, processor capacity, and network bandwidth) in WSNs raise new theoretical and practical challenges, drawing great attention in the research community. This chapter focuses on the practical aspects of WSN power efficiency, which is critical in real-world WSN deployments for environmental monitoring. It is essential that domain scientists and engineers who include WSNs in their studies and experiments have a fundamental and comprehensive understanding about the WSN power efficiency and characteristics for different applications under dynamic operational environments.

Multiple research works have attempted to identify key features in WSN deployments (Beutel et al., 2010; Barrenetxea et al., 2008; Teo et al., 2006). While good qualitative results and guidelines are available, the lack of more quantitative results and descriptions is often a major obstacle in today's WSN design, implementations, and deployments. In practice, WSN deployments can lead to various important tradeoffs among quality of service, network performance, power consumption, and operational cost. Understanding these tradeoffs in a quantitative way is fundamental for any successful WSN practice. This chapter quantitatively investigates the power efficiency and battery savings for WSNs with various network characteristics, using the popular MICAz wireless motes for environmental monitoring.

WSNs often use individual node power efficiency as a key performance metric, because the battery power of individual nodes can lead to multiple failure types within the network. For example, Beutel et al. (2010) examined 14 environmental WSN field deployments from year 2002 to 2008, ranging in scale from 3 to 98 nodes. Their examination showed that battery power can affect every level of wireless networking problems which they classified into four categories: node, link, path, and global failures. Node failures can occur when the battery power drops below the level at which the mote can still operate. Link failures can occur when low battery power reduces the range that a mote can transmit its data, effectively removing it from the network's communication topology. Path failures can result when an important node or nodes that route transmitting data have node or link failures forcing the network to use less efficient paths for transmitting data to the base station. Global failures can occur when a critical or bottleneck node experiences node or link failure which cuts off the transmission of data from the whole or part of the network. It is important to identify these problems and their causes such that they can be avoided in field deployments.

For this investigation, the commercially available Memsic's WSN platform was used. It is one of today's most widely used WSN platforms (previously developed by Crossbow Technology), in which motes are programmed in nesC language (Gay et al., 2003), linked together with specific data acquisition boards, to form a mesh WSN for various applications. The application code, compiled and installed on wireless motes, runs under the motes' TinyOS operating system (TinyOS, 2010). Memsic's XMesh routing protocol (Crossbow, 2007c), which features TrueMesh technology, is utilized as the networking foundation for these applications. It provides a mesh network which is self-healing and self-organizing where each mote acts as both a sensor node and a router for its neighbor's data. The ad-hoc formation of the network is based on link estimates made between node neighbors and routes data down a path of lowest transmission cost to the base station where data is stored. Parameter assignment in the nesC code, along with some runtime argument passing during compilation, allows users to have partial control over the mote's power efficiency.

Using Memsic's WSN application platform, motes can be programmed in either a high power (HP), or low power (LP), or extended low power (ELP) operation mode. Note that the ELP operation mode does not support routing. Since multi-hop networks are considered, the ELP operation mode is not included in this study. LP and HP operation modes implement different power efficiencies and thus have different battery savings associated with them. In addition to the two operation modes, the transmission frequency of data packets can also be manually adjusted. The transmission rate also plays an important role on the battery life of the motes, because data transmission is the most power consuming operation, as shown in Beutel et al. (2010, Table 3), Mainwaring et al. (2002, Table 2), and Crossbow (2007b, Table 6-2). Adjustments to the operation mode and transmission frequency can reduce the power

consumption of the motes which increases their operating life. The power consumption of wireless motes is an important consideration when deploying networks as it may affect the sensor readings (Davis et al., 2012b, also Chapter 3.0) and network connections (Beutel et al, 2010).

The battery savings is investigated by implementing the two operation modes over various data collection intervals with Memsic's XMesh WSN. Battery life is measured over the transmission life of the motes and is used to compare the actual battery savings of each configuration. In addition, experiences and lessons learned from the experimental study regarding the WSN gateway Stargate Netbridge's operational robustness are presented. It has been found that the Stargate Netbridge, a Linksys NSLU2 device specially modified by Crossbow, has severe robustness and reliability issues for practical WSN deployments. To address this problem, a solution is presented of how a general integrated network and data management system can facilitate the new deployment of a WSN gateway. It was developed to successfully replace the Stargate Netbridge and significantly improve the operational robustness of WSN deployments.

4.1 LABORATORY METHODS

To determine the effect of different sampling frequencies on battery life, a series of experiments were conducted using the MICAz wireless mote, equipped with a MDA300 data acquisition board, both manufactured by Memsic Corporation (previously Crossbow Technology). The mote's operation mode, sampling, and transmission intervals were investigated in these experiments. The mote's operation mode is set using the XMesh high power (HP) or low power

(LP) configuration. The difference between these two settings is mainly in the bandwidth consumption and latency of the transmissions. The LP mode, which utilizes a sleep function that powers off all unnecessary electronics between operations, has a high latency (i.e. long transmission delay) and low bandwidth consumption (i.e. low data capacity) compared to the HP mode.

The mote's onboard radio power is adjustable within a range of 0 to -25 dBm, the maximum and minimum power allocations, respectively. The MICAz CC2420 radio transceiver (Chipcom AS, 2004) uses IEEE 802.15.4 protocol (IEEE, 2003) and transmits data in the 2.4 GHz frequency band with a maximum data rate of 250 kbits/s. Other investigations have already looked at the effects of the transmission power on mote connectivity and battery power (Teo et al., 2006). The previous study was interested in understanding how the antenna power affected the transmission distance and how implementing various duty cycles increased the battery life. In contrast, this study is more focused on how the data collection rate and multi-hop network functionality affects battery life. The default 0 dBm (1 mW) radio power setting in the 2.405–2.425 GHz frequency channel was used in this investigation. The data message interval (DMI) and the data transmission interval, which are designated to be the same value, can be set to any multiple of seconds. Intervals were chosen from 10 to 900 s.

4.1.1 Experiments and analysis on basic battery capacity

In an effort to reduce uncertainties in the sampling interval study that may be caused by the variability of the individual motes or batteries, eight random motes were selected and tested four times each (tests A–D) under the HP mode using a 1-s DMI. To test the effect of the battery capacity, half of the motes tested were powered with two 2500 mAh AA batteries and the other

half were powered with two 2450 mAh AA batteries. The batteries used in this experiment were all rechargeable nickel-metal hydride (NiMH) AA batteries. In this experiment, motes were placed on workspace benches in close proximity to the base station. A summary of the results of these tests for the eight motes is given in Table 5. Note that these results correspond with those presented in Table 4 (Section 3.2.1).

Table 5. Summary of the mote and battery experiment of four tests on eight motes in HP mode at 1-s transmissions. (Davis et al., 2012c)

	Mote	Test	Start (volts)	End (volts)	Pkts (thou.)	Life (min.)
2500 mAh	1	A	2.55	1.97	102.9	3481
		B	2.79	1.97	156.1	5342
		C	2.72	1.97	147.1	5094
		D	2.85	1.97	143.6	5090
	2	A	2.46	1.87	100.7	3468
		B	2.71	1.65	87.9	3038
		C	2.64	1.87	145.1	5094
		D	2.70	1.87	143.5	5095
	3	A	2.63	1.99	99.7	3488
		B	2.87	1.99	154.4	5383
		C	2.80	1.99	147.1	5230
		D	2.97	1.99	149.9	5381
	4	A	2.47	1.86	98.4	3458
		B	2.69	1.88	150.1	5296
		C	2.59	1.87	135.5	4872
		D	2.76	1.86	142.3	5191
2450 mAh	5	A	2.72	1.95	43.7	1570
		B	2.77	1.91	78.0	2773
		C	2.75	1.93	89.6	3182
		D	2.74	1.94	53.0	1910
	6	A	2.70	1.93	56.5	2043
		B	2.75	1.93	112.9	4056
		C	2.72	1.93	71.7	2576
		D	2.80	1.92	103.3	3744
	7	A	2.75	1.96	84.2	3072
		B	2.79	1.96	85.6	3093
		C	2.76	2.11	63.8	2318
		D	2.85	1.95	103.9	3810
	8	A	2.77	1.91	118.1	4351
		B	2.73	1.91	78.5	2881
		C	2.71	1.91	63.9	2346
		D	2.80	1.88	106.4	3942

The starting and ending battery voltage (based on the two AA batteries) for the experiment are shown in Table 5. The starting voltage is based on the earliest stable battery level recorded in the mote's data packets. This is typically within the first 5–10 transmissions. The ending voltage is the last recorded battery level received by the mote before transmissions ceased. Also included in Table 5 is the number of data packets that were received by the base station from each mote during its operation life. Each data packet represents one sample transmitted and collected from a mote to the base station.

Results, as shown in Table 5, indicate that the batteries from both series had variances in their starting voltages. There is a large amount of variability in the starting voltages of each test. This is due to the rechargeable nature of the batteries. While there are varied differences in the voltage drop for each mote, individual motes have a specific lower-end power requirement for operating as shown by the similar ending voltages in each mote's series of tests. The average battery life for the four motes with 2500 mAh batteries compared to the four motes with the 2450 mAh batteries is 4625 min to 2979 min (a 55% longer battery life in the 2500 mAh batteries). Given these results, it can be concluded that variability in the batteries (e.g. differences in charging time, age of the battery, etc.) is the major component in the variability of the battery life in the motes. This is not to say that the variability in the battery life is caused by the batteries alone. There is evidence of both individual variability (as seen in the ending voltage for each individual mote) and inter-mote variability (as seen in the differences in the ending voltages between groups of motes). However, the variability of the individual mote and inter-mote are less than the variability caused by the batteries.

4.1.2 Experiments and analysis on data message intervals

The battery life for various DMIs in the two power modes were analyzed on 12 motes. The 12 motes that were selected were separated into four groups of three motes. The battery life for a given DMI and power mode (HP or LP) pairing was tested on one of the four groups of motes. The three motes in each group were used to determine statistical parameters of the results. Each DMI and power mode pairing was tested three separate times. To reduce the battery variability, the motes were powered by two Panasonic Industrial (AM3) AA batteries, rated at 2870 mAh. The motes were once again tested using the default radio power level, 0 dBm (1 mW).

In HP mode, the DMIs tested were 10 s, 30 s, 60 s, and 900 s. The same DMIs were tested in LP mode except for 10 s. Due to the battery saving nature of the LP mode's sleep functionality, the 10 s DMI was deemed inappropriate (personal communication with Crossbow technician) and was intentionally excluded from the LP tests.

A summary of the battery tests is given in Table 6. The recommended operating voltage for motes is 3.6–2.7 V (Crossbow, 2007b, Table 6-1), however it has been shown that the mote can continue to collect data down to 2.2 V and transmit messages down to 2.1 V (Musaloiu et al., 2006). Included for each DMI and power mode pairing is the average number of packets received and the average battery life over the total range of battery voltages. Next to each of the average values is the standard deviation based on the nine results (i.e., three groups of three motes per test).

Table 6. Average battery life and transmission packets from nine motes tested in HP and LP modes at various sampling intervals with the standard deviations. (Davis et al., 2012c)

Power mode	Samp. Int.	Life (min.)	Data packets
HP	10 s	5545 (±341)	29991 (±1909)
	30 s	5618 (±345)	10889 (±645)
	60 s	4610 (±217)	4499 (±213)
	900 s	4597 (±139)	308 (±10)
LP	30 s	20502 (±446)	21157 (±3213)
	60 s	20859 (±790)	14035 (±2736)
	900 s	19952 (±1942)	1179 (±112)

Table 6 shows that for both the HP and LP modes, there is a small difference between the total battery life of the various DMIs. In the HP mode results, the 10 s and 30 s DMIs have a slightly longer battery life than the 60 s and 900 s sampling intervals. The battery life for all four of the HP tests are similar to the results obtained in Table 5 for the 2500 mAh batteries. The variability in the battery life results in Table 6 is also similar to the variability in the battery life of a single mote tested in Table 5. Therefore, based on these results it can be concluded that the sampling interval has no significant influence on the battery life, which is contrary to what was expected. These results would suggest that the sensed data transmissions, regardless of sampling DMIs, may not be a major cause of the battery energy consumption in the XMesh network. Further investigation was then conducted, as described in the next section, to determine the possible major causes for the battery life results shown in Table 6.

4.1.3 Experiments and analysis on the impacts of health messages

One possible cause for the battery life results shown in Table 6 would be the XMesh network's health messages. The network's health messages are sent to the base station periodically for

updates regarding each individual mote's neighbor list and its own physical health statistics. The default configuration is to send the health messages every 60 s and 600 s for motes in HP and LP modes, respectively. The health message interval (HMI) is adjustable in the nesC application code. The mote alternates sending the neighbor information and its own statistics health messages at each transmission interval.

The transmission of the health messages could effectively reduce the mote's transmission interval to that of the HMI if the HMI is smaller than the mote's DMI. These health messages sent to the base station can influence the speed in which the mote's battery power is exhausted. Thus, an experiment was conducted in attempt to quantify the effect of different HMIs on the mote battery power.

The same four groups of three motes that were used in the previous battery life measurements were once again used to test the effect of the HMIs on mote battery life. All four groups of motes were programmed to collect data in the LP mode and sample at a 900 s interval. The motes were powered with two Panasonic industrial 1.5V AA batteries. The HMI for each of the four groups of motes was set to 120 s, 300 s, 600 s and 900 s, respectively. The experiment took place indoors in a laboratory environment. In the laboratory, two groups (i.e., six motes per group) were positioned approximately 45 cm away on either side of the centrally located base station. Motes within each group were spaced a few centimeters apart.

The initial results from the health message testing revealed little insight on the effect of these message transmissions on battery life. The average battery life, shown in minutes, over the total range of battery voltages is displayed in Table 7. Similar to the results shown in Table 6, there is again little difference between each group's average battery life. There is, however, a large difference between the average LP battery life between Table 6 and Table 7. The battery

life results in Table 7 are over 200% longer than the LP results in Table 6. There was a single LP test completed that was not included in the results in Table 6 because it was deemed an outlier but had an average battery life of 68000 min. It seems now that the outlying results from the DMI tests (as shown in Table 6) are in accordance with the results presented in Table 7 for the LP motes. This poses a new question concerning the reason for the low LP mote battery life presented in Table 6.

Table 7. Battery life and transmission packets for various health message intervals at a 900-s sampling interval in LP mode over the total battery life. (Davis et al., 2012c)

Health Int.	Life (min.)	Data packets
120 s	67304 (±2174)	3655 (±130)
300 s	67956 (±641)	3672 (±11)
600 s	68974 (±1665)	3826 (±63)
900 s	68112 (±3697)	3818 (±305)

To see how the motes are performing over their battery life, the number of samples taken by each mote in the four groups, separated into voltage bins, is plotted in Figure 21. There is a distinct "double hump" feature noticeable at around the 2.3V and 2.65V bins. The cumulative number of samples collected (as shown in Figure 22) shows that all the motes begin and end collecting approximately the same number of samples over their battery life. The variance present throughout the middle region may be explained by individual mote variability.

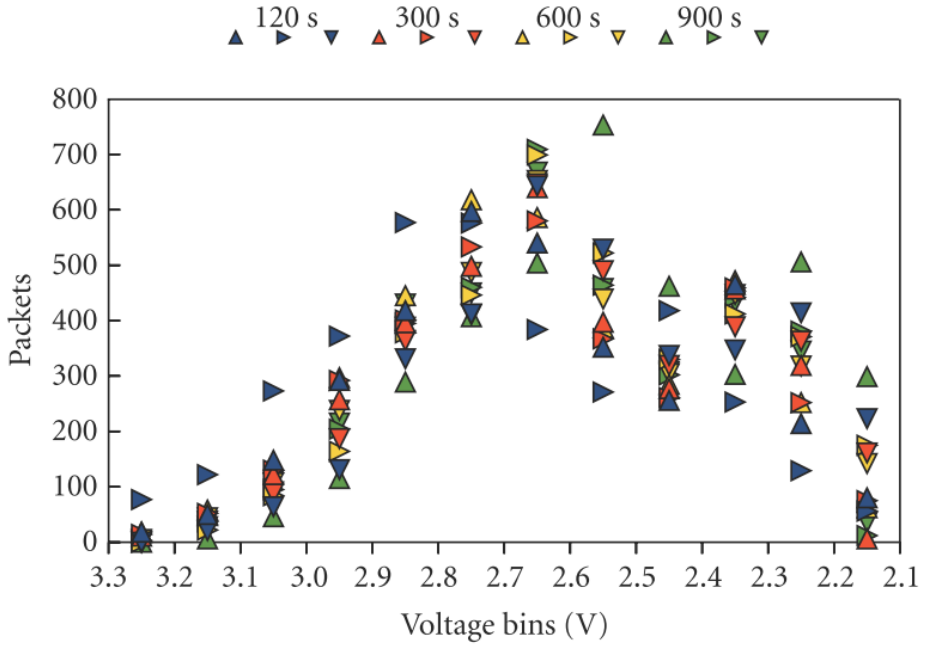

Figure 21. The number of data packets collected for motes in LP mode sampling at 900 s at various health message intervals. The four groups of three motes are plotted, each at their respective health update interval, i.e. 120 s, 300 s, 600 s and 900 s. (Davis et al., 2012c)

The similarities in battery life (e.g., Table 7) and samples collected (e.g., Figure 22) prompted further investigation on the mote's DMI over its battery life to understand these similarities. Figure 23 is an example of a typical mote's behavior over time. The plot of the battery voltage curve is based on the measurements from the data packets received by the base station. The DMIs are taken as the time difference, in minutes, between two successively received data packets. The number in parentheses indicates the count of data packets received at each given interval. It can be seen in Figure 23 that the majority of the packets collected were made at the prescribed interval, that is, 15 min or 900 s. A large number of packets were also collected at twice the programmed interval, that is, 30 min. This indicates that the mote dropped

one data packet between two successful transmissions to the base station. Fewer packets are shown to have been collected at intervals of 45, 60, 75, and 90 min, signifying that the mote dropped 2, 3, 4, and 5 data packets between successful transmissions to the base station. Using the time intervals between successfully received data packets, of the 3744 data packets received, there were an estimated 714 dropped packets (almost 20%). This is a considerably large number with respect to both the loss of data and the wasted power for transmissions. Therefore the dropped packet rates of the wireless motes were more closely examined.

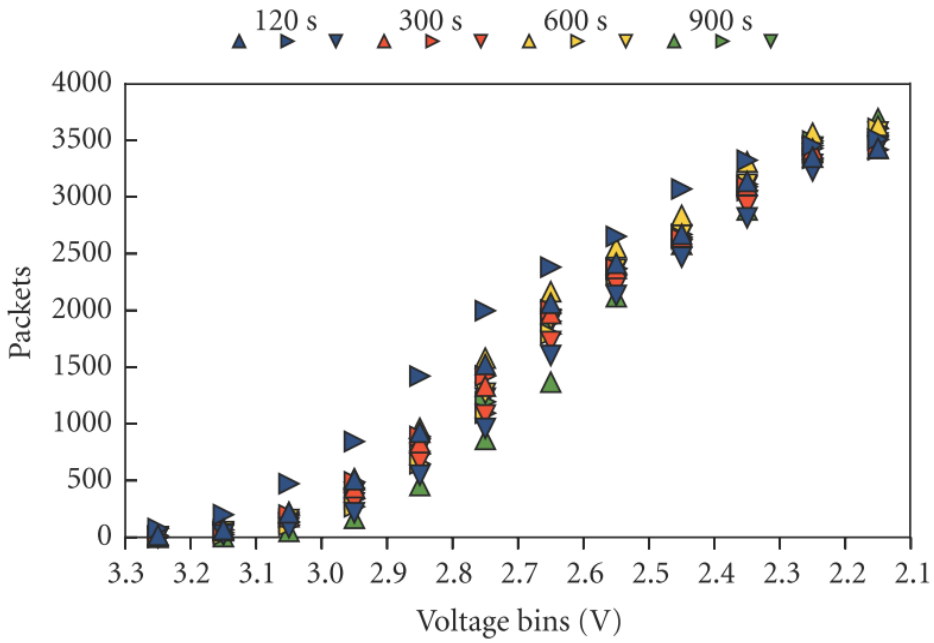

Figure 22. The cumulative number of packets collected over the battery life of motes in LP mode sampling at 900 s at various health message intervals. The four groups of three motes are plotted, each at their respective health update interval, i.e. 120 s, 300 s, 600 s and 900 s. (Davis et al., 2012c)

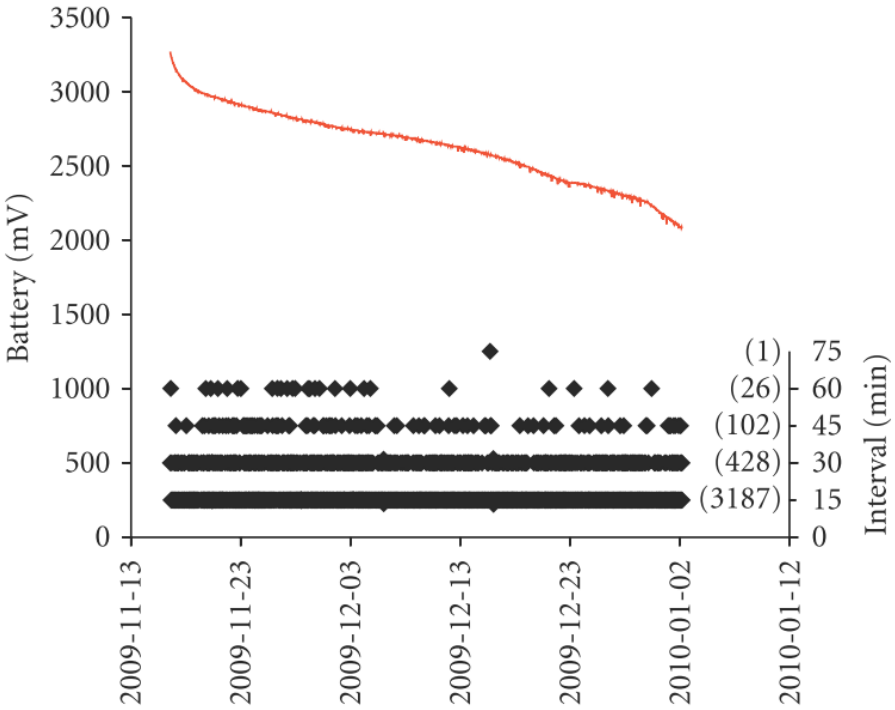

Figure 23. Mote battery voltage (in millivolts) over time (red line) in HP mode with a 15 min DMI and the time interval (in minutes) between received data packets (black diamonds). The number of samples collected at each time interval is indicated in parentheses on the right side of the plot. (Davis et al., 2012)

4.1.4 Experiments and analysis on the impacts of dropped packets

A dropped packet occurs when a mote's data packet is unsuccessfully delivered through the network to the base station. This can occur at any individual node through the multi-hop network. The unsuccessful delivery of a data packet can be due to one of the four failure modes previously defined. In the transmission of a data packet over each consecutive hop, an acknowledgement message is returned to a mote by its parent to signify that the packet was received. If a mote does not receive an acknowledgement, it will attempt to resend the packet up to eight times until either an acknowledgement is received or the maximum number of retries is

reached (Crossbow, 2007a, Section 10.1.4). If the maximum number of retries is reached without successful acknowledgment, the packet is "dropped" and results in lost data.

An analysis of the mote transmission performance is summarized in Table 8. The received packets in Table 8 are for the total transmission life of the motes as given in Table 7. Dropped packets were identified based on the time interval between successful receipts of data packets, as described above. Table 8 breaks down the number of data packets received into the number of packets received on-time (no packet loss), the number of packets received with only one retransmission, the number of packets received after two retransmissions, the number of packets received after three or more retransmissions, and the number of asynchronous packets received.

Asynchronous packets occur when a mote does not receive an acknowledgement from its parent that its packet was received and therefore retransmits the same packet again. In some cases, the link quality from a mote to its parent is better than the link quality from the parent to the mote. This is also known as link quality asymmetry (Teo et al., 2006). Under this circumstance, asynchronous packet delivery can occur and the base station will receive duplicate packets from a mote. The amount of this occurrence can be calculated by identifying two or more duplicate data packets received by the same node within a few seconds of each other.

The number of packets expected, shown in Table 8, is an estimated value based on the sum of received and dropped data packets. The success rate, shown in the last row of Table 8 as a percentage, is the ratio of received to expected data packets. In each case, only about 80% of the packets expected were received. Furthermore, each of the four sets of motes had a similar number of dropped packets. To better understand the nature of packet drops, an investigation into the multi-hop structure of the network was conducted.

Table 8. Transmission analysis of LP motes sampling at 900 s for various health message intervals. (Davis et al, 2012c)

Health Int.	120 s	300 s	600 s	900 s
Packets recv.	3655	3672	3826	3818
0 drop	84.3%	84.0%	84.9%	86.2%
1 drop	12.4%	12.8%	11.8%	11.2%
2 drop	2.4%	2.3%	2.4%	1.9%
2+ drop	0.6%	0.7%	0.7%	0.5%
Async.	0.3%	0.2%	0.2%	0.2%
Packets drop	713	728	735	638
Packets exp.[a]	4367	4399	4561	4456
Success rate	83.7%	83.5%	83.9%	85.7%

[a] Estimates based on summation of data packets received and data packets dropped.

4.1.5 Experiments and analysis on routing usage

This WSN's architecture uses XMesh which features TrueMesh technology. This means that the network is self-healing and self-organizing and each mote acts as both a sensor node and a router for its neighbor's data. The ad-hoc formation of the network is based on link estimates made between node neighbors in order to send data down a path of lowest transmission cost to the base station. As a consequence, certain motes may be exploited as a relay due to their low path cost.

A mote's parent ID is included in the data packet sent to the base station. The parent node is defined as a mote's neighbor with the lowest transmission cost (Crossbow, 2007d, Section 4.2). Thus, by analyzing the parent data, it can be determined if there are any motes being exploited and therefore having their batteries drained at a higher rate.

Table 9 shows the results of this analysis. The number of health packets that each group sent is estimated based on the average battery life of each group and the HMI. The packets forwarded represent the additional transmissions motes make relaying neighbor data (i.e., data and health packets) through the network. To determine the multi-hop forwarding through the network, the distribution of each mote's connection with their neighbors was calculated based on the parenting information collected in the data packets. The same distribution was used to determine the routing of the data and health packets through the network. The number of packets generated is the summation of the estimated number of data packets expected (see Table 8) and the estimated health packets (based on the battery life and HMI). The route-utilization is the percentage of additional forwarding transmissions compared to those generated by the mote (i.e., data and health packets) and is defined as the ratio of packets forwarded to packets generated as shown in Equation 4.1:

Equation 4.1

$$"Route-Utilization" = \frac{Fwd_{data} + Fwd_{health}}{Gen_{data} + Gen_{health}}$$

where: Gen_{data} = Estimated total data packets generated by a mote

Gen_{health} = Estimated total health packets generated by a mote

Fwd_{data} = Estimated number of data packets forwarded by a mote

Fwd_{health} = Estimated number of health packets forwarded by a mote

It can be seen that the route-utilization increases as the HMI increases. The motes with the 900 s HMI were utilized considerably more than those with the smaller HMIs.

Table 9 shows the scaled battery life of each set of motes, as given in Table 7, based on the route-utilization percentage. The results of scaling the mote battery life show the expected trend in increasing battery life with decreasing the number of samples taken.

Table 9. Route-utilization analysis of the LP mote health message interval experiment (from Table 7 and Table 8). (Davis et al, 2012c)

Health Int.	120 s	300 s	600 s	900 s
Data Pkts. Gen.[a]	4367	4399	4561	4456
Health Pkts. Gen.[b]	33652	13591	6897	4541
Data Pkts. Fwd.	1123	2790	2454	12714
Health Pkts. Fwd.	4575	11745	10191	52682
Route-util.[c]	15.0%	80.0%	110%	727%
Life (min.)[d]	77390	122859	145088	563209

[a] Expected data packets based on the received and estimated dropped packets (see Table 8).
[b] Estimates based on the battery life and health message interval.
[c] Route-utilization is the ratio of forwarded data and health packets to generated data and health packets.
[d] Total battery life, see Table 7, scaled based on the route-utilization percentage.

The MICAz mote is expected to have a battery life up to one year (Crossbow, 2007c). This corresponds to the estimated battery life of the 900 s HMI results. Given the results in Table 6 and Table 7 for LP motes, the scaling may be an exaggerated battery life adjustment. It is more likely that the increase in battery life compared to the default 600 s would be closer to 14% (logarithmic trend) or 29% (linear trend). Regardless, this test shows that decreasing the HMI from the default 600 s to either 300 s or 120 s will affect approximately 15–46% of a mote's battery life.

The analysis of scaling the battery life according to the mote's route-utilization can be applied to the results in Table 6. The adjusted HP results in Table 10 show no change in the battery life trend from the original results in Table 6. Given that the HMI for motes in HP mode is 60 s, it is expected that the data sampling rates of 60 s and 900 s HP results would be similar, that is, motes sending data every 900 s are also sending health packets every 60 s. The HMI, however, does not explain why with 10 s and 30 s sampling rates the HP motes' battery lives are about 25% longer than those with the 60 s and 900 s sampling rates. A more detailed analysis of the power consumption of HP motes of XMesh-based WSNs is necessary to understand this difference which is beyond the scope of this work. The adjusted battery life for the LP results in Table 10 show the expected trend in battery life with increasing the transmission interval.

Table 10. Route-utilization analysis of the HP and LP mote sampling interval experiment (from Table 6).

(Davis et al., 2012c)

Sampling Interval	HP 10 s	HP 30 s	HP 60 s	HP 900 s	LP 30 s	LP 60 s	LP 900 s
Health Packets[a]	5545	5618	4610	4597	2050	2086	1995
Packets dropped	156 (±77)	12 (±3)	2 (±2)	0 (±0)	5655 (±4672)	1229 (±472)	112 (±41)
Packets Fwd.[b]	1855 (±4099)	64 (±157)	1 (±1)	0 (±0)	11525 (±3962)	12294 (±12498)	9531 (±18069)
Packets Gen.[c]	35692	16519	9111	4905	28862	17350	3286
Route-util.	4.6%	0.5%	0.0%	0.0%	42.2%	66.5%	289.6%
Success rate	99.5%	99.9%	99.9%	99.9%	78.9%	91.9%	91.3%
Life (min.)[d]	5800	5646	4610	4597	29154	34730	77732

[a] Estimates based on battery life and the default health message intervals, i.e., 60 s for HP motes and 600 s for LP motes.

[b] Forwarded packets are estimates of both data and health packets, taking into account the multi-hop network functionality, based on the statistical distribution of mote routing.

[c] Packets generated is equal to the sum of a mote's total data and health packets.

[d] Total battery life, see Table 6, scaled based on the route-utilization percentage.

The route-utilization shown in Table 10 indicates that the LP motes with the largest data transmission interval, that is, 900 s, suffer the highest routing relays. This is similar to what was seen in the results of Table 9. For the HP motes, however, it can be seen that the smaller transmission interval (10 s) leads to the highest routing relays. From both tests, the HP motes have a consistently higher rate of successful transmissions (above 90%) while the LP motes have a generally lower rate of successful transmissions over a varied range (78–92%).

While increasing the HMI can save some battery life, it reduces the amount of information collected regarding node link quality and path cost through the network. For purposes of monitoring the network mesh, the HMI may be more important than the battery savings. However, if network monitoring is not being considered, the health messages can be disabled completely to maximize battery savings.

4.2 FIELD STUDY

In addition to the laboratory tests, a pilot test bed consisting of 11 nodes was deployed in a residential backyard (see Figure 24) located in western Pennsylvania (40.5436° N, 80.0638° W). The nodes range 6–60 m away from one another and are located mainly along the perimeter of the yard which is an approximately 60 m by 30 m rectangle. The wireless motes, data acquisitions boards, and battery packs were all housed inside polycarbonate high-impact enclosures (Bud Industries Inc.; part no. PN-1337). An external antenna (Pulse Electronics; part no. W1038) was attached to the outside of each enclosure. The enclosures were mounted on wooden stakes, placing the antenna approximately 0.3 m above the ground.

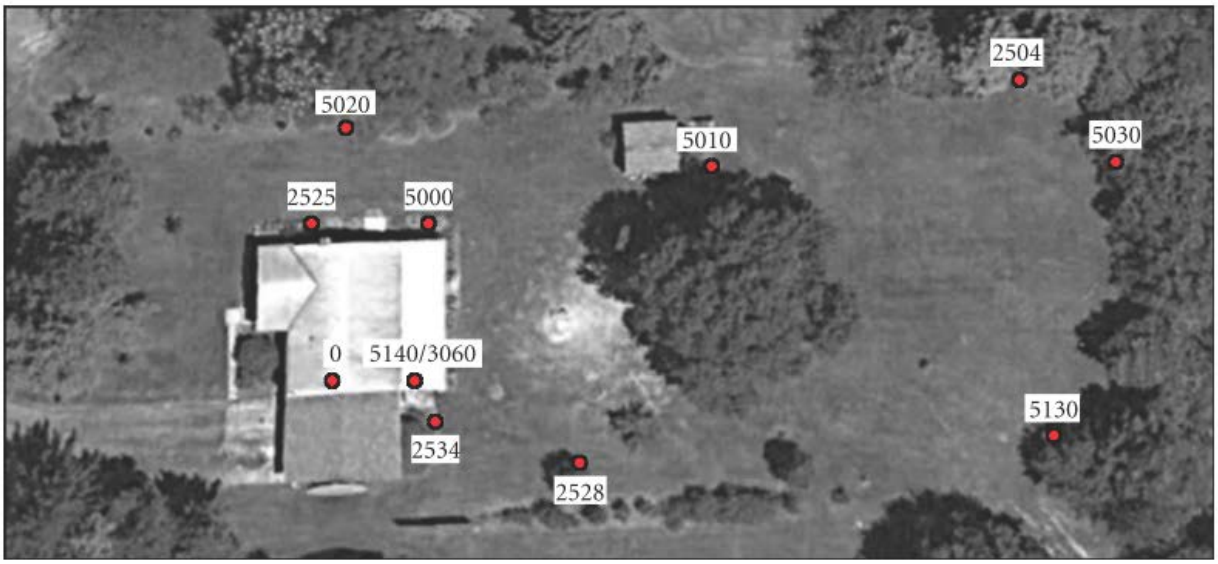

Figure 24. Residential backyard WSN pilot test bed node locations during the late summer and autumn of 2009. (Davis et al., 2012c)

The network collected data starting from the summer of 2009 until the summer of 2010. All nodes were programmed in LP mode with the default radio transmission power (0 dBm) and health message interval (600 s). Motes were powered by two sets of two AA rechargeable batteries connected in parallel. The LED indicator lights on the mote were turned off due to their power consumption (6–8 mA) on the mote's batteries (Crossbow, 2007c; Krämer and Geraldy, 2006). From 22 July to 20 October 2009, all nodes sampled data at a 15 min interval. An analysis of the network's battery life and behavior was completed during this time. Figure 25 shows the received packets from each of the 11 nodes over the three-month period.

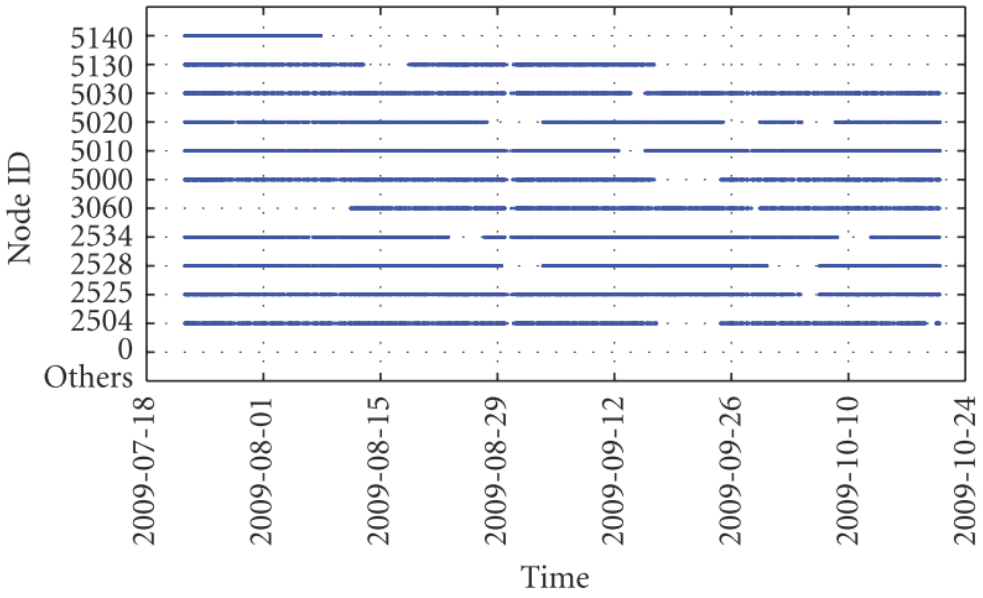

Figure 25. Data packets received by the base station for each node over the three-month monitoring period. (Davis et al., 2012c)

Table 11 shows the battery life and packet analysis for each of the 11 nodes. The start and end dates correspond to the time that the transmission of data packets began and ended (see Figure 25). It should be noted that mote 5140 was replaced by mote 3060 due to technical difficulties. Data was excluded from the analysis during periods where the motes did not complete a full battery cycle, that is, motes whose batteries were replaced before 20 October 2009, but were not fully depleted until after this time.

Table 11. Battery life and transmission statistics of the 11 node pilot test bed network. (Davis et al., 2012c)

Node ID	Start date	End date	Life (min)	Pkts. rcvd.	Pkts. drop	Packets fwd.	Parent changes	Success
2504	07/22	09/17	81294.9	2152	3136	1080	1149	40.7%
	09/24	10/19	35322.6	895	1399	321	370	39.0%
2525	07/22	10/04	106103.0	2889	4018	371	495	41.8%
2528	07/22	08/29	54618.2	1432	2113	2029	686	40.4%
	09/03	09/30	38652.5	1121	1382	1946	477	44.8%
	10/06	10/20	20038.7	529	768	835	235	40.8%
2534	07/22	08/23	45425.7	1136	1802	1931	518	38.7%
	08/27	10/08	60846.7	1646	2287	2753	634	41.9%
3060	08/11	09/28	69075.5	1934	2567	3492	741	43.0%
5000	07/22	09/16	80909.1	2194	3116	2636	783	41.3%
	09/24	10/20	37022.7	945	1491	601	318	38.8%
5010	07/22	09/12	74773.6	1944	2947	1517	1007	39.7%
5020	07/22	08/27	52074.3	1312	2110	689	570	38.3%
	09/03	09/25	31147.3	923	1122	625	277	45.1%
5030	07/22	09/13	76867.8	2014	3013	663	938	40.1%
5130	07/22	08/12	30834.1	790	1281	363	410	38.1%
	08/18	09/16	42237.0	1214	1611	380	581	43.0%
5140	07/22	08/07	23470.4	752	807	674	213	48.2%

It can be seen from Table 11 that the outlier nodes (e.g., 2525, 2504, and 5030) have less forwarded packets and longer battery lives. While the outlier node 5130 did not forward many packets, its battery life was not as long as the other outlier nodes. This could be due to its isolated position (i.e., not having a good connection to its neighbors). Nodes with the closest contact with the base station (e.g., 3060, 2534 and 2528) forwarded the most packets with battery lives, on average, less than the outlier nodes.

The success rate is approximately the same, around 40%, for all the nodes tested in the field. The high dropped packet rate may be attributable to the motes' close proximity to the ground. From the laboratory experiments the dropped packet rate for LP motes was found to be 10-20% compared to the 60% dropped packet rate in the field. The network's stability can be analyzed by examining the number of times a mote's parent changes. The high number of parent changes for the motes suggests that the link quality between the motes and their neighbors was low and fluctuated around the threshold for new parent selection. A large number of parent changes is representative of low network stability. This may have also attributed to the high dropped packet rate.

The pilot study results are comparable to the laboratory experiments for motes programmed in LP mode. In Table 7 and Table 9, the results for motes with the smallest health message interval (120 s) are close to the pilot study results for nodes with higher traffic (e.g., more forwarded packets) such as motes 2504 and 5010. The comparable battery life (approximately 70000 min) and forwarded packets (approximately 1400 packets) show that the laboratory experiments were completed in a relatively high traffic condition compared to other motes in the field. It should be noted that under the same HMI, the scaled battery life from the laboratory experiment (Table 9) is about double the results in the field. The original battery life

(Table 7) is similar to what was seen from the motes in the test bed. This would suggest that route utilization affects approximately 50% of the battery life. While there does not appear to be any correlation between mote location and the transmission success rate, battery life does show some dependence on the network's topology.

An analysis on the route utilization in this pilot study was then performed. Based on the motes' parent IDs included in the data packets (see Figure 26), the link selection probability distribution for a packet to be forwarded by a mote through any of the mote's neighbors was calculated. Based on the conditional probabilities, a graph of the network topology was created where the vertices represent motes and each edge corresponds to a communication link associated with its selection probability. Considering the possibility of asymmetric links, the topology graph is a directed graph and its adjacency matrix is presented in Table 12. In this graph, the data for mote 5140 was included in the data for mote 3060.

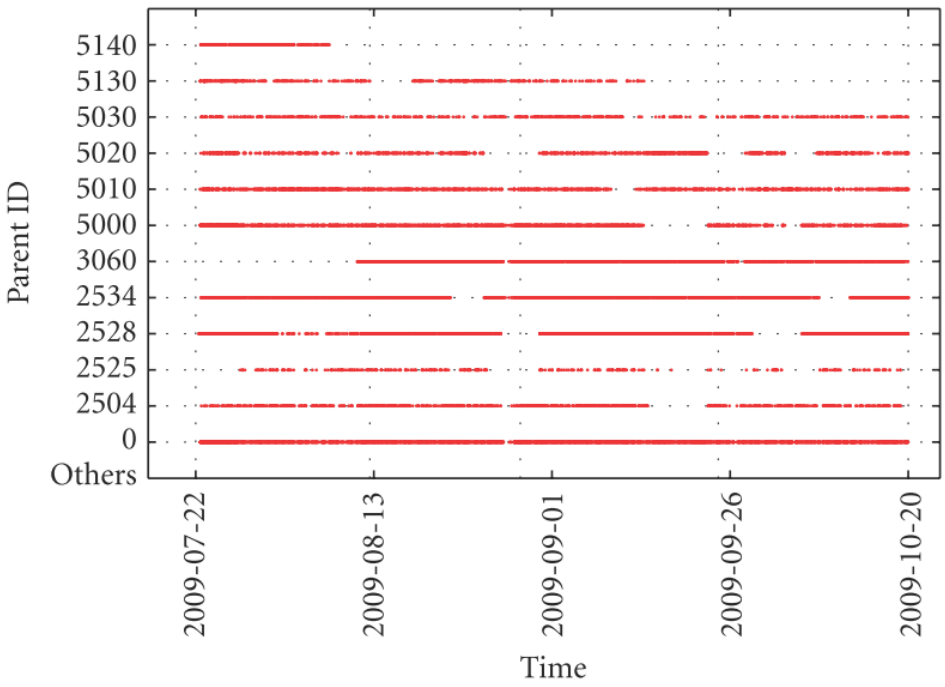

Figure 26. Parent ID for the received data packets for each node over the three-month monitoring period. (Davis et al., 2012c)

For this graph, each individual mote's link selection probability distribution is calculated based on the parent IDs included in its generated data packets, which represents the behavior of the mote's first hop of its route towards the base station. In order to understand the overall routing behavior of the multi-hop network, it is reasonable to assume that the forwarded packets at each relay mote will follow the same link selection distribution of the generated packets at this mote for there is no parent ID information for relayed packets. Thus, in the following analysis, the link selection in routing for both originally generated packets and relayed packets at each individual mote follows the same link selection distribution.

Table 12. Pilot study adjacency matrix for the network topology graph. (Davis et al., 2012c)

	0	2504	2525	2528	2534	3060	5000	5010	5020	5030	5130
0	0	0	0	0	0	0	0	0	0	0	0
2504	0	0	0	0.162	0.242	0.303	0	0.117	0.013	0.104	0.059
2525	0	0	0	0	0	0	0.631	0.002	0.367	0	0
2528	0.274	0.030	0	0	0.261	0.253	0.046	0.076	0.007	0.029	0.025
2534	0.478	0.025	0	0.170	0	0.221	0.008	0.046	0.019	0.022	0.011
3060	0.528	0.032	0	0.130	0.222	0	0.026	0.042	0.007	0.007	0.006
5000	0	0	0.038	0.348	0.134	0.188	0	0.191	0.094	0.001	0.007
5010	0	0.074	0.001	0.214	0.271	0.318	0.038	0	0.012	0.045	0.028
5020	0	0.001	0.107	0.129	0.168	0.310	0.230	0.041	0	0.003	0.011
5030	0	0.150	0	0.210	0.329	0.079	0.001	0.129	0.017	0	0.084
5130	0.005	0.185	0	0.228	0.168	0.135	0.021	0.119	0.015	0.124	0

The results of route utilization for each mote are presented in Table 13. The average route utilization for the field in this test is 105%. Given that the network is configured for a DMI of 900 s with a default HMI of 600 s, these results are comparable to the corresponding lab test with the same DMI and HMI (see Table 5). These results can be justified based on the higher traffic loads and route utilization in some specific motes. The analysis for the route utilization in each individual mote shows that motes 2528, 2534, and 3060 have the highest route utilization in the network followed by motes 5000 and 5020. The group with the smaller route utilization consists of motes 5010, 2504, 5130, 5030 and 2525. These three groups of motes are highlighted in Figure 27 according to their route utilization status.

Table 13. Route-utilization percentage for each mote and network average. (Davis et al., 2012c)

Node ID	Generated packets	Forwarded packets	Route utilization
2504	1524	742	48.7%
2525	2889	181	6.3%
2528	1027	2417	235%
2534	1391	2689	193%
3060	1343	2584	192%
5000	1570	2270	144%
5010	1944	1106	56.8%
5020	1118	1343	120%
5030	2014	444	22%
5130	1002	385	38.4%
Network avg.	1582	1416	105%

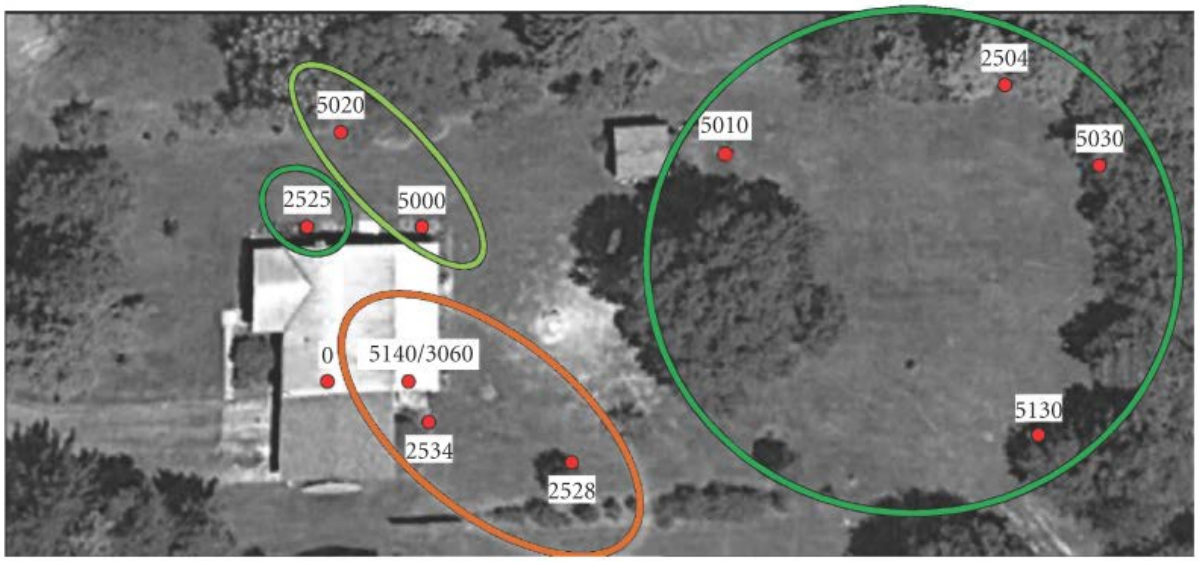

Figure 27. Locations of three mote groups in the pilot test bed according to their route utilization status. Groups highlighted in green correspond to the smaller route utilization and orange corresponds to the highest. (Davis et al., 2012c)

Lastly, aiming to provide a major insight in the network dynamics, the highest probability routing links between nodes were selected based on the network topology graph (see Table 12). Based on these probability links, the most possible paths through the network for each node are presented in Table 14. From these results, it can be confirmed how motes with higher route utilization are used to forward packets from other motes. Since their main route is directly connected to the base station (mote 0), these motes are associated with a much higher route selection probability.

Table 14. Routes with highest probability for the different motes in the pilot test bed. (Davis et al., 2012c)

Node ID	Route 1	Route 2
2504	3060–0 (16%)	2534–0 (11.5%)
2525	5000–2528–0 (6%)	5000–2528–2534–0 (2.7%)
2528	0 (27.4%)	2534–0 (12.4%)
2534	0 (47.8%)	—
3060	0 (52.8%)	—
5000	2528–0 (9.5%)	2528–2534–0 (4.3%)
5010	3060–0 (16.7%)	2534–0 (12.9%)
5020	3060–0 (16.3%)	5000–2528–0 (2.1%)
5030	2534–0 (15.6%)	—
5130	2528–0 (6.2%)	2528–2534–0 (2.8%)

While the results from the field study show comparable results with experiments done in the laboratory setting, there is still a considerable amount of missing information with respect to the effect of transmissions on the battery life of motes. The information transmitted via the health messages, which include a large portion of the missing information (e.g. retransmissions, dropped packets, forwarded packets, link quality, neighbor nodes, etc.), is not stored by default in the gateway; therefore further investigation is not possible.

One lesson learned from this pilot study is that the gateway in a WSN plays a critical role in the WSN reliability and robustness, which could be easily overlooked at the beginning of WSN deployment. The gateway that was used in both the laboratory and field studies presented here was the Stargate Netbridge, a Linksys NSLU2 device specially modified by Crossbow. One attractive feature of the Stargate Netbridge is that it has a local web-based monitoring and management interface called MoteExplorer. MoteExplorer allows easy visualization of the network status, the current mesh topology and a live stream of mote data collected by the base station, which made the Stargate Netbridge a favorable choice for the WSN gateway in deployment and operation. On the other hand, the Stargate Netbridge runs a Debian Linux operating system from an attached 4 GB flash drive. The deprecated version of Debian Linux came with limited and mostly outdated software. Due to the limited resources, in computational speed and memory storage, and an unconventional architecture (ARM), updates and upgrades to the device were not possible. With no direct access to the computer (monitor, keyboard and mouse are not supported), all operations had to be performed through an SSH connection (for more information regarding the limitations on the Stargate Netbridge, see Stanjo et al., 2008). It has been discovered, in the pilot study, that the Stargate Netbridge has severe reliability issues, and its strong limitations in hardware and extensibility make it an inconvenient solution for a practical WSN deployment. Following several unrecoverable crashes to the Stargate Netbridge, it was decided to move the gateway platform to a Linux x86 computer running Ubuntu Linux 10.04 operating system. While the new gateway platform would allow for faster, easier, and more reliable network operations, it raises an important challenge due to its lack of the convenient WSN management tool, MoteExplorer. In order to facilitate the change in the gateway platform to fundamentally improve the WSN system's robustness and reliability, a

general integrated network and data management solution (Navarro et al., 2011) was applied to the WSN test bed in this study. The following section describes the network and data management system and its application to the deployed outdoor WSN pilot test bed. An example of its operation is also shown.

4.3 NETWORK AND DATA MANAGEMENT SYSTEM

A web-based integrated network and data management system called INDAMS, presented in Navarro et al. (2011), has been applied to the WSN pilot test bed to facilitate the management and monitoring needs of real-world WSN deployments. This management system employs the following fundamental features: 1) separates the WSN management functions from WSN applications, 2) utilizes an accessible web-based user interface for management functionalities, and 3) systematically supports multiple WSNs from multiple platforms and technologies. Multiple users can not only independently, in real-time, remotely access the WSN test bed operations via this management system, but they can also retrieve and monitor sensing data and network management information with unified web-based management tools that may not be available at local gateway system(s). This eliminates the complexity involved with users dealing with the complex commands and configurations of the WSN gateway (e.g., Memsic's XServe). Instead, this management system enables users to access the important and critical WSN management data for conducting network analysis. Integrated with the new gateway platform in this study, INDAMS not only successfully addresses the issue of regular network management needs in the operations of the WSN test bed, but also collects and saves the network health statistics that were previously discarded, most likely due to storage requirements of the Stargate

Netbridge, such that comprehensive power and routing analysis can be conducted. For a detailed description of the general INDAMS development, see Navarro et al. (2011).

The developed management system is deployed and tested in the pilot test bed described in Section 4.2 during the spring and summer of 2011. The management system offers a geographical and topological monitoring feature. This allows users to see a map of the mote locations in the WSN. Each node on the map shows the connections to each of the mote's neighbors and highlights the mote's parent. The neighbor and parent information is received from the data and health packets transmitted by each mote. The map is updated continuously as the information is being received by the server. This provides a useful tool for a fast and updated view of the network's state and routing for each of the motes.

The management system also includes a live feed of the information received from the WSN gateway. All packets received from the motes (i.e., data and health) and base station (i.e., heartbeats) are displayed. Filters were implemented to specify the types of packets (e.g., data or health) or packet fields (e.g., specific node IDs) to be displayed.

4.4 CONCLUSIONS

This chapter thoroughly examines the power efficiency and battery savings, as well as some key characteristics of wireless mote transmissions in both laboratory and field settings, based on a pilot study and following a quantitative approach. The major contributions are summarized as follows:

(1) This study shows that motes in LP mode drop between 10–20% of their sampled measurements (see Table 8 and Table 10). This is one of the drawbacks of operating the Memsic

MICAz motes in LP mode, as opposed to the HP mode. In contrast, in the HP mode, mote measurements at 15 min intervals have less than 1% packet loss occurrence. While the LP mode decreases the success rate of data transmission compared to using the HP mode, it supplies around four times the battery life in comparison with that in the HP mode (see Table 6).

(2) It can be seen that by adjusting the health message interval, an additional 15–46% in battery life can be gained (see Table 9). In addition, the multi-hop capability of the network can exploit certain motes which have relatively low path cost to the base station. This leads to motes serving as relays to have a shorter battery life than those that do not. This investigation shows that route-utilization can result in over 50% reduction in the battery life in motes in LP mode in the laboratory setting (see Table 10). Little to no route utilization was found in HP motes.

(3) The field study shows that the battery life of the LP motes was comparable to those tested in the lab. Route utilization was identified in motes with good link quality with the base station. The battery life of these motes was variable but generally less than motes located on the outer edge of the network which forwarded less packets. The dropped packet rate was significantly higher in the field than that in the laboratory setting, about 60% compared to 10–20%, respectively. The higher rate of dropped packets may be due to the low clearance of the motes above the ground and the stability of the network, which was found to be questionable due to the high occurrence of parent changes.

(4) This study reveals the practical vulnerability of the original Stargate Netbridge gateway in the WSN test bed deployment, when the deployed WSN moderately scales up. To address this critical issue, a new gateway platform integrated with a general web-based WSN management solution is presented. This solution not only succeeds in replacing the Stargate Netbridge gateway for reliable WSN deployment and operations, but also enables the network

management data collection of all the important network health and neighbor statistics information. This was not previously possible due to the limitations of the Stargate Netbridge gateway. Indeed, this has allowed for a more comprehensive study of the power consumption of WSN deployment, as discussed in Chapter 5.0 and Navarro et al. (2013). The network management system also further improves upon the functionality of the original web-based monitor provided by Crossbow on their Stargate Netbridge. This pilot test bed study demonstrates the merits of the general integrated WSN network and data management system INDAMS presented in Navarro et al. (2011).

In conclusion, this work provides new quantitative insights into the power characteristics of a practical WSN using current, popular, and commercially available WSN technology, which would be useful in energy-efficient WSNs for environmental monitoring in real-world practice. Knowledge gained from the experience of installing and studying this pilot test bed was applied when deploying and maintaining the ASWP test bed, which is discussed in Chapter 5.0

5.0 EVALUATION OF A WIRELESS SENSOR NETWORK FOR ENVIRONMENTAL MONITORING

Portions of Section 5.1 are taken from the paper "ASWP: A long-term WSN deployment for environmental monitoring" which was submitted and is under review for the 12th ACM/IEEE Conference on Information Processing in Sensor Networks (IPSN '13). References to this conference paper are cited as Navarro et al. (2013).

As noted in Section 1.2, previous studies have used WSN technology in a variety of environmental fields. Despite the growing potential of this technology, problems still exist with the implementation, deployment, and maintenance of these systems. Also, WSNs present a steep learning curve in usability (Trubilowicz et al., 2009). An important step in improving this technology is to understand how it operates at its current state. Early investigations were presented on a pilot WSN deployment in Section 4.2.

This chapter evaluates a two-year deployment period of a WSN for environmental monitoring at the ASWP (Audubon Society of Western Pennsylvania) test bed. The objective is to determine if using WSNs for real-time environmental monitoring is a practical option, considering the cost and complications of maintaining the network. The goal is to have a low-maintenance network that can be affordably scaled to meet any research need.

This chapter also presents an overview of subjective experiences and objective data analysis for the purpose of presenting a well-rounded evaluation of the ASWP test bed. Included

in this analysis is data collected from the network (i.e., sensor and health data). It should be noted that health data is not available before the installation of the Linux gateway (i.e., pre-July 2011). In addition to examining the network data, a statistically-based network data transmission path simulation was performed. The information regarding the path of data transmissions through the network cannot be directly retrieved from either the sensor data or health data, thus a simulation was conducted. Results from these analyses are presented in Section 5.2.

5.1 TEST BED DESCRIPTION

This project's WSN is deployed at the Audubon Society of Western Pennsylvania's (ASWP) Beechwood Farms Nature Reserve (BFNR), which is located in Fox Chapel in northern Allegheny County, Pennsylvania (see Figure 28). The BFNR is 134 acres of protected land, which is owned by the Western Pennsylvania Conservancy. The reserve facilitated the power and internet needs for the project's network gateway and it had an accessible wooded area where nodes could be installed. The hiking trails of the reserve are closed from dusk to dawn and attract mostly nature enthusiasts, so the equipment was relatively well protected from human interference.

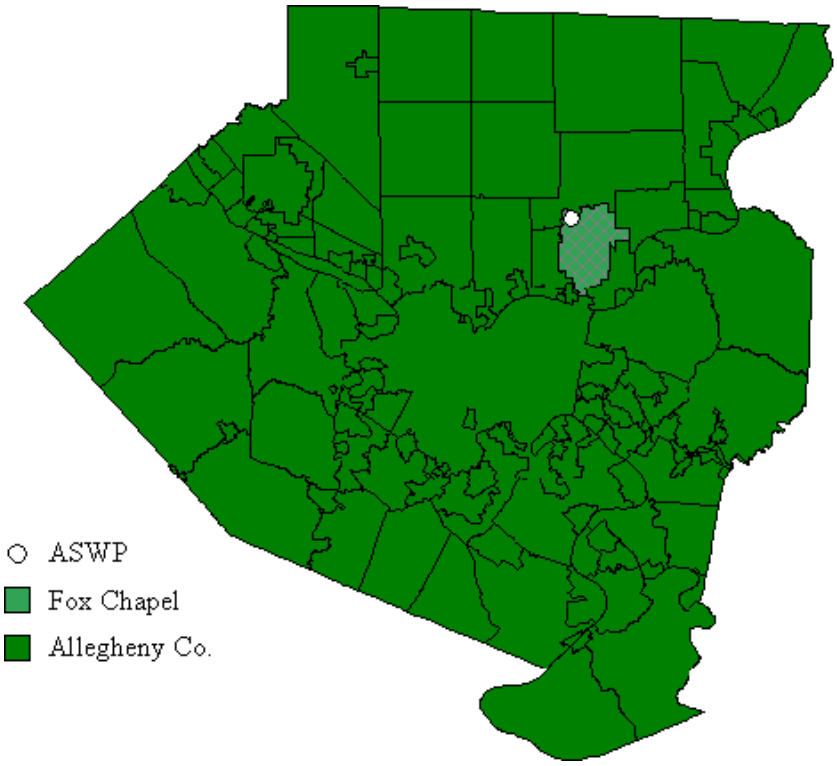

Figure 28. The location of the ASWP study site within Fox Chapel, Allegheny County, Pennsylvania**.**

5.1.1 Test bed requirements

There were two requirements for this study. First, a WSN gateway (i.e., computer for data collection and network management) and base station (i.e., sink node) needed to be placed in a location where power and internet were accessible. Second, the ability to place wireless nodes throughout the study region was necessary. Provisions agreed upon before the first phase of deployment began were that power and internet access would be provided within the Nature Center for the WSN gateway and the visibility of the motes throughout the study region would be kept to a minimum.

5.1.1.1 Base station and gateway

The first requirement for this network was the placement of a base station and gateway computer within the Nature Center. The original gateway was a Stargate Netbridge NSLU2 (i.e., Sluggo) device. Its small size and low power requirements were attractive for this project. A corner of an office desk was made available for the Sluggo unit and the base station was located in the office window via a USB cable. Wired internet and outlet power was provided.

During the operational time of the Sluggo gateway, remote connectivity was made using a simple web service manager. The Apache AXIS2/C web service was chosen because it was compatible with the Sluggo gateway's ARM architecture and it had low computation power/memory requirements. The Sluggo gateway operated the server, while a computer on the University grounds, running the client software, requested daily data downloads from the server.

The Sluggo gateway remained in operation from April 2010 until July 2011 when it was replaced with a new gateway platform (i.e., the Linux gateway). The need to create a new gateway platform stemmed from the shortcomings of the Stargate Netbridge gateway (Stanjo et al., 2008; Davis et al., 2012c, also Section 4.2). The Sluggo gateway was replaced by the Linux gateway, described in Section 4.2, which featured the web-based WSN management system, INDAMS, as described in Section 4.3. It is important to note that the pilot test bed described in Section 4.2 overlaps with the deployment period of the ASWP test bed and that INDAMS was not fully tested until July 2011 when the new Linux gateway was incorporated into this study.

The new gateway was built using an older untested and unqualified version of XServe (the gateway middleware provided by Crossbow). This version of XServe is compatible with open source Linux operating systems running x86 computer architecture. This not only made the gateway more reliable, but provided the convenience of being able to run the gateway on a

laptop or PC. XServe was successfully integrated with a compatible postgreSQL database for collecting and storing WSN data. The Linux gateway was equipped with TinyOS, which provided the ability to write, compile, and program wireless mote applications. Crossbow's XMesh library files were also added to the gateway. The Linux gateway was installed on an AOpen BB10 mini-PC running an Ubuntu 10.04 LTS operating system. The AOpen mini-PC was chosen in order to maintain a small form factor (19×7×20 cm), while providing the performance and power of a desktop computer.

A suite of software programs were developed for the Linux gateway for the purposes of backing up the WSN data, monitoring battery status of motes, tracking IP addresses of the gateway, recording the internal CPU temperature of the gateway, and blocking unwanted remote access to the gateway. Following the move to the new Linux gateway, Dropbox, a file synchronization and backup service, replaced the AXIS2/C. Dropbox provided a more sophisticated and user-friendly alternative for backing up and providing remote connectivity to the WSN data.

5.1.1.2 Node enclosures

To accommodate the patrons of the reserve, the deployment of the network needed to meet certain parameters to keep with the natural aesthetic. In order to not take away the experience of the landscape, the nodes were placed in camouflaged enclosures. For relay nodes, the enclosures used are similar to those used in the pilot field study described in Section 4.2, i.e. polycarbonate waterproof enclosures (Bud Industries, part no. PN-1337). All relays were originally mounted on 3/4" PVC pipe to ensure that the antenna height was adequate for transmitting over distances (a concern presented during the pilot study in Chapter 4.0). To hide the appearance of the enclosures and pipes, they were spray painted with camouflage paint.

In locations that received heavy foot traffic, nodes were arranged in trees, and other areas where understory growth was present, to make them less visible. This limited the locations where nodes could be placed. To accommodate this, wire hangers were attached to node enclosures such that they could be hung from tree branches. This method of placing motes limited the possibly of human tampering. It also provided antenna height over 2 m (over twice that reachable using the PVC pipes) which provided a longer transmission range for the nodes in the network, since antenna height is proportional to transmission distance (Teo et al., 2006). Section 6.3.2 discusses the node enclosures used for the sap flow sensors.

Figure 29. (Left) Camouflaged node enclosure on PVC pipe. (Right) Tree-hanging camouflaged node enclosure.

5.1.2 Hardware description

At the time this project started, the main competitors for WSN technology were Crossbow (MICAz) and MoteIV (Tmote Sky). Crossbow was selected because they had reliable and effective technical support, even though their products were slightly more expensive than their competitors. The MICAz MPR2400 radio and processor board was used for the base station and wireless motes.

Figure 30. (a) Crossbow MICAz and (b) MoteIV Tmote Sky wireless modules.

For data collection, the MDA300 data acquisition boards were purchased. The onboard temperature and humidity sensors on the MDA300 allow for monitoring a mote's environment (checking for overheating and wet conditions), so they were used on all the wireless motes (both sensor and relay nodes). In addition, the MDA300 data acquisition boards provide seven analog sensor inputs. The MDA300 analog inputs can measure sensor readings ranging from 0 to 2.5 V at a resolution of 0.6 mV (Crossbow, 2007e). All of the analog inputs on relay nodes and any unused inputs on sensor nodes were grounded (measurements were zeroed) to avoid confusion (noise looking like data) and reduce storage requirements (less data to save).

The MICAz motes were sold with 1/4 whip tail antennas attached directly to the mote using an MMCX connector. To improve transmission performance, an external antenna, similar to the one used in the pilot study in Section 4.2 (i.e., Pulse Electronics Corp. part no. W1038), was purchased. The 6" antennas, which provide a 4.9 dBi gain, have an RP-SMA style connector. In order to attach the antennas to the motes, an auxiliary antenna wire was needed to bridge the MMCX and RP-SMA connectors. This was solved using a Laird Technologies Wireless M2M 5" coax MMCX to RP-SMA cable.

5.1.3 Software description

The ASWP test bed uses the same XMesh networking protocol that was used in the pilot test bed discussed in Chapter 4.0 . XMesh provides TrueMesh, a self-healing and self-organizing network service (Crossbow, 2007c). The XMesh application code for the wireless motes is compiled for specific mote architectures (e.g., MICAz or Iris) and data acquisition boards (e.g., MDA300). There are three power modes offered by XMesh: high power (HP), low power (LP) and extended low power (ELP). The power mode assigned to each mote is defined during the application compilation. For consideration of long-term monitoring applications, the LP mode was used in this study. The LP mode turns off a mote's unessential electronics while the mote is idle and is still able to forward messages from neighboring motes.

XMesh uses three types of data packets: sensor data, health statistics, and route updates. The first two types of packets are transmitted through the network and stored at the gateway. The third type of packet (i.e., route update) is only transmitted to nodes within radio range of the motes (i.e., it is not stored at the gateway). Sensor data packets store related sensor data based

on the data acquisition board as well as the origin node's ID, the parent node's ID, and the mote's battery status.

There are two kinds of health statistics packets: node health and neighbor health. These packets are transmitted alternatively by motes, one after the other, at the set health message interval (HMI). This makes the effective transmission interval for each health packet type twice the defined HMI. Node health packets also report the origin node's ID, parent node ID, and battery status. In addition, node health packets transmit accumulating counters which track the number of generated health packets and node packets, as well as the number of packets forwarded, dropped, and retransmitted. Link quality estimates and the path cost of the parent node are also included. Neighbor health messages report up to five neighboring nodes for each mote and include link quality estimates and the path cost associated with each neighbor.

The transmission frequency of sensor data, health data, and routing messages are defined in the application code. The routing messages are defined by the route update interval (RUI), which is also definable during the compilation of the XMesh applications. The RUI specifies the time between route update messages, which are mainly responsible for mesh formation time. This relates to the time it takes for the network to form a topological mesh in which data is passed through the network from the farthest away node to the base station. For a multi-hop ad-hoc network, such as the one in this study, the mesh formation time is typically estimated by the RUI multiplied by the average number of hops in the network.

XMesh's multi-hop routing uses the Minimum Transmission (MT) cost metric which is similar to the Estimated Transmission Number (ETX) method used in other protocols (Crossbow, 2007a). The MT cost is calculated in the network based on two cost values: node cost and link cost. In order for messages to be routed to the base station, the base station node

cost is set to zero. Link costs are calculated between a node and its neighbor based on the transmission link quality from the node to its neighbor and the receiving link quality from the neighbor to the node. The link costs for each node are added to the respective node cost of each mote. This, in turn, creates a gradient in node costs based on a mote's relative location to the base station.

Each mote in the network always selects the neighbor with the lowest cost as its parent node. These cost values are updated by routing and health messages. Routing messages, which are transmitted more frequently, keep track of neighboring nodes that are within the radio range of the mote. Health messages are sent through the network to the base station. XMesh uses an end-to-end acknowledgement with health messages. This means that a mote must receive an acknowledgment from the base station for each health message sent through the network. The purpose of this end-to-end acknowledgement is to update the network's routing table (i.e., neighbor and parent selection).

5.1.4 Test bed timeline

In April 2010, the first phase of the network deployment commenced. This consisted of a two-node network and a base station, which was connected to a gateway, located in the BFNR Nature Center. The gateway was located in an office within the Nature Center, where it had connections to electricity and the internet. The base station was connected to the gateway and set on the inside ledge of an office window, facing the field where the first two nodes were located. The purpose of the first phase was to test the connectivity between the base station inside the office and the nodes in the field.

Following the success of the first phase, the deployment moved into the second phase, which consisted of placing additional nodes in what is called site 1. The nodes in site 1 are proximal to the Nature Center and are for the purpose of relaying information only, not collecting environmental data. The purpose of the second phase was to build relays out to sites 2 and 3, which were designated as areas to conduct field sensor measurements. Sites 2 and 3 are located in the forested hill-sloped region of the reserve. Site 2 is located approximately 200 m from the base station and site 3 is located approximately 300 m from the base station. The organization of the network communication is setup such that site 3 relies on site 2 to communicate to site 1, thus the sites are successively located from the base station. At the end of the second phase, which was finished by late June 2010, there were 15 nodes in place.

The third phase consisted of filling sites 2 and 3 with additional relays and sensor nodes, which were used for collecting environmental data. By late July of 2010 all 25 sensor nodes (12 in site 2, 13 in site 3) were in place. At this time the network consisted of a total of 43 nodes, including both sensor and relay nodes, collectively. Following the third phase, network modifications consisted of relay changes, addressing both the quantity and location of relays, to increase network connectivity. As a result of these changes, the node count dropped to 40.

The first major reconfiguration of the relay nodes in site 1 occurred in late August 2010. The reconfiguration was done to accommodate a construction project taking place near the Nature Center. The reconfiguration did not, however, have a notable impact on the network. Connectivity before and after the relay node location changes remained relatively the same.

Originally, each wireless node was given an arbitrary identifier (i.e., node ID). To improve the management of the network, a new node ID labeling system was established in July 2011. In the new (and current) system, node IDs were associated with location and sensor

configurations. The naming scheme uses a four-digit identifier which relates three pieces of information regarding the node. The thousands place digit holds the node's site number (currently 1, 2 or 3). This allows for a quick reference of the node's location in the network. It also allows for up to six additional sites to be created. The tens and hundreds place digits provide a unique node number to each individual mote. This two-digit identifier allows each site to hold up to 100 nodes. The one's place digit characterizes whether the node is a relay or sensor node. It further indicates between relay nodes on the ground (i.e., ends in "1") or in a tree (i.e., ends in "0"). For sensor nodes, the ones digit distinguishes between nodes that are outfitted with only soil moisture sensors (i.e., ends in "3") or are attached to a tree for soil moisture and sap flow (i.e., ends in "5"). This naming scheme allows a user to quickly know where a node is located and what it is measuring without having to travel to the site. For example, the node ID 3113 conveys the following information: 1) the "311" indicates that this is the 11th node in site 3 and 2) the fourth digit, "3" indicates that this is a sensor node that is located on a PVC pipe and it is attached to soil sensors. This naming scheme is presented in the network map in Figure 31. In addition, in July 2011, the RUI of the network was updated from 36 s to 128 s.

From August 2010 until March 2012 the network consisted of 40 nodes. Figure 31 shows a map indicating the relative locations of the nodes within sites 1–3 during this time. The next major configuration change occurred in mid-March 2012 when the network was increased to 42 nodes. Both of these additional nodes were relay nodes, not sensor nodes, positioned near the interface of sites 1 and 2. The purpose of these additional relays was to increase the connectivity between the sensor beds (i.e., sites 2 and 3) and the base station. The last major reconfiguration took place at the end of March 2012. Nodes in site 1 were relocated once again due the end of the construction taking place near the Nature Center.

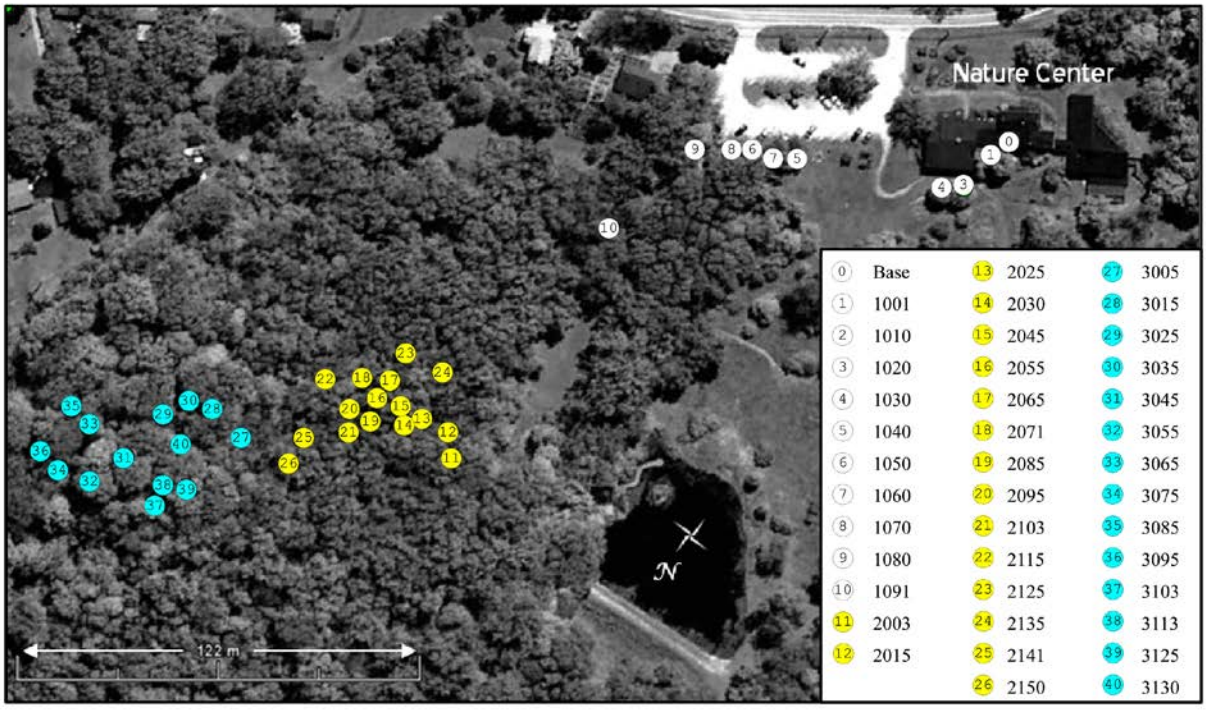

Figure 31. Map of the study area indicating the relative locations of the base station (which is housed within the Nature Center) and the nodes in sites 1–3. The nodes are divided into three sites which are color coded as: site 1 in white, site 2 in yellow, and site 3 in cyan. The node locations are representative to the network configuration in April 2011.

In May 2012, select MICAz motes in the network were replaced with the newer Iris module from Memsic (previously Crossbow Inc.). The Iris motes are more energy efficient and have a stronger antenna power compared to the MICAz motes. The Iris mote has a smaller current draw for receiving, transmitting, and sleeping modes (16 ma, 17 ma and 8 μA, respectively) compared to MICAz (19.7, 17.4 and <15 μA, respectively). Iris also has a strong radio, with a maximum power of 3 dBm whereas MICAz has a maximum radio power of only 0 dBm. The Iris radio is also more sensitive to low power signals (-101 dBm compared to MICAz at -94 dBm).

Nodes were chosen for Iris motes based on network pathways through the test bed. In addition to the base station, the other seven nodes which were replaced with Iris motes were: 1010, 1100, 2003, 2071, 2141, 3005, and 3113. The network pathways that were used for selecting these motes will be discussed in Section 5.2.2.

5.2 NETWORK EVALUATION

In this section the ASWP test bed will be evaluated by cost, both in terms of financial considerations and maintenance, and by examining complications incurred while deploying, maintaining, and collecting data from the network.

5.2.1 Financial considerations

One main consideration for deploying a network of environmental sensors is the cost of purchasing the necessary materials for the scope of the study (i.e., startup costs). For this

network, costs are broken down into the following categories: 1) WSN hardware. Hardware consists of the wireless motes, data acquisition boards, antennas, and antenna wires. This also includes the interface board for the base station; 2) Power. This includes the batteries that power the wireless motes, as well as additional batteries for environmental sensors; 3) Enclosures. This includes the waterproof enclosures, insulation, and any additional hardware needed for connecting enclosures as described in Section 6.3.2; 4) Mounting. This includes PVC pipes, wire, and screws for attaching enclosures to trees; 5) Environmental sensors. This includes the soil moisture sensors (EC-5), soil water potential sensors (MPS-1), sap flow sensors, wireless sap flow circuit, and other peripherals needed.

Table 15. Summary of associated costs for the 2012 WSN.

Category	Description	Cost
Hardware	Wireless motes, antennas, gateway, etc.	$12,980
Sensors	Soil moisture, water potential and sap flow	$9,360
Power	Batteries (AA, D, 12 V)	$1,340
Enclosures	Water proof boxes and insulation	$1,290
Mounting	PVC pipes, wiring, nuts and screws	$30
		$25,000

5.2.2 Maintenance costs

The maintenance costs include the deployment and the upkeep of motes and the network. The deployment time, as seen in Section 5.1.4, shows that the total installation took approximately three months to complete. It took three months because the deployment was broken down into three phases. At this time, three undergraduate students were working as interns and volunteers on the project and they assisted with the various stages of installation. After the network was deployed, regular maintenance visits were necessary to change batteries and perform other tasks.

Each node had its batteries changed between 15 to 25 times within the first two years of deployment (see Figure 32). Based on this, batteries needed changed, on average, every 29 to 49 days. This is consistent with the average time between maintenance visits as presented in Navarro et al. (2013), which presented a mean maintenance interval of approximately 38 days. These battery changes were conducted during 70 field visits within the two-year period. This suggests that network maintenance (i.e., of all the nodes) is required more often (on average every 10 days) than individual node maintenance.

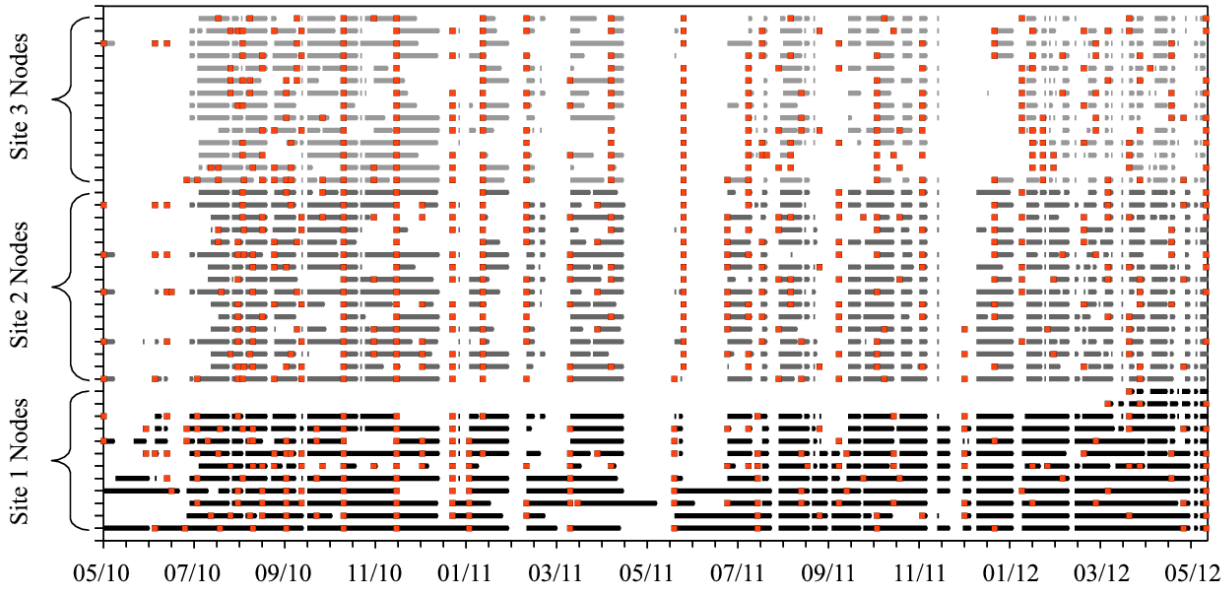

Figure 32. A schedule of mote battery changes (orange dots) for all the nodes in site 1 (dark gray lines), site 2 (medium gray lines), and site 3 (light gray lines) over the course of the two-year study period. Each line of information represents one node (12 nodes in site 1, 16 nodes in site 2, and 14 nodes in site). Gaps present in each node's line represents days where no data was collected from the nodes.

Network maintenance consists of individual node maintenance, gateway maintenance, and mesh maintenance. Individual node maintenance includes battery changes, as noted above, and also includes addressing hardware failures (i.e., faulty sensors, broken antennas, electrical short circuits) and software failures (i.e., errors that can only be resolved by reprogramming the mote). As noted in Section 5.1.1.1, a program was written to check and record the battery status of the wireless motes once per day. When a node's battery power drops below a certain threshold, an e-mail notification is sent to a list of users. This means that the researcher will know in advance which motes will need to be addressed prior to planning a field visit. Because of these notifications, battery changes can be planned for, however, not all maintenance concerns

can be known in advance. In order to be able to address the myriad of unpredictable problems that may arise, the researcher prepares for each visit to the ASWP test bed by carefully packing a variety of supplies.

These supplies include: spare antennas, as antennas become broken or dislodged; PVC pipe to situate any mote enclosures that have potentially fallen from tree branches; additional batteries; wire, in case the battery packs to the mote need reconnected; antenna wires, which are prone to breakage; and spare motes in case a mote has shorted and/or ruined by water damage. A box of these materials is always taken to the test bed in preparation for unexpected problems. An engineer's bag is also taken into the field. This bag holds the following contents: safety glasses, screwdrivers, voltmeter, marker, ruler, multi-tool, spare nuts and bolts, and silica packets (which are used to mitigate humidity within mote enclosures). It is essential to be prepared to maximize the amount of tasks that can be performed in one site visit. On average, site visits can last between 3 to 4 hours and can last upwards of 8 hours if time consuming problems arise, such as network outages.

Another aspect of network maintenance is the upkeep of the gateway. While most of the programs that were developed to run on the gateway have been automated, there are still some general tasks that need to be done manually (i.e., operating system and application updates). For security purposes operating system and application updates are applied on a regular basis. However, TinyOS is an exception. Updates cannot be made to TinyOS, because the newer version of TinyOS is not compatible with XMesh applications. Linux kernel updates are also a concern because these updates require a system reboot. Reboots can cause loss of data (i.e., network measurements made while the gateway is offline). Another issue with kernel updates is that the gateway computer has a limited amount of space allotted for Linux kernels. Because of

this limitation, old kernels must be deleted intermittently to make room for installing new kernels. If updates are downloaded and there is not enough room to completely install them, then a fragment of a file will be created, which will not load properly. This will prevent the system from booting. It is an issue to be conscious of when addressing gateway updates, though it is not a common problem.

Another update that is of concern is for the postgreSQL database. The postgreSQL server has to be shut off while it is being updated. Since it is responsible for collecting all of the wireless measurements, having it turned off for updates means that measurements will not be stored. These updates typically last a short duration thus minimizing the amount of lost data.

An additional maintenance consideration for the gateway remote connectivity is that port forwarding must be used to route incoming messages over the internet to the gateway via the network router. Because most local networks use dynamic host configuration protocol (DHCP) it is possible for the gateway's IP address to change. Depending on the type of router available, port forwarding management may require manual updates.

XMesh maintenance involves maintaining a connected network of motes within the field. It may require relay node relocations or additions. Node relocations address both how they are spatially arranged and their height above the ground. This can be challenging, given that site 1 has less trees than sites 2 and 3 and there is also a maintenance road that separates sites 1 and 2. Heath statistic data is collected from the network, which shows how well the nodes are communicating with one another. This helps determine if/when relay nodes need to be added or relocated. For example, in the spring of 2012 site 2 nodes were having difficulty connecting to site 1 nodes. To improve the communication (as noted in Section 5.1.4), an additional two relay

nodes were installed at the interface between sites 1 and 2, which improved the connection. This type of maintenance is addressed as needed.

Following network outages or changes to the network, there is a waiting period before those changes are propagated through the network, known as the mesh formation time. This waiting period before network changes are realized adds to the maintenance time of the network. As mentioned in Section 5.1.3, the mesh formation time is a function of the number of transmission hops (i.e., the number of parent nodes between the origin node and the base station) through the network.

To understand the mesh formation time, two different parameters will be examined. First, are the path costs reported by each node in the node statistics packets, which are stored by the base station. Path costs are four times the estimated transmission number (Crossbow, 2007a). Therefore, the estimated transmission number (i.e., hops through the network) can be calculated based on the path cost values recorded by each node.

In the following two figures (Figure 33 and Figure 34), the inverse distance weighted spatial distribution of the average annual path costs (as reported in the nodes' health statistics data) for nodes in sites 1–3 are shown.

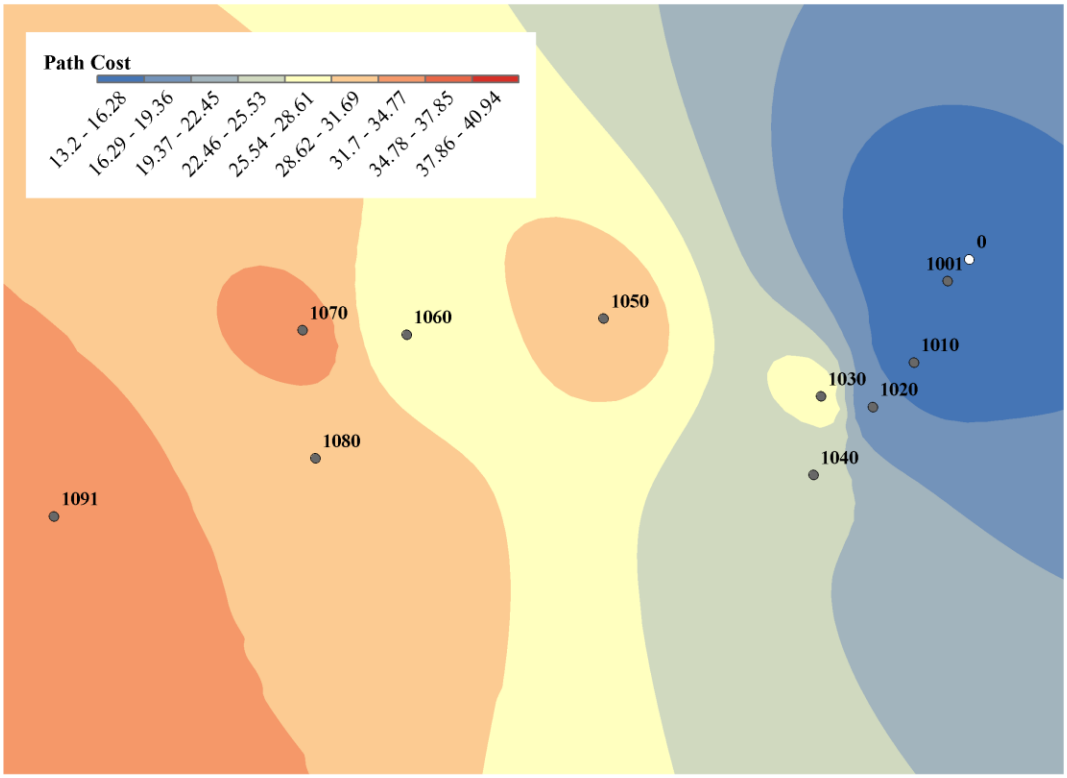

Figure 33. Spatial distribution of one-year average path costs over site 1 nodes.

The second parameter is an estimation of the actual paths that packets take through the network. This is necessary because the actual path that packets take through the network (beyond the parent node) is not recorded in any data by the nodes. Therefore, to estimate the paths that packets take through the network, the probability of node-parent pairs can be used. Node-parent pairs are recorded in the sensor data collected by each node. By examining all the node-parent pairs within the network, an adjacency matrix can be formulated (as seen in Table 12 in Section 4.2). Based on the adjacency matrix, the path of any packet generated in the network can be simulated.

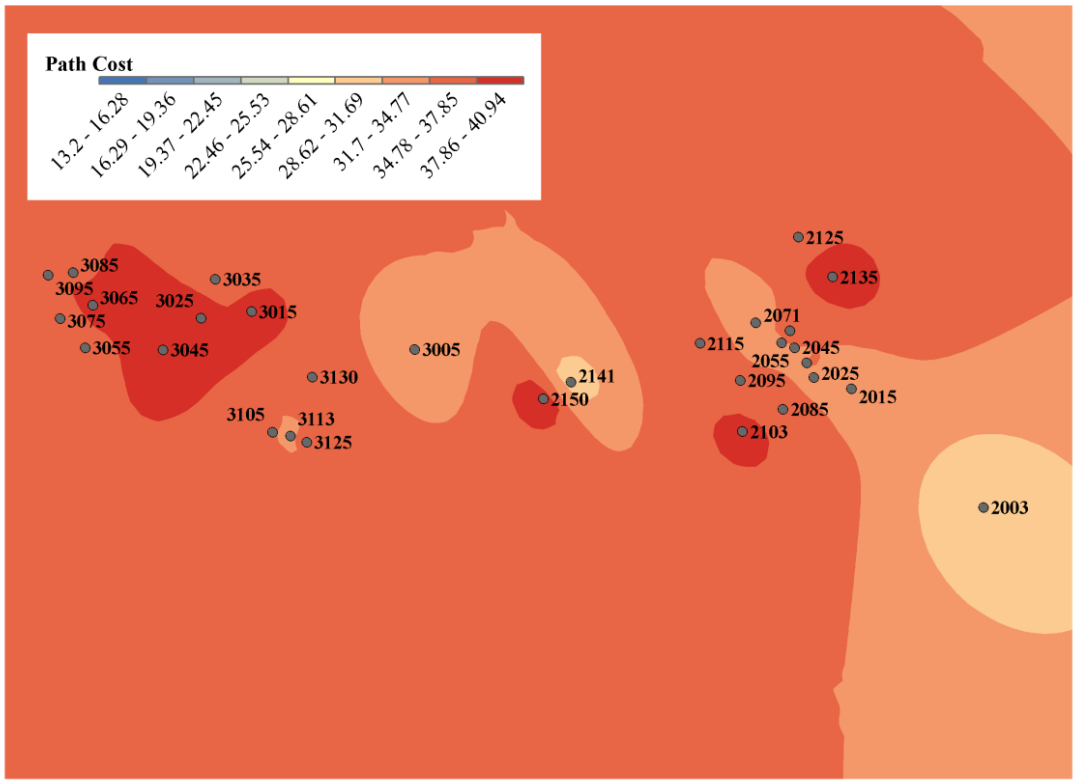

Figure 34. **S**patial distribution of one-year averaged path costs over nodes in sites 2 and 3.

The node–parent pairs shown in Table 16 are from a two-year period (12 May 2010 through 12 May 2012). It should be noted that the table results do not adequately reflect the impact of the two newly added relays (nodes 1100 and 1110), especially node 1100 which takes the majority of transmission forwarding between sites 2 and 1.

Table 16. Top two node–parent pairs, percent time paired, and distance between pairs (5/2010 to 5/2012).

SITE 1 Node ID	Parent (node)	1 Paired (%)	Dist. (m)	Parent (node)	2 Paired (%)	Dist. (m)
1001	0	55	4	1020	15	20
1010	1001	38	12	0	27	16
1020	1001	40	20	0	29	24
1030	1020	21	6	1001	19	25
1040	1001	24	68	1020	20	55
1050	1040	28	14	1020	22	69
1060	1030	25	53	1010	22	65
1070	1060	35	13	1050	22	4
1080	1070	36	14	1010	15	100
1091	1070	30	45	1050	27	49
1100	1010	50	120	1050	28	52
1110	1100	76	17	2135	9	57

SITE 2 Node ID	Parent (node)	1 Paired (%)	Dist. (m)	Parent (node)	2 Paired (%)	Dist. (m)
2003	1070	30	134	1091	25	94
2015	1091	30	88	2003	22	21
2025	2003	35	26	2055	16	6
2030	2003	39	28	2055	12	2
2045	2003	28	30	1091	23	88
2055	1091	63	89	2003	15	32
2065	2003	36	32	2055	15	2
2071	2003	35	36	2055	13	4
2085	2003	37	28	2055	15	9
2095	2003	36	34	2055	15	9
2103	2003	49	33	2055	9	15
2115	1091	18	98	2003	15	40
2125	2003	29	40	2055	17	16
2135	1091	44	76	2003	15	34
2141	1091	59	116	2003	14	56
2150	2141	45	5	2095	23	29

Table 16 (continued)

SITE 3 Node ID	Parent (node)	1 Paired (%)	Dist. (m)	Parent (node)	2 Paired (%)	Dist. (m)
3005	2141	21	21	2003	16	77
3015	2141	13	43	2115	12	62
3025	3035	22	6	3125	15	20
3035	2141	28	48	3005	10	27
3045	3005	33	35	3105	15	17
3055	3005	20	46	3105	13	26
3065	3105	19	27	3113	16	29
3075	3125	20	34	3005	17	49
3085	3113	22	33	3105	19	31
3095	3113	27	36	3105	16	33
3105	2141	35	44	3005	23	26
3113	2141	32	42	2103	22	64
3125	2141	62	40	2103	18	62
3130	3005	17	16	2141	15	36

Adjacency matrices were generated based on the daily sensor data collected at the gateway. For each packet generated in the network, the path it takes to the base station was simulated 25 times to determine the most probable path taken. Based on the daily simulated transmission paths for each node, the average number of hops that packets take to the base station was calculated. Over the two-year period, the average number of hops for site 1 was 2.99; for site 2, the average was 4.63; and for site 3, the average was 5.49.

Based on the information presented in Figure 33 and Figure 34, the average number of hops in: site 1 is 5.65, site 2 is 7.78, and site 3 is 8.23. It can be noted that the number of hops estimated, based on the path costs, is about 60% larger than the number based on the network simulation. This is because the network simulation assumes that packets are transmitted over the most probable paths, when in reality less efficient paths are utilized (i.e., additional hops are

necessary). Based on these hops and the RUI of the network, the formation time for site 1, on average, is between 6 and 12 min. For site 2, the average time is between 10 and 17 min and for site 3 the average time is between 11 and 18 min. Therefore, for the entire network, formation time is between 28 and 46 min.

These values are most likely low estimates, because this assumes that every transmission a mote makes to a parent is successful, which is not always the case. Based on results of this network's performance (Navarro et al., 2012), and the results obtained from the pilot test bed presented in Section 4.2, the transmission success of motes is below 50%. This means that actual values for formation time could be double.

The purpose of analyzing the formation time is to show that if, for example, a change is made to a node in the network or if the network is restarted following an outage, there is a long wait time before the mesh becomes stabilized. When it comes to maintenance, changes at the node level, within this particular network, are typically seen in the scale of hours and at the network level changes are seen in the scale of days.

A final consideration for maintenance involves the node IDs. As discussed in Navarro et al. (2013), it was observed that a few nodes occasionally and arbitrarily changed their node ID after restarts. However, in other respects, aside from the node ID change, these nodes continued to function properly. This situation is likely to be related to a software bug in XMesh. This poses a problem with respect to network and data analysis. This is due to the fact that data will be stored on the gateway under a false node ID. In most cases, it is possible to identify the real node ID based on the data and node configuration. However, under certain circumstances (e.g., lack of network data) it is not possible to discern the original node ID. Additionally, this adds preprocessing maintenance because the original node ID must be identified and corrected.

5.2.3 Complications and considerations

Numerous complications and considerations for improvement were addressed during the life of the ASWP test bed. These complications could also be considered maintenance costs, since it took time to address each of the problems. However, for organizational purposes these maintenance issues will be addressed as complications and considerations under the following topics: gateway connectivity, WSN communication, physical damage to hardware, and battery life. These broad points encompass the challenges encountered when deploying and maintaining the network.

5.2.3.1 Gateway connectivity

During the time of the Sluggo gateway, the login credentials were, by default, generic. The port forwarding was setup for the remote connectivity to the gateway to allow anyone the access to login to the gateway via the internet. The only information someone needed to login is the IP address, username, and password. In late September 2010, after five months of operation, the Sluggo gateway was found to be inoperable (i.e., the whole system was in read-only mode). After investigating the authentication logs on the gateway, it was found that numerous repeated login attempts were made to the gateway. Because an unauthorized user hacked into the system, the Sluggo went into read-only mode in order to protect itself. It then became clear that safeguards needed to be put in place to prevent unauthorized users from hacking into the system.

The first line of defense established was to change the password into something more complex than it was previously. For the second line of defense, the authentication logs were searched and unauthorized IP addresses were identified and blocked. The Linux command iptables was used to restrict access to blocked users. Within the first year of operation, 175

unauthorized attempts were recorded, which grew to over 600 unauthorized attempts by May 2012. Following a change in the internet service provider for the BFNR in May 2012, the rate of failed attempts significantly reduced. As of October 2012, a total of 724 blocked IP addresses have been logged.

Another complication in maintaining the WSN is the limited access to the gateway. The Nature Center, which houses the Linux gateway, is open during standard hours of operation for the Audubon Society. It is closed on Mondays, has reduced hours on Sundays and is closed in observation of holidays. There have been several occasions over the life of the testbed when accessing the gateway was not possible, which delayed necessary maintenance from being performed. For example, following power outages the gateway does not automatically turn back on and it must be manually restarted. When the gateway is off, data cannot be collected and the network begins to cycle in search of the base station, which (as observed in field deployments) drains mote batteries quickly.

Once real-time monitoring (i.e., INDAMS) was established in the network, XServe became controlled by the management system. In addition to real-time monitoring, INDAMS provided remote control network functionality. During lapses in internet connectivity, the agent (part of the management system running on the gateway) continued to operate the network gateway offline. A problem occurred if the agent stopped running (due to power outages, system reboots, updates, etc.) while there was no internet connection to the gateway.

Starting on 24 May 2012, remote connectivity to the gateway was lost due to an internet service provider change at the BFNR Nature Center. The internet was still available, but port forwarding was not (i.e., no outside access to the gateway). The agent could no longer be used, which meant that XServe had to be managed manually. This was a new problem, as the Linux

gateway had never been run manually up until this point. To address this problem, XServe, similar to the other monitoring scripts that were currently running on the gateway, was also monitored. The XServe monitoring script would check to see if XServe was running every hour; if it was not running the script would attempt to restart XServe. XServe's status was logged in Dropbox, which provided hourly updates on the XServe status throughout the remote connectivity outage. The outage lasted until 25 July 2012, which was the day remote connectivity was restored and the agent came back online. Had there not been a backup plan in place for outages like this, two months of data would have been lost.

5.2.3.2 WSN Communication

There are several aspects of network deployments which affect WSN communication. The first concern deals with network stability (i.e., the efficiency of network communications). Noted previously, there are two link qualities (transmission and receive) which determine the cost of transmitting between two motes. These two link qualities are often not equal and may affect network communication. Environmental conditions as well as radio frequency interference (i.e., RF noise) influence the link quality between nodes.

There are numerous ways to characterize network stability including the transmission success and battery life of motes. Another metric for determining network stability is the number of parent changes a node undergoes over time (Teo et al., 2006). Nodes select their parents based on the lowest transmission cost to the base station. Ideally, there should be a well-defined path for message transmission through the network which would cause few parent changes to occur over time. The number of parent changes increases when either nodes enter or leave the network (e.g., battery changes or node relocations) or link qualities are poor between the node and all its neighbors. Figure 35a shows the average daily parent changes per each node

in the network. Zero parent changes indicate either static network connections or network outages. Network outages (i.e., no data received from any nodes in a site) are indicated by the number of transmission paths present in the network (see Figure 35b).

Figure 35a shows the stability of the network in terms of the number of node parent changes. Each individual node's sensor data was processed to count the daily number of parent changes. Then, all nodes' daily parent changes were accumulated and divided by the total number of nodes in the network. Figure 35b shows the unique paths through the network based on the simulation results (see Section 5.2.2). All the unique paths simulated per day were divided by the number of nodes in the network. There is a strong correlation between the number of parent changes and unique paths in the network. As the number of parent changes increases more unique paths are created. Therefore, time periods of low network stability can be identified from increases in the number of parent changes and subsequently more unique paths through the network.

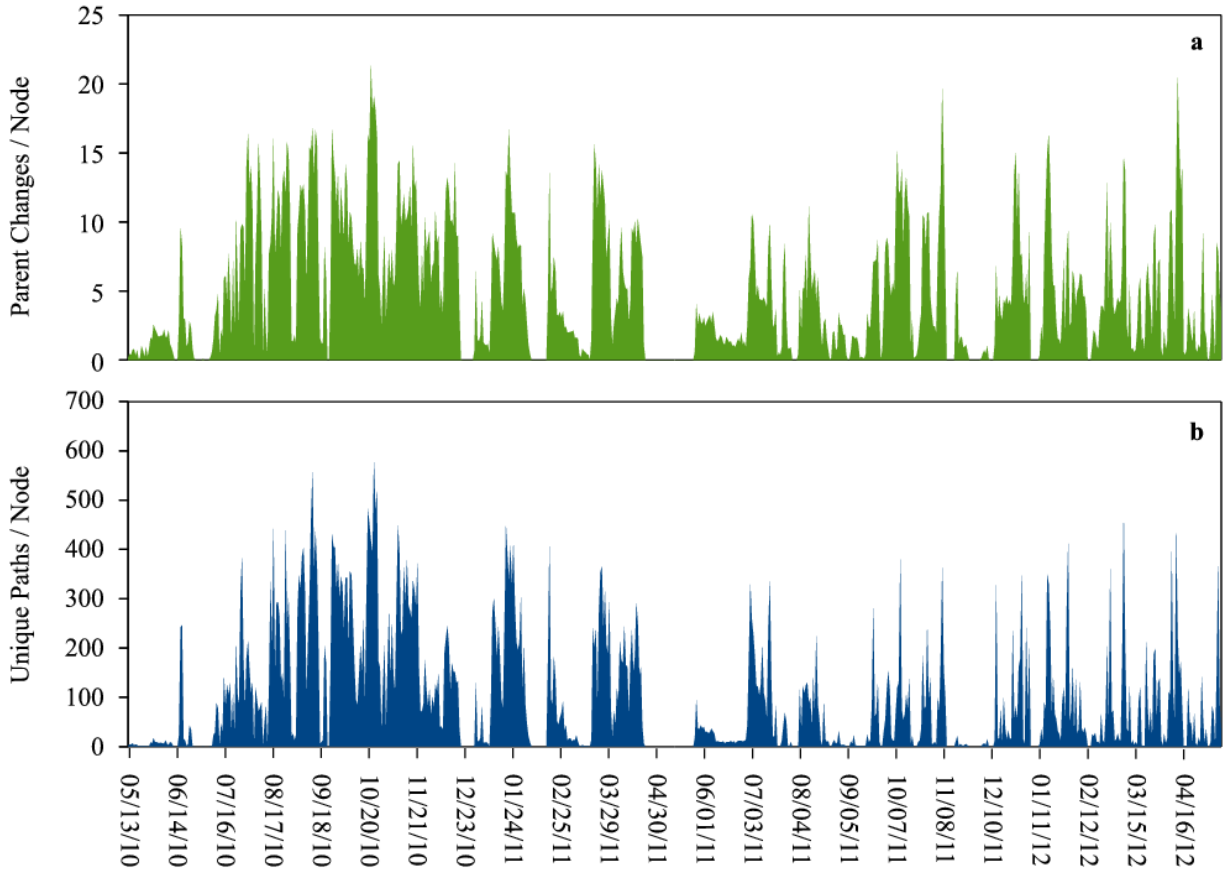

Figure 35. (a) The daily number of parent changes per node in the network; (b) daily unique transmission paths through the network.

Network stability problems can result in partial to whole network outages. Figure 36 shows network performance during the two-year operation of the network in terms of node connection as a percent of total nodes in the network and percentage of total expected data (based on the sampling rate and number of nodes in the network).

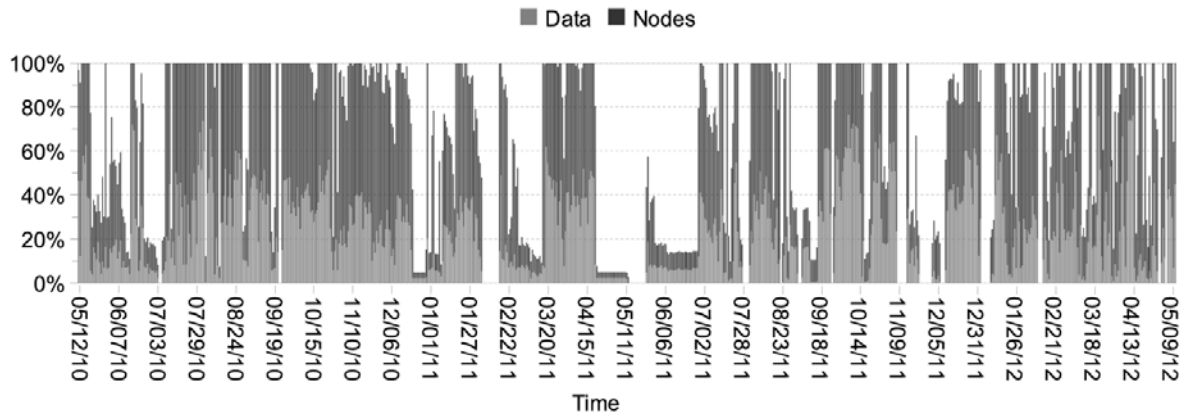

Figure 36. Bar graph of network performance of nodes (dark grey) and data (light grey). The network performance is calculated for each day of the network life where the percentage of nodes is based on the number of nodes heard from the field (at least one packet received) compared to the total number of nodes deployed and the percentage data is the number of lines of data received compared to the total expected data (based on the sampling rate and number of motes).

Figure 36 shows that following the establishment of sensor nodes in sites 2 and 3 (i.e., the end of the third phase of deployment) all nodes are connected within the network. Connection levels remain high for over five months (i.e., July–November 2010). Starting at the end of December 2010, for a period of ten days, the network experienced its first connectivity issues due to a high number of dead batteries within a large group of motes. This is indicated by the

drop in connectivity over the ten-day period. Extensive battery changes were made. Following the next few weeks, in January 2011, connectivity in the network was restored. Between 5–16 February 2011, the gateway was offline. No data was collected at this time and it represents the first substantial gap in the dataset. Network connection remained low until March. For the month of April, connectivity was again restored. In the beginning of May 2011, connectivity dropped again due to battery power losses throughout the network. Maintenance at the end of May restored limited connectivity, but it is not until July that complete connectivity is restored. At this time, sap flow monitoring sensors were deployed and substantial battery replacements were made.

Connectivity reduced again in August 2011 due to a second gateway outage. Even though the gateway was restarted within a day, it is not until September that full connectivity was restored to the network. This is because battery changes were also taking place throughout August and, as previously discussed, mesh formation time can be long. In mid-October there was another low connectivity due to batteries discharging in the network motes. In November and December 2011, there were a number of gateway outages, resulting in gaps within the dataset. From January to May 2012, there was periodic maintenance, which resulted in highly variable network connections.

In terms of network efficiency, the communication quality between nodes and their parents can be analyzed via the received signal strength indicator (RSSI) values and the communication quality between nodes and the base station can be analyzed via the path cost, both of which are recorded in the health statistics data. For comparison purposes, the average network path costs, RSSI values, and number of network paths are presented in Figure 37.

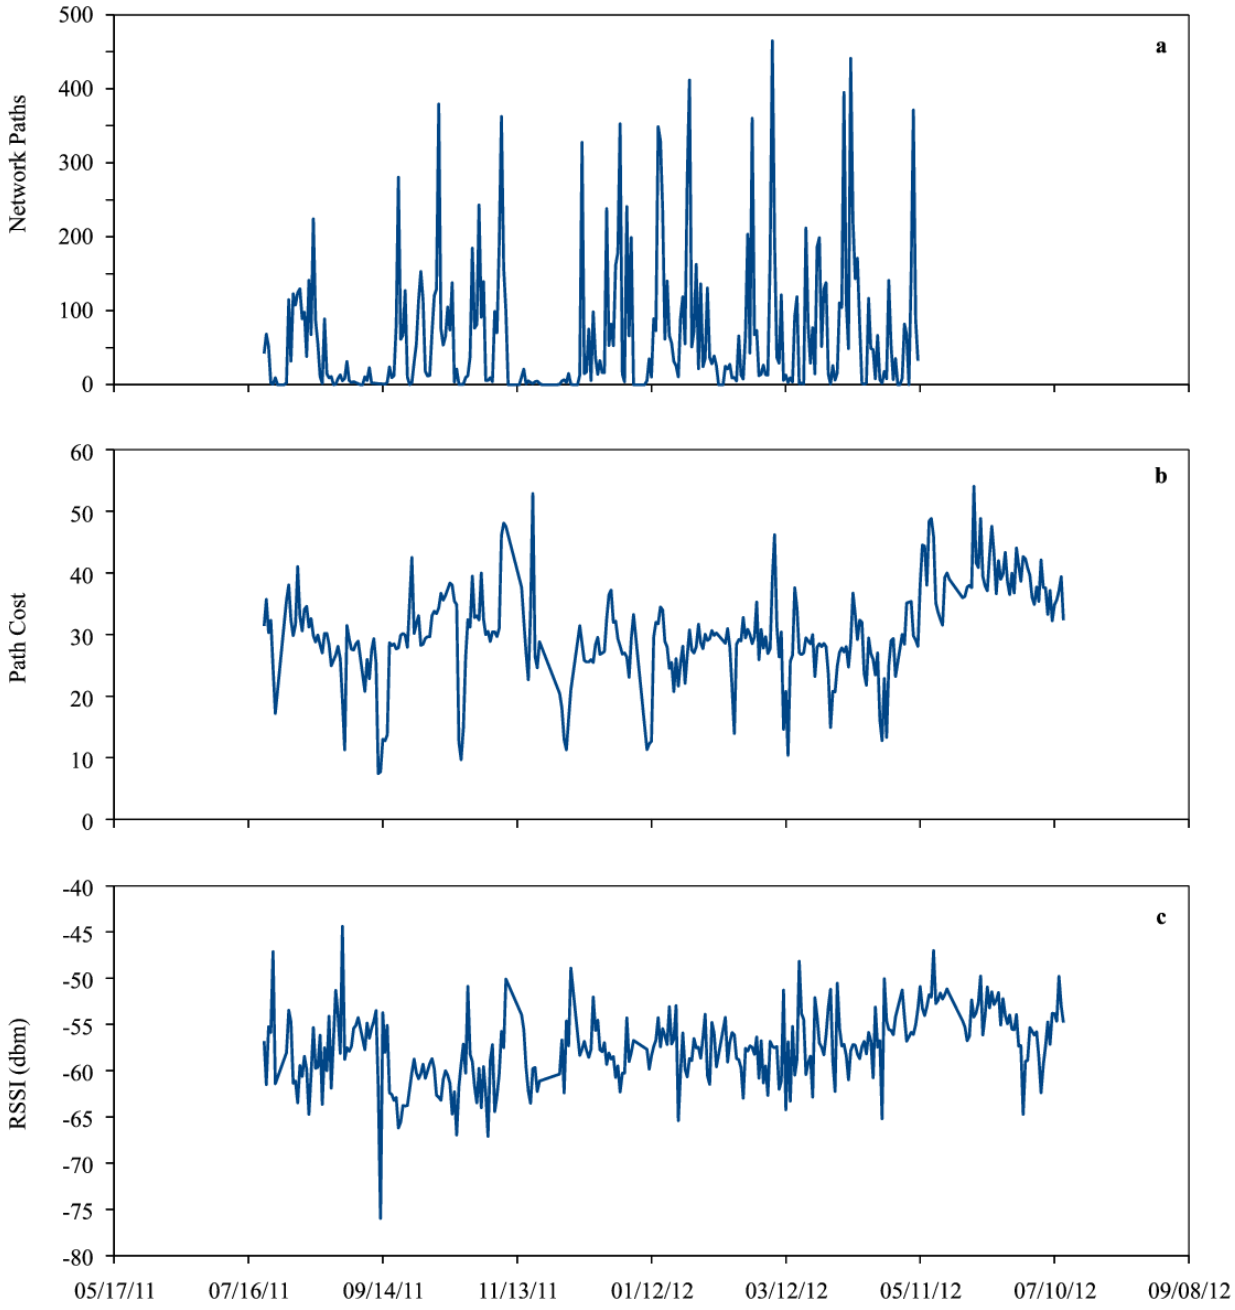

Figure 37. Time series of (a) number of network paths, (b) network-averaged path cost and (c) network-averaged RSSI. Note that path costs and RSSI values have been extended through July 2012.

Figure 37 shows the number of network pathways (based on the simulation results) and the network-averaged path costs and RSSI values over time. It should be noted that the path costs and RSSI values have been extended beyond the study period (which ends on 12 May 2012) to show results over one-year. Drops in the path cost time series correspond with drops in the number of network paths (i.e., network outages). Network RSSI tends to peak during periods of low network path costs. Again, these occurrences are during periods with low network paths (i.e., network outages) when only nodes in site 1 are in the network. Throughout the one-year period, RSSI values remain fairly steady with an average over -60 dBm (this is greater than the threshold for receiving signals on both the MICAz and Iris motes). There is a noticeable jump in the network path cost in May 2012 when the two new relay nodes between sites 1 and 2 were added to the network. However, this increase in cost is seen to be gradually decreasing as RSSI values improve.

The link quality between nodes and their parents are not symmetric. This problem leads to circumstances where an origin node can communicate its packet of information to its parent, but cannot receive an acknowledgement from its parent that the message was received. Asymmetrical links within the network causes duplicate packets to be created (as was identified in Section 4.1.4). This is the case when a node sends its packet of information to its parent multiple times. In each case, the parent node received the packet from its child, but due to poor link quality, the child node never receives an acknowledgement, and continues sending its parent the same information. XMesh defines the number of retries (i.e., eight) each mote will make in attempt to successfully transmit its packet to its parent. Based on the network analysis performed by Navarro et al. (2013), on average between 3–4% of packets received at the base station are duplicates. The significance of duplicate packets on network statistics monitoring is

discussed in Navarro et al. (2013) and the significance of duplicates on sensor measurements is discussed in Chapter 6.0

A lesser known problem affecting network communications is interference in the transmission frequency band. Motes communicate with one another on a specific frequency channel. The motes' antenna can only receive messages if they are perceived at a power level above a certain threshold (discussed in Section 5.1.4). The power that is used to transmit messages decreases over the transmission distance. Radio interference can also effect the distance a message travels. Therefore, an investigation on radio interference was performed. The goals of the investigation were to identify radio interference on the 16 available frequencies motes can transmit messages over and compare the levels of interference to determine if the channel being used (to transmit messages) is the best option for the network.

The MICAz motes can be programmed to transmit in one of 16 frequencies in the 2.4 GHz frequency band. The frequency channels range from 11 to 26 (2.4-2.48 GHz). The motes in this study operate on channel 12 (2.41 GHz). The available channels and their respective radio frequencies are shown in Table 17.

Table 17. IEEE 802.15.4 channels and their respective frequency ranges (Crossbow, 2007c).

Channel	Frequency (MHz)
11	2405 ± 1
12	2410 ± 1
13	2415 ± 1
14	2420 ± 1
15	2425 ± 1
16	2430 ± 1
17	2435 ± 1
18	2440 ± 1
19	2445 ± 1
20	2450 ± 1
21	2455 ± 1
22	2460 ± 1
23	2465 ± 1
24	2470 ± 1
25	2475 ± 1
26	2480 ± 1

A background study of the radio frequency noise was conducted at the ASWP test bed over the spring and summer of 2012. The threshold sensitivity of the MICAz radio is -90 dBm (Crossbow, 2007f) and therefore there should not be background noise higher than -90 dBm. The most common sources for radio interference in the 2.4 GHz band include cell phones, Bluetooth, microwave ovens, video cameras, and Wi-Fi networks. Of those listed, the only static source located onsite is Wi-Fi. Wi-Fi networks are most commonly configured on channels 1 (2412±11 MHz), 6 (2437±11 MHz) and 11 (2462±11 MHz). It should be noted that the band width of the Wi-Fi is larger than the band width of the MICAz transmission channels. Therefore there is some overlap between Wi-Fi and MICAz channels except for channels 15 (2425 MHz), 20 (2450 MHz), 25 (2475 MHz) and 26 (2480 MHz). While these four channels are often recommended for Zigbee devices (e.g, MICAz motes), the dynamic nature of radio interference

in the natural environment prompted the actual sampling of the background noise at each channel.

The sampled background noise is plotted in terms of the RSSI which is the measure of radio power heard by the mote antenna. Each channel was sampled in approximately the same location in site 1 of the ASWP test bed. The RSSI sampling program was created by HyungJune Lee from the Department of Electrical Engineering, Stanford University and is available for download on SourceForge.net (Lee, 2007). The application samples the RSSI at the mote's antenna every 1 ms (i.e., 1 kHz frequency) for approximately 18 min.

Figure 38 shows the 1000-point averages of the 250,000 samples of RSSI for each of the 16 channels. The average RSSI values vary between -94.2 and -92.0 which are below the perceivable noise level of the MICAz mote. All the channels had peak RSSI values greater than the mote threshold (i.e., -90 dBm). In Figure 38, channel 12 is highlighted in red and the two solid black lines represent the highest and lowest levels of noise out of the four recommended channels. Based on the noise sampling, it is seen that channel 12 falls between the two solid lines.

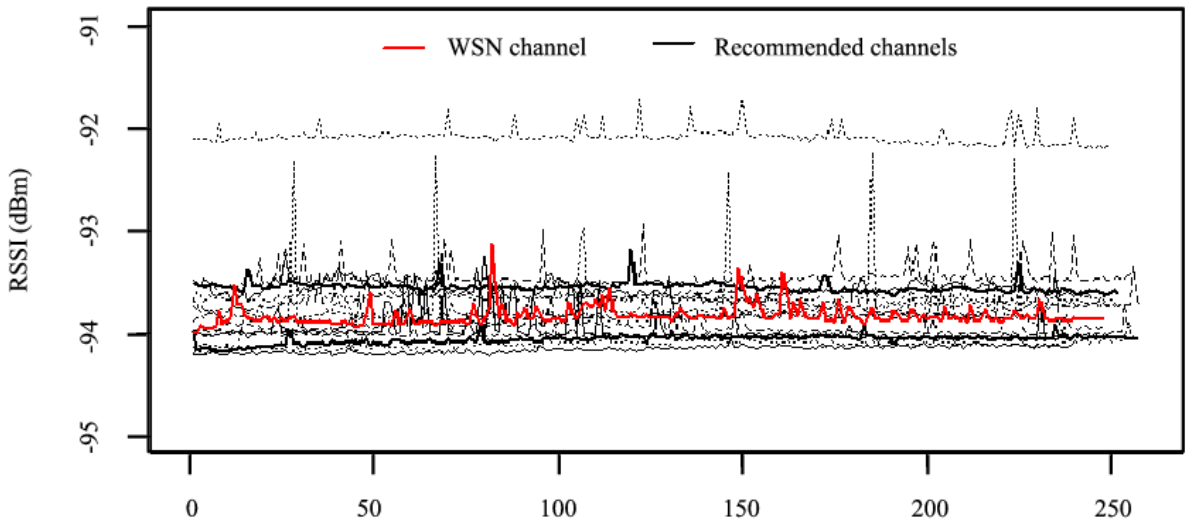

Figure 38. Average radio frequency noise (dBm) measured at 16 channels in the 2.4 GHz band (2.405–2.480 GHz). Noise measurements in each channel collected approximately 250,000 samples at 1 MHz. Averages are based on 1000 points (N ≈ 250).

Table 18 shows the average and standard deviation of the total RSSI samples and perceivable RSSI samples (i.e., greater than -90 dBm) as well as the peak RSSI (based on the results shown in Figure 38).

The transmission frequency of this network (i.e., channel 12) was arbitrarily selected. However, based on the results in Table 18, the mean perceivable RSSI and peak RSSI for channel 12 outperforms two of the four recommended channels (i.e., channels 15 and 20). The only other channels that are better than channel 12, in terms of the level of background noise, are the two other recommended channels (i.e., 25 and 26). These results suggest that channel 12 is suitable and unless further information is found that shows background noise is effecting the WSN motes, changing the channels is not necessary.

Table 18. Averages and standard deviations of the total and perceivable (> -90 dbm) RSSI samples as well as peak RSSI of the 16 frequency channels.

Channel	Mean RSSI (dbm)	Mean Perceivable RSSI (dbm)	Peak RSSI (dbm)
11	-93.56 ± 0.78	-86.66 ± 7.23	-30
12	-93.83 ± 0.80	-86.82 ± 4.94	-31
13	-92.08 ± 1.01	-86.32 ± 9.38	-29
14	-93.45 ± 1.27	-84.90 ± 10.9	-26
15	-93.55 ± 0.76	-61.31 ± 21.1	-26
16	-93.93 ± 1.08	-83.61 ± 6.62	-25
17	-93.65 ± 0.86	-86.31 ± 5.27	-27
18	-93.68 ± 0.84	-59.60 ± 16.0	-24
19	-93.83 ± 1.00	-85.32 ± 8.80	-23
20	-93.96 ± 1.55	-84.11 ± 12.9	-21
21	-93.86 ± 0.62	-71.20 ± 18.7	-27
22	-93.75 ± 1.32	-72.47 ± 20.2	-35
23	-94.02 ± 0.54	-74.07 ± 16.5	-24
24	-94.10 ± 0.48	-80.42 ± 10.7	-45
25	-94.15 ± 0.44	-86.30 ± 2.58	-80
26	-94.06 ± 0.46	-80.61 ± 5.59	-70

5.2.3.3 Physical damage to equipment

Another complication encountered was the physical damage done to equipment due to short circuits, mote box fatigue, water damage, and animal interaction. Short circuiting was encountered for two different reasons. Motes rely on a battery pack for power and the battery pack is wired either directly to the mote or through the data acquisition board. The battery pack is equipped with an on/off switch. There were some cases where battery pack on/off switches malfunctioned, resulting in a short circuit of the batteries, which in turn destroyed the battery pack. The second way that nodes short circuit is due to the insulation used within the enclosures. This problem only effects enclosures that house the sap flow circuit board. The insulation is a

reflective metal foil. If care is not taken with placing the circuit board within the enclosure, the circuit board could come in contact with the insulation, which can cause a short circuit.

Another concern for equipment is mote enclosure damage. Mote enclosures (i.e., the PN-1337 polycarbonate enclosure) are subject to fatigue caused by thermal stresses (brought on by freezing temperatures) and vegetative growth. When nightly temperatures drop below freezing, water (which can reside within screw holes or in cracks in the plastic) freezes and expands which breaks apart the enclosure. Mote enclosures attached to trees for sap flow monitoring can also become fatigued due to seasonal tree growth. The tree trunk presses against the enclosures and can fracture the box. Fatigued and cracked enclosures are more susceptible to water intrusion.

Despite the fact that enclosures are watertight, numerous holes had to be cut into the boxes to allow for external antennas and to connect two enclosures together for the sap flow mote box designs. These holes are locations where water has the potential to intrude. Rubber gaskets and silicon adhesive were used around these potential problem areas, however, in some instances, water was still able to get into the boxes. Silica desiccant packets were placed inside each box to help protect the mote hardware and circuitry from water. However, motes submerged in water became inoperable and the motes and their enclosures required replacement.

Alternatively, animals were also responsible for some of the damage incurred to the equipment. By far, the most damaged part of the enclosure boxes was the antenna. The 6" long antenna protrudes from the side of each mote box. In the case of enclosures attached to the sides of trees, mote boxes provided a place for squirrels and chipmunks to sit and chew on the ends of the antennas.

5.2.3.4 Battery life

The life of a wireless network is highly dependent on the battery life of its individual nodes. Factors that influence battery life include the packet size of data being transmitted, the sampling rate in which data is being sent, and the distances over which packets are transmitted. Due to a node's autonomous nature they depend on finite energy sources, therefore the size and type of batteries will also influence network life.

Packet size is based on the information being transmitted through the network (i.e., sensor data, health data, and routing data). The maximum packet size that MICAz motes can transmit is 128 bytes. This study does not investigate changing the packet size or its influence on battery life.

As mentioned in Section 5.1.3, there are two types of packets that are transmitted through the network from nodes to the base station (i.e., sensor data and health data) and one type that is transmitted from a node to its neighbors (i.e., route updates). Each of the three packet types (e.g., sensor data, health data, route updates) have their own transmission rates. Data messages for all nodes in this case study were set to transmit every 15 min in LP mode. Health messages, by default, transmit every 10 min. The RUI for the motes for the first year of deployment was incorrectly set to 36 s (default RUI for motes in HP mode). The default RUI for motes in LP mode is 360 s. However, due to a bug in the application code compiler, motes in LP mode received the same RUI as motes in the HP mode setting. This bug was addressed with a patch released by Crossbow that corrected the error. As a compromise between power savings and mesh formation time, a new RUI was set to 128 s.

A previous study has shown the relationship between antenna power and transmission distances (Teo et al., 2006). The antenna power can be adjusted for the MICAz mote. There is a

tradeoff between using more power to transmit a message over a long distance and thereby saving the number of relays necessary verses using less power to transmit over short distances but then requiring more relays to send the message. This study attempts to reduce the number of relays because of the limitations on where they can be placed, therefore uses the highest power setting for antennas to transmit over longer distances. By default the antenna power of the motes in this case study are set to a power level of 0 dBm. The average distance motes in sites 1, 2 and 3 transmit is 33.6 m, 40.1 m and 37.9 m, respectively. The average distances between nodes in Table 16 are calculated based on the 2011 GPS survey and weighted by the percentage of parent connection based on the adjacency matrix. Transmission distances have been recorded in excess of 170 m with numerous accounts of node transmissions over 100 m. The MICAz outdoor range is between 75 and 100 m (Crossbow, 2007f). As mentioned in Section 5.1.2, the external antennas used by the motes provide a 4.9 dBi gain, which is likely the cause of the above average transmission distances experienced at the ASWP test bed (i.e., 170 m verses the published potential of 100 m).

The batteries used in this case study are rechargeable AA NiMH (nickel-metal hydride). There are two different battery capacities used: 2450 mAh and 2500 mAh. All relay nodes are outfitted with a pair of battery packs that hold three AA batteries each. This is because relay nodes primarily forward messages from other nodes. Therefore the extra battery pack is in place to compensate for the extra draw on their power supply due to the higher transmission rates from forwarding packets. Sensor nodes are equipped with only one battery pack which holds three AA batteries. The battery packs that hold three AA batteries were chosen instead of the standard packs which only hold two AA batteries because the rechargeable AA batteries hold a smaller voltage (i.e. 1.2 V nominally). Non-rechargeable batteries have a voltage of 1.5 V.

Rechargeable batteries were chosen over the non-rechargeable due to cost concerns and environmental considerations. The three rechargeable AA batteries will produce a voltage between 3.6 and 4.0 V, which is an acceptable voltage for operating the motes, despite being over the recommended operating voltage (2.7 to 3.3 V). This was determined via personal communication with a Crossbow technician.

Despite the cost savings and environmental benefits, there are setbacks to using rechargeable batteries. Within the first six months of study, it was found that the batteries recharged to different voltages. After the batteries are recharged, they are left to rest and then tested. This is done because individual rechargeable batteries discharge at different rates. The resting period was to filter out batteries with high discharge rates. Batteries with a voltage of between 1.25 and 1.31 V were used in the field. Batteries were grouped by voltage, meaning that only batteries with similar voltages were put together in the same battery pack. By testing and grouping batteries together, it helped reduce the number of dead batteries being used in motes, and therefore it increased network life (see Figure 32).

In an effort to save battery power, the three LED indicator lights on the motes were turned off, which was also done in the pilot test bed discussed in Chapter 4.0 The LED lights consumed approximately 2.5 mA each (Crossbow, 2007c; Krämer and Geraldy, 2006).

5.3 DISCUSSION

WSNs are more affordable in terms of startup costs than a traditional data logger network. However, this savings is offset by the high maintenance costs of WSNs. While the

complications and considerations for the WSN were initially found to be problematic setbacks, once these problems were addressed and/or mitigated they became more manageable.

Some of the most significant issues encountered revolved around internet availability. The first issue was how to access data remotely. The current solution uses INDAMS, which requires router access at the host facility. It utilizes port forwarding to transmit data to the management system. As a backup, the gateway also saves daily records of network data to a Dropbox account. This does not require port forwarding. An issue with port forwarding is knowing where the computer is located on the local network (i.e., the local IP address). While some routers associate port forwarding to a specific computer, most associate forwarding to a local IP address. Power outages, router restarts, and computer restarts can all lead to a new local IP address being assigned to the gateway. Therefore, it is essential to know the gateway's local IP address. This problem was solved using a simple script which logs the computer's local IP address and e-mails a notification if the IP address has changed.

The channel the WSN operates on does experience background noise. The background noise is characterized as being highly variable over daily, monthly, and seasonal time periods. Therefore, the noise sampling conducted over the spring and summer of 2012 may not offer a complete view of background noise at the ASWP test bed. Additional noise studies should be done in order to get a well-rounded view of noise levels in the 16 frequencies.

Gateway and maintenance issues are the main causes for low network connectivity in the test bed. Gateway issues were resolved by using the web-based management system INDAMS. Maintenance issues were addressed using automated battery monitoring on the gateway. INDAMS and the automation of battery monitoring have both helped improve network function.

However, issues still exist (e.g., electrical outages and internet outages) which future work can address.

The analysis of the health statistics data, recorded by the motes, can be analyzed to improve network stability. In regards to the ASWP test bed, network stability was analyzed in terms of link quality, path cost, and RSSI. These metrics provide useful insight on the efficiency of individual nodes (e.g., message routing and mote location) and can also be used to address the performance of the entire network.

With all outdoor monitoring networks, physical damage to equipment will be incurred. It should be an anticipated cost in maintaining an outdoor WSN. Nothing experienced in this study was beyond expected, with the exception of a lightning strike, which is further discussed in Chapter 6.0

An unforeseen issue was voltage variation in the rechargeable batteries. It was discovered that the rechargeable batteries have inconsistent voltages after recharging. Therefore, batteries were checked and grouped by voltage after each recharge and before use in the field. When batteries were not examined after recharge, often uncharged batteries were used in the field. This resulted in significantly shorter battery lives than expected.

Another challenge encountered involves power usage of certain motes. Motes that are frequently utilized to transmit data to neighbors are expended of their battery power more quickly than motes that are not (Davis et al., 2012c). This was apparent in motes from site 1, which relay data to the base station from motes in site 2 and site 3. To combat this problem highly utilized motes were given an additional battery pack, making the total number of batteries six per mote. Figure 32 shows the longer battery lives of nodes in site 1 (i.e., relay nodes).

6.0 ANALYSIS OF SAP FLOW, SOIL MOISTURE, AND SOIL WATER POTENTIAL MEASUREMENTS USING WIRELESS SENSOR NETWORKING SYSTEMS

The purpose of this chapter is to analyze the environmental data collected from the ASWP WSN test bed. Analysis is performed on the two main sensor types: soil sensors and sap flow sensors. There are two types of soil sensors that are analyzed: the MPS-1 soil water potential sensors and the EC-5 soil moisture sensors. Data analysis for these sensors includes an examination of seasonal spatial distributions, a comparison of time series to data logger measurements, and an examination of the impacts of duplicate packets.

There is one type of sap flow sensor being analyzed: the thermal dissipation method sap flow sensor (Granier, 1985). Analysis includes the examination of temporal variations in sap flow and an examination of noise in sap flow sensor measurements.

6.1 SITE DESCRIPTION

This study takes place at the ASWP's BFNR (see Section 5.1). The reason for choosing this location is because BFNR is located within the East Little Pine Creek (ELPC) watershed, a sub-basin of the Pine Creek watershed in northern Allegheny County, PA (Figure 39). This study's watershed of interest is ELPC due to its size, close proximity to Pittsburgh, and because it

consists of mostly undeveloped land. Since this study focuses on the natural environment, rather than the human impacts on hydrology, finding an undeveloped location was important. Based on the 2006 National Land Cover Data (NLCD), the ELPC watershed is 71% undeveloped and less than 10% developed with impervious land (Figure 40). Due to the nature of this study, frequent maintenance visits were required, which made it unrealistic to consider watersheds outside the greater Pittsburgh area. Also, the watershed has a small catchment area (15 km^2), which limited the number of complexities added to the study. For all of these reasons, the ELPC watershed was chosen.

Once the watershed was chosen, the next step was to find a location where a test bed could be deployed. The study site at BFNR was selected because it met all of the study needs. These needs included: 1) internet access, 2) minimal human interaction to ensure the protection of the equipment, and 3) a forested hill-sloped area for study. BFNR, where offices, a nature store, and education facilities are located, had an internet connection, which could be used for this study's WSN. The reserve is also home to hiking trails and nature programs, which attracts mostly nature enthusiasts and is closed from dusk until dawn, providing some protection to the research equipment. Included in the nature reserve is an oak forest, which is proximal to the Nature Center. This forest provided an area for deploying the WSN and environmental sensors.

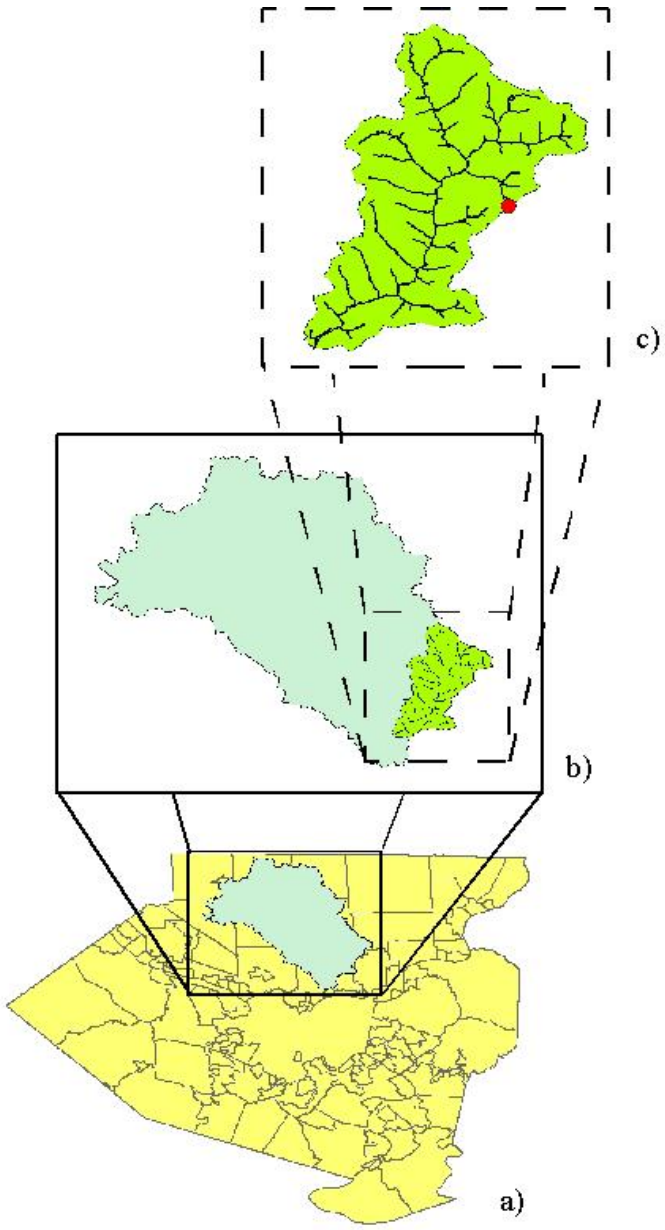

Figure 39. (a) Pine Creek watershed, located in northern Allegheny County, PA. (b) East Little Pine Creek watershed, a tributary of Pine Creek. (c) Beechwood Farms Nature Reserve (red dot) located along the eastern edge of East Little Pine Creek watershed.

6.1.1 Land cover and soil characterization

The soil type over the test bed area was found using the USDA Web Soil Survey (http://websoilsurvey.nrcs.usda.gov). The monitored areas of sites 2 and 3 are classified as Gilpin-Upshur complex with slopes ranging from 8–25%. Gilpin-Upshur complex is classified as a well-drained soil and consists of 25.2% sand, 52.2% silt, and 22.7% clay (Pennystone, 2012). The map unit for this area shows roughly 45% Gilpin silt loam and 35% Upshur silty clay loam. During sensor installation, the soil profile was observable to approximately 30 cm depth. Site 2 (along the hill slope) has a thick carpet of fibrous roots sitting on top of approximately 30 cm of dark silt loam due to heavy leaf litter. Site 3 (at the top of the hill) has a thinner layer of dark brown silt loam (approximately 10–12 cm) which is mixed with a variety of rocks ranging in size from 2–20 cm in diameter. Underneath the silt loam is a layer (at least as deep as 30 cm) of brown silty clay loam.

The United States Geological Survey (USGS) NLCD 2006 map of the ELPC watershed is shown in Figure 40. Based on the image classifications, the area of site 2 and site 3 is mainly deciduous forest with some developed open space.

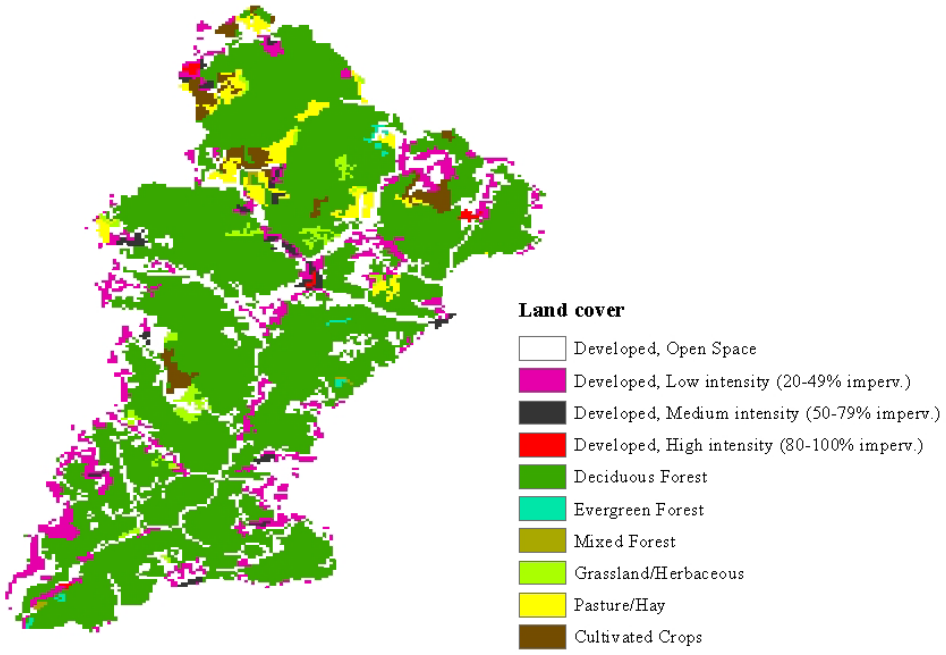

Figure 40. 2006 National Land Cover Data (30 m × 30 m) for the East Little Pine Creek watershed.

6.2 OVERVIEW OF ENVIRONMENTAL SENSORS

To study the three different variables, three different sensors were used in the ASWP test bed: the EC-5 soil moisture sensor, the MPS-1 soil water potential sensor, and the Granier-style thermal dissipation sap flow sensor. For a description of the design, construction and testing of the sap flow sensor see Chapter 2.0 . The EC-5 and MPS-1 sensors are both from Decagon Devices, Inc. The EC-5 is a capacitance-style soil moisture sensor. The EC-5 sensor is rapidly charged and discharged using a 70 MHz frequency to create an electromagnetic field. This magnetic field uses the surrounding soil as a capacitor. The charge time of the capacitor (i.e., the soil) is related to the dielectric permittivity of the soil. The dielectric permittivity of the soil

(which is made up of air, solids, and water) mostly responds to changes in the water content of the soil (i.e., air and solids have low dielectric permittivity). This allows the dielectric permittivity to be related to soil moisture and therefore, the capacitance of the soil can be related to soil moisture. This is done using a linear calibration equation. Equation 6.1 shows the conversion equation for mineral soils as calibrated by Decagon Devices, Inc. which is used in this study.

Equation 6.1

$$SWC = 0.00119 \cdot (mV) - 0.401$$

where: SWC = volumetric soil water content, $m^3 \cdot m^{-3}$

mV = raw sensor measurement using a 2.5 V excitation, mV

Using this equation, sensor measurements ranging from 337–841 mV will result in values between 0–60% volumetric water content. That is assumed acceptable because the upper range (i.e., 0.60 $m^3 \cdot m^{-3}$) is above the expected porosity of the test site soil based on the soil type.

The MPS-1 also measures the dielectric permittivity, but rather than using the soil as a medium it uses a ceramic disc. Moisture travels from the soil to the disc and the dielectric permittivity is measured inside the disc, rather than within the soil. The MPS-1 only measures the matric potential of the soil, which involves the attraction of water to the soil particles. It does not consider the pressure, gravitational, or osmotic potential. Measurements are then related to water potential using an exponential conversion equation. Equation 6.2 shows the relationship between the sensor measurements and the soil water potential, as calibrated by Decagon Devices, Inc., used in this study.

Equation 6.2

$$\psi = -\exp\left(0.000048 \cdot (mV)^2 - 0.0846 \cdot (mV) + 39.45\right)$$

where: ψ = matric soil water potential, kPa

mV = raw sensor measurement using 2–5 V excitation, mV

Using this equation, sensor measurements ranging from 591–841 mV produce soil water potential values from -500 to -10 kPa, which is the range of valid measurements for this sensor.

The range of validity for both the MPS-1 and EC-5 sensors was determined based on the sensor specific conversion equations. The MPS-1 measurements of water potential are valid from -500 to -10 kPa which corresponds to measurements between 591 and 841 mV (see Figure 41). The EC-5 measurements of volumetric water content are valid from 0 to 0.6 $m^3 \cdot m^{-3}$ which corresponds to measurement voltages ranging from 337 to 841 mV (see Figure 41).

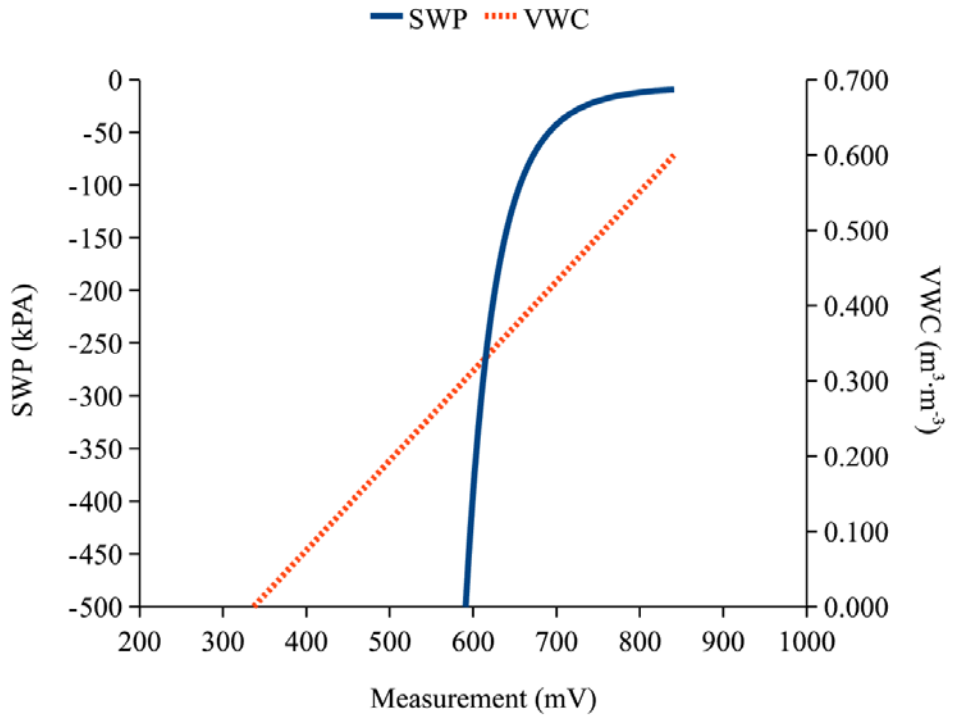

Figure 41. Conversion of measurement voltage (mV) to units of soil water potential (kPa) and volumetric water content ($m^3 \cdot m^{-3}$) in the wireless network.

6.3 INSTALLATION

6.3.1 Soil moisture sensor and soil water potential sensor installation

There are 12 nodes in site 2 and 13 in site 3 that are outfitted with both soil moisture and soil water potential sensors. Each node has three soil sensors: two EC-5 sensors and one MPS-1. EC-5 sensors are located at 30 cm and 10 cm and the MPS-1 is located at 30 cm. The reason there is not an MPS-1 at 10 cm is due to budgetary concerns. (At the time these sensors were purchased, in April 2009, the cost per EC-5 sensor was $65 whereas the MPS-1 cost $105 per sensor.)

Holes were dug or augured at each location to the appropriate depths. All three sensors were placed in the same hole in the ground. Placing the soil moisture sensors in the same hole allowed for monitoring the vertical soil profile. The EC-5 sensors were installed horizontally into the sidewall of the augured hole with the prongs vertically aligned.

6.3.2 Sap flow installation

Traditional methods of sap flow monitoring use large swaths of reflective insulation wrapped around sections of tree trunks to protect the sensors from radiant heating. This was unacceptable due to the second provision of this study. Attempts were made at designing a small enclosure that wrapped partly around the tree trunk to provide protection to the sensors while minimizing visibility. Early attempts were unsuccessful because they were not waterproof. Waterproofing was necessary, since the enclosures would contain sensitive electronics and wiring. Because of this need, waterproof enclosures were purchased. These enclosure types include: Bud Industries Inc. polycarbonate enclosure (part no. PN-1337) and McMASTER-CARR polycarbonate enclosures (part no. C9945K126 and part no. 69945K128). Four configuration designs were developed using these three enclosures types (see Figure 42). Each of these configurations used a combination of two enclosures (one to house the wireless node and one to house the sap flow sensors) and were attached directly to the trees being monitored.

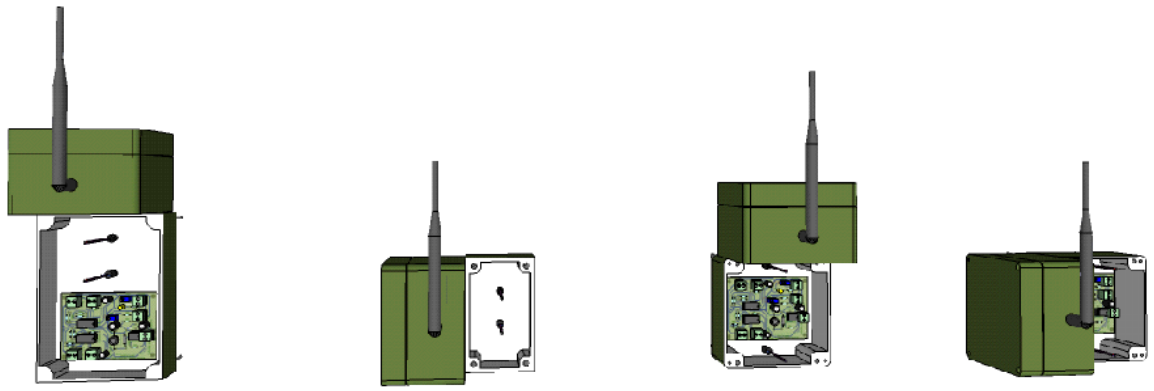

Figure 42. The four designs for waterproof enclosures, which consist of a mote box shown with attached antenna and a sap flow box shown with examples of sensors and circuitry inside. (a) Mote box (PN-1337) and sap flow box (K128). (b) Mote box (PN-1337) and sap flow box (K126). (c) Mote box and sap flow box (PN-1337) vertically aligned. (d) Mote box and sap flow box (PN-1337) horizontally aligned.

Twenty-two individual trees were chosen for monitoring within the study area (10 in site 2 and 12 in site 3). Before sensor installation could occur, a portion of the tree's bark was removed to create a flat surface on the tree. This was done using a drawknife. The flat surface was sized to fit the enclosure box design for the sap flow sensor (Figure 42). The goal was to make sure the enclosures fit securely against the side of the tree. Enclosures were attached to trees using 1-in wood screws to allow for easy removal of the box. Depending on the enclosure size, holes were drilled between 6 and 10 cm apart (vertically aligned) for the sap flow probes to fit into the trees.

The thermal dissipation sap flow sensors used in this study consist of two probes (see Chapter 2.0). The first is a reference temperature probe, which is installed in the upstream

location (i.e., the bottom hole in the enclosure) and the second is the constant heat probe, which was installed in the downstream location (i.e., the top hole in the enclosure). Before installing the two probes, the holes that were drilled were first cleaned using hydrogen peroxide. This is done to prevent introducing a disease or infection into the tree.

Probe needles were greased using petroleum jelly to ease the probe insertion and removal at the end of the monitoring season. This also helped maintain the contact between the probe and the wood of the tree. The individual probes were glued in place using a silicone adhesive, which helped keep the probes stationary and prevented water from intruding into the box. Once secured, the probes were wired to their circuit board for collecting measurements.

6.4 DATA ANALYSIS

6.4.1 Soil sensor analysis

This section analyzes soil sensor data (soil water potential and soil moisture content). There are three types of analyses in this section. The first analysis will look at seasonal spatial variations of the soil sensor measurements located in sites 2 and 3. The second analysis examines the impact of duplicate packets, generated in the wireless network, on measurements taken on both types of soil sensors in sites 2 and 3. The third analysis is a comparison of the time series measurements between the WSN and data loggers.

The data used in the spatial analysis and duplicate packet analysis is similar to the data studied in Chapter 5.0 , with some notable differences. First, the time periods examined vary. In Chapter 5.0 the study period was from May 2010 to May 2012. Because soil sensors were not

installed at the ASWP test bed until late July 2010 (when nodes in site 2 were outfitted) to early August 2010 (when nodes in site 3 were outfitted), the data for this analysis will start from the time that the soil sensors were installed in site 2 and end in May 2012. Another difference between this and the Chapter 5.0 data is that this study only examines nodes that have environmental sensors, which includes 25 nodes (12 in site 2 and 13 in site 3). The time period examined for the time series analysis corresponds to the time period when data loggers were installed in sites 2 and 3 (April 2012 to September 2012).

6.4.1.1 Spatial analysis of soil moisture and soil water potential measurements

To examine the spatial variability of the soil moisture and soil water potential over the study region, measurements were seasonally averaged with respect to the USGS water year, which begins on the first day in October (see Table 19). Note that the total number of lines of data (i.e., packets) processed in each season is included as a reference to periods when data was limited (see Figure 36). In Table 19 there is an apparent seasonal fluctuation in the number of received packets with highs during the autumn months and lows during the spring months.

Table 19. Seasonal periods used for averaging soil sensor data with their respective dates and total packets processed.

	Season	Dates	No. of Packets
1	Summer 2010	1 Jul 2010–30 Sep 2010	38388
2	Autumn 2010	1 Oct 2010–31 Dec 2010	44454
3	Winter 2011	1 Jan 2011–30 Mar 2011	25359
4	Spring 2011	1 Apr 2011–30 Jun 2011	16325
5	Summer 2011	1 Jul 2011–30 Sep 2011	27851
6	Autumn 2011	1 Oct 2011–31 Dec 2011	43415
7	Winter 2012	1 Jan 2012–30 Mar 2012	24739
8	Spring 2012	1 Apr 2012–30 Jun 2012	13894

The raw soil sensor measurements (mV) were converted to their respective units of measure using Equation 6.1 and Equation 6.2. Calculations were only made on valid measurements. Erroneous measurements occur whenever the mote battery voltage drops below a certain threshold, resulting in measurement values of 1250 mV or when a sensor/data acquisition board malfunctions resulting in measurements of either 0 mV or 2499 mV. Therefore, valid measurements were limited to the range discussed in Section 6.2.

For the purpose of visualization, a grid was placed over the study regions of site 2 and 3. Each grid cell covers an area of 1.5 m^2. An algorithm was used to select grid cells located within 0.5 m of each node location. Node locations were geo-located using a Topcon GR-3 GPS device. Each node may be designated between one and four grid cells depending on how it is located within the grid. Table 20 shows the number of grid cells that have been allocated to each of the 25 sensor nodes.

Table 20. Sensor node IDs and the number of grid cells allocated to each.

Site 2 Nodes	No. of Cells	Site 3 Nodes	No. of Cells
2003	3	3005	4
2015	4	3015	3
2025	4	3025	3
2045	2	3035	4
2055	4	3045	4
2065	4	3055	4
2085	4	3065	3
2095	4	3075	2
2103	3	3085	4
2115	1	3095	4
2125	4	3105	2
2135	4	3113	4
		3125	4

The following 24 figures (Figure 43 to Figure 66) highlight the spatial variability of the seasonally-averaged soil water potential and volumetric soil moisture measurements. The node ID labels have been removed from the figures to minimize the clutter of the images. The node locations are the same as defined in Figure 31. The orientation of the field has been rotated with north facing upwards, such that site 2 nodes are located in the bottom left and site 3 nodes are located in the top right of the figures.

The variation in soil water potential at 30 cm over eight seasons from 2010–12 is presented in the following figures (Figure 43 to Figure 50).

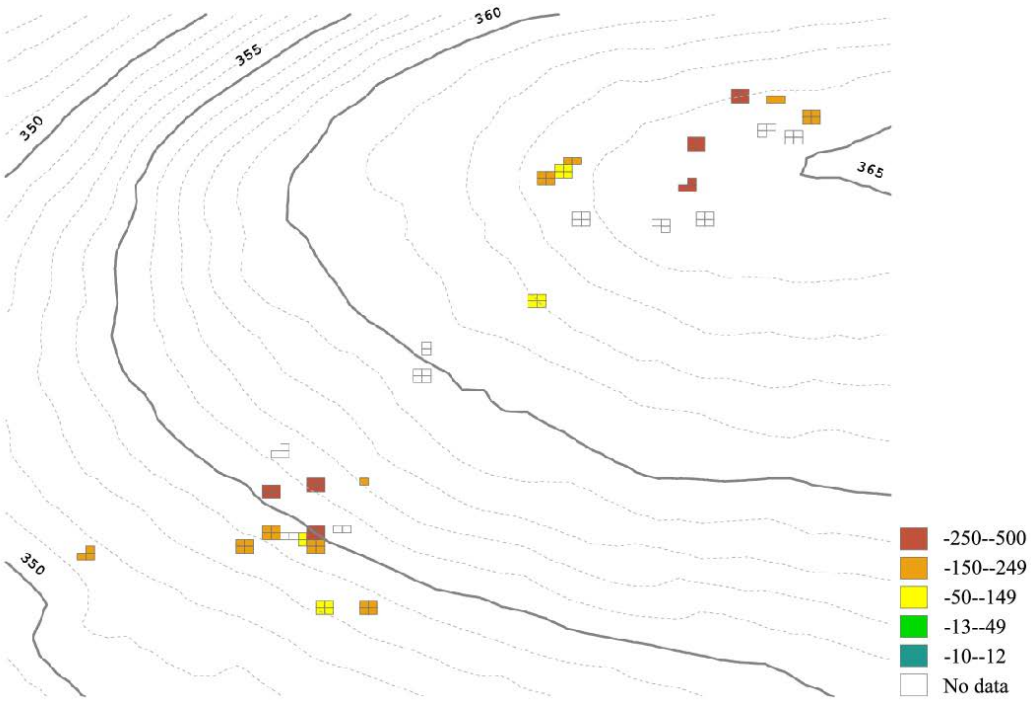

Figure 43. Spatial variation of soil water potential during the summer of 2010. Water potential measurements are at 30 cm and range from -10 to -500 kPa. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements. Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

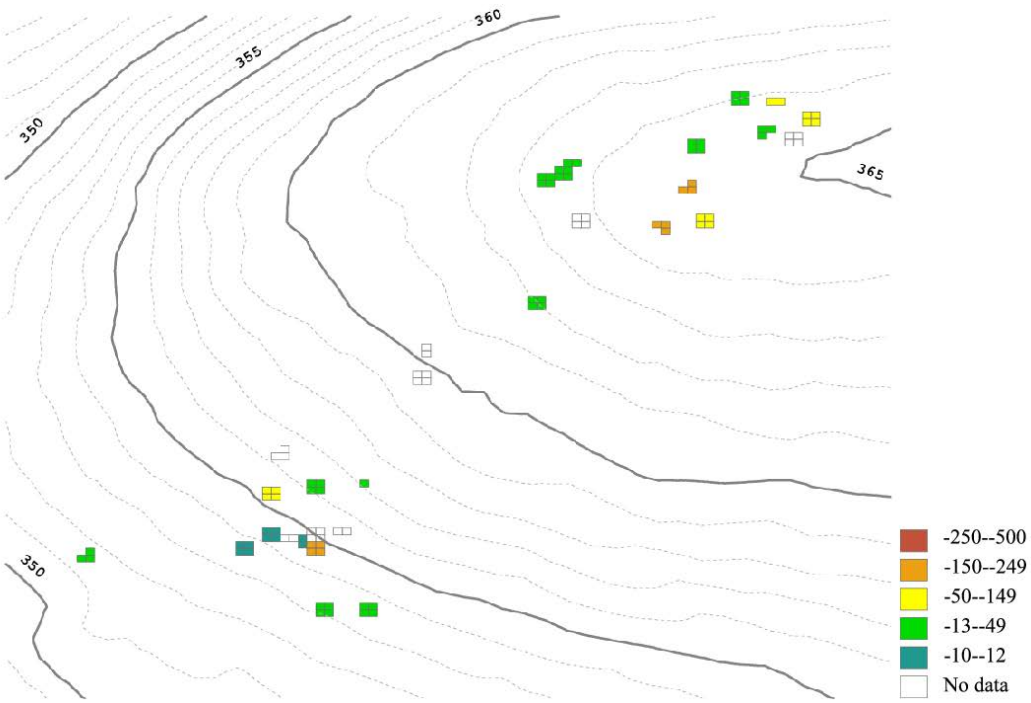

Figure 44. Spatial variation of soil water potential during the autumn of 2010. Water potential measurements are at 30 cm and range from -10 to -500 kPa. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements. Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

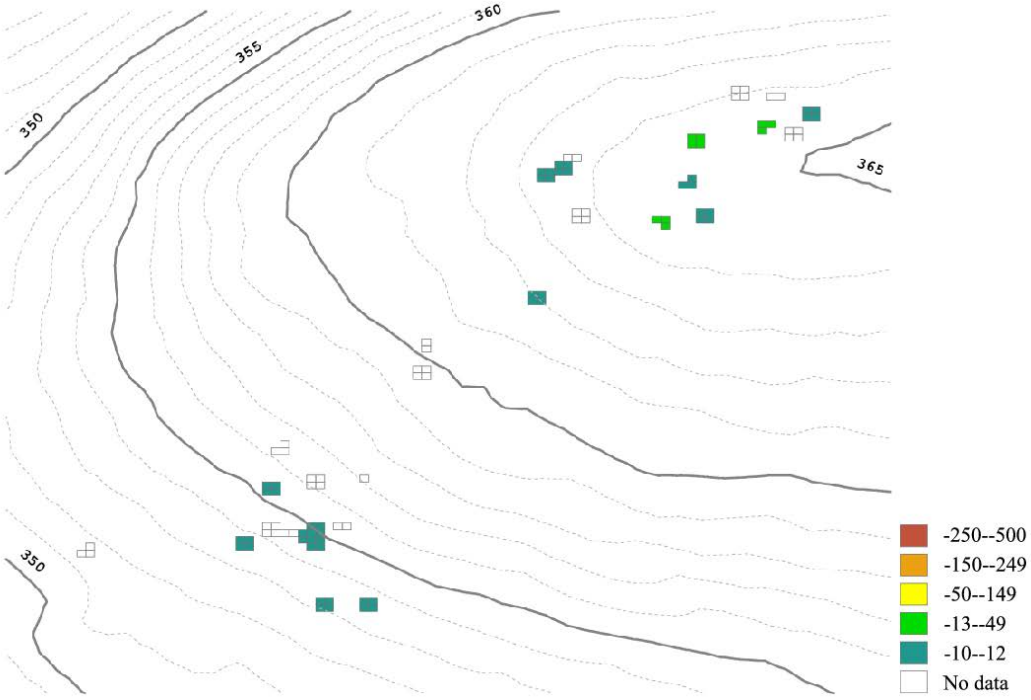

Figure 45. Spatial variation of soil water potential during the winter of 2011. Water potential measurements are at 30 cm and range from -10 to -500 kPa. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements. Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

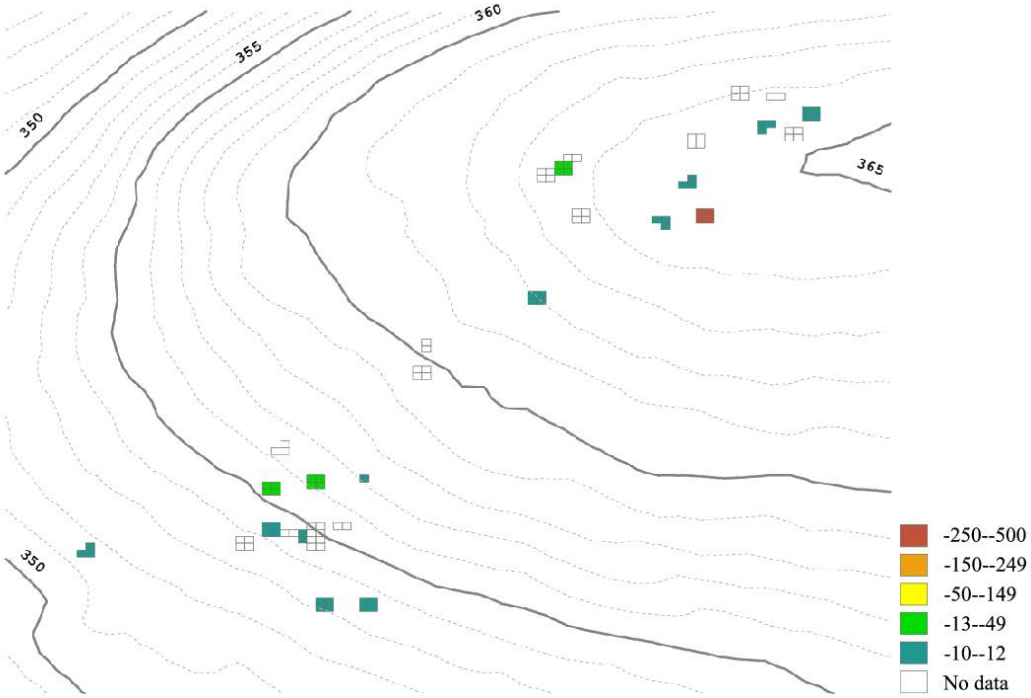

Figure 46. Spatial variation of soil water potential during the spring of 2011. Water potential measurements are at 30 cm and range from -10 to -500 kPa. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements. Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

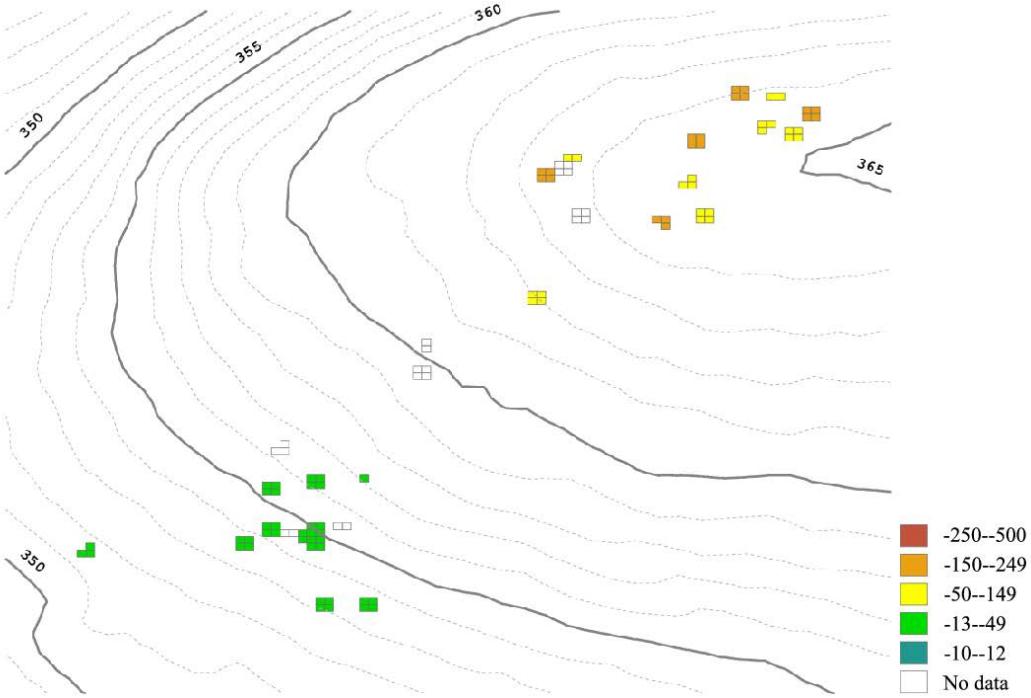

Figure 47. Spatial variation of soil water potential during the summer of 2011. Water potential measurements are at 30 cm and range from -10 to -500 kPa. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements. Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

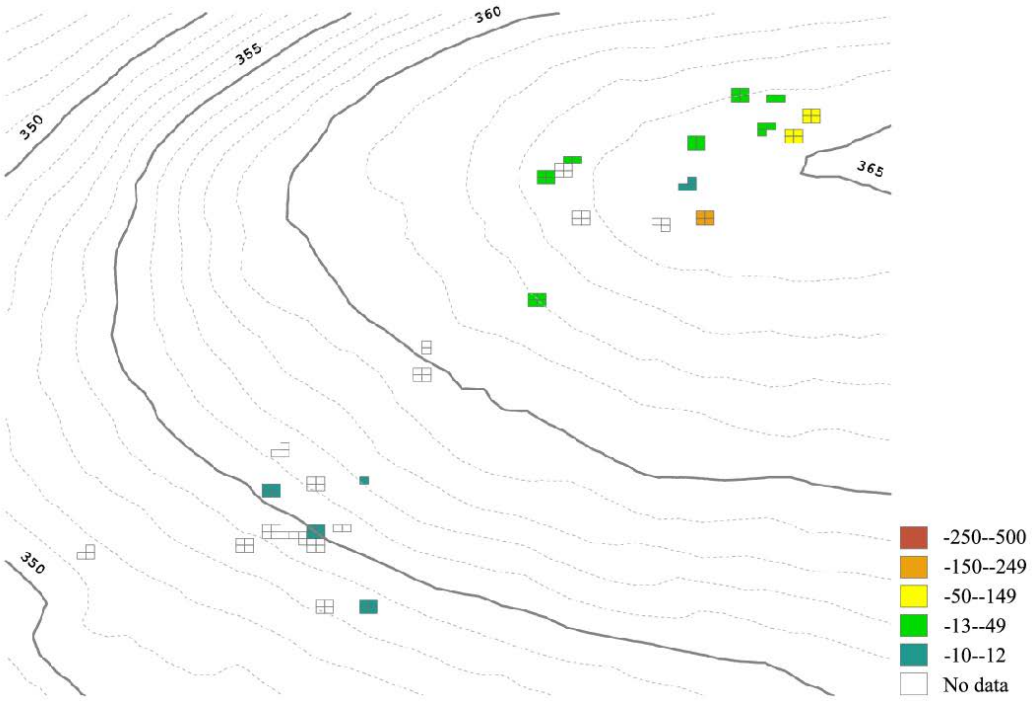

Figure 48. Spatial variation of soil water potential during the autumn of 2011. Water potential measurements are at 30 cm and range from -10 to -500 kPa. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements. Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

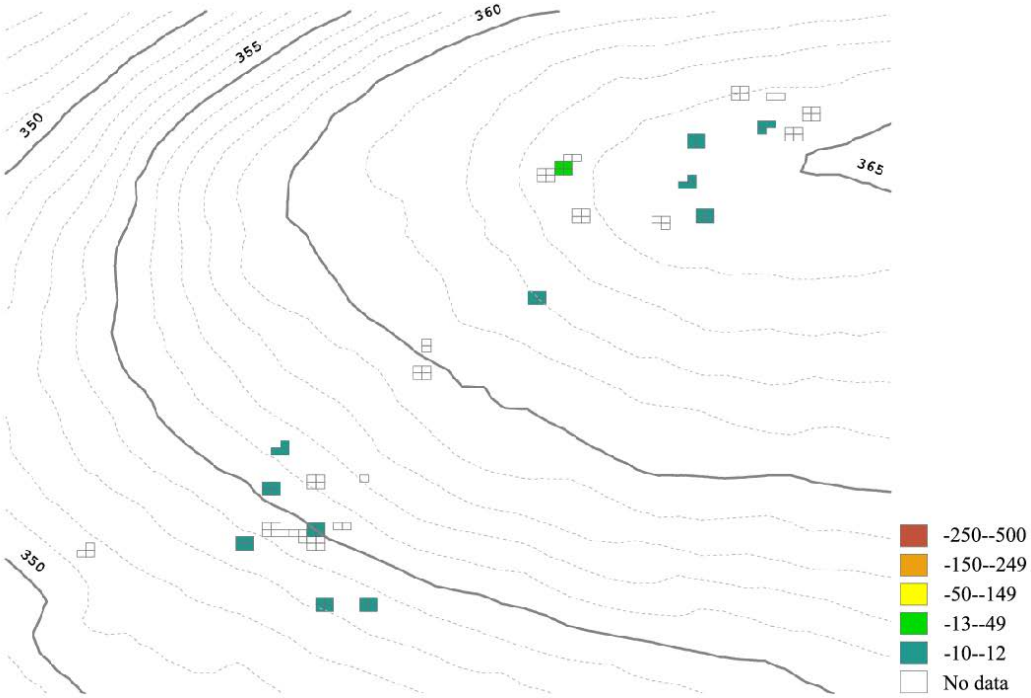

Figure 49. Spatial variation of soil water potential during the winter 2012. Water potential measurements are at 30 cm and range from -10 to -500 kPa. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements. Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

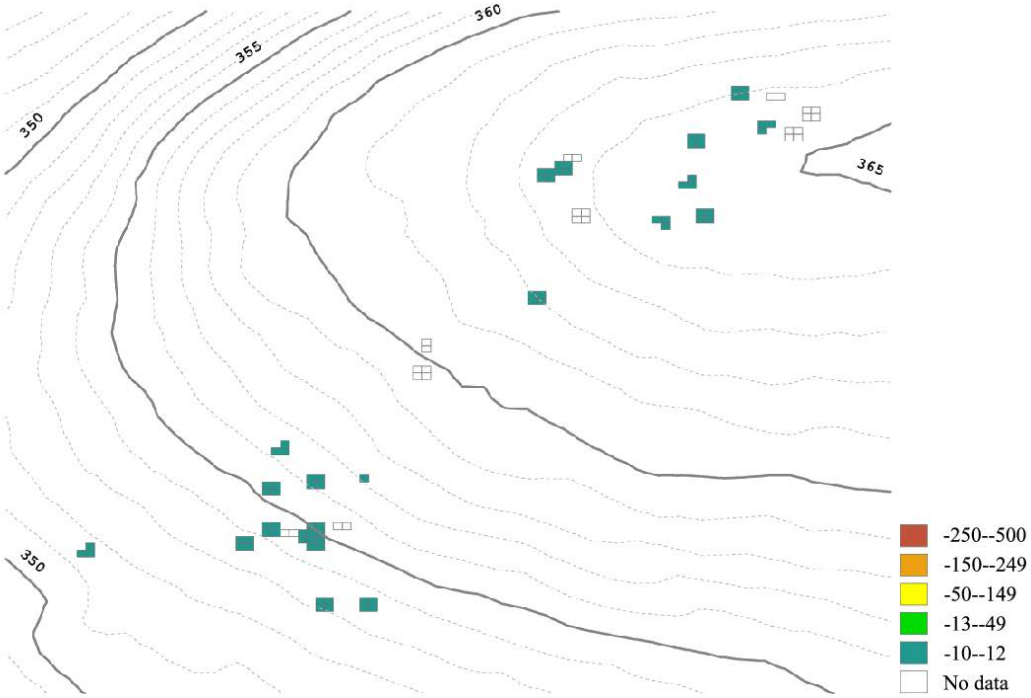

Figure 50. Spatial variation of soil water potential during the spring of 2012. Water potential measurements are at 30 cm and range from -10 to -500 kPa. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements. Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

After examining the eight seasons of soil water potential at 30 cm across sites 2 and 3, seasonal trends can be seen. Starting in the summer of 2010, soil water potential is low, which represents dry soils across sites 2 and 3 (Figure 43). The summer of 2010, which received only 210 mm of rainfall, represents a dry summer and is reflected in the dry conditions. As compared to the summer of 2011, which received 342 mm of rainfall. There is also some distinction between the uphill and downhill measurements in site 2, where downhill locations are slightly wetter than uphill.

In the autumn of 2010, moisture content in both sites 2 and 3 increased (Figure 44). However, there are a few exceptions. Some locations remain dry, mostly near the east to southeast section of site 3. There are also two locations in site 2 that remain dry. Both of these examples are noteworthy because they show the heterogeneity of soil moisture within each site.

In the winter of 2011, all locations in both sites, for which there is valid data, present wet conditions (see Figure 45).

In the spring of 2011, conditions in site 2 remain wet, with some drying near the top of the hill (see Figure 46). One location in site 3 is shown as being very dry. The one very dry area in site 3 appears to be an outlier. The outlier may represent either an extreme heterogeneous location or a monitoring error.

In the summer of 2011, the bottom of the hill in site 2 displays some moderate drying; however, significant drying is seen throughout site 3 (see Figure 47). There is now a spatial differentiation in soil water potential between sites 2 and 3, because site 2 is wet and site 3 is dry.

In the autumn of 2011, only four locations in site 2 represent valid data (see Figure 48). All of which show high levels of moisture. In site 3, the northwest and southern regions get

wetter, while the eastern and southeastern areas remain dry. These are the same regions that are dry in the autumn of 2010.

In the winter of 2012, site 2 has high water potential, which represents wet conditions (Figure 49). The few valid results in site 3 show similarly wet conditions. From winter to spring 2012, wet conditions remain throughout sites 2 and 3 (see Figure 50).

The following figures (Figure 51 to Figure 58) present the spatial variation of volumetric soil moisture at 30 cm.

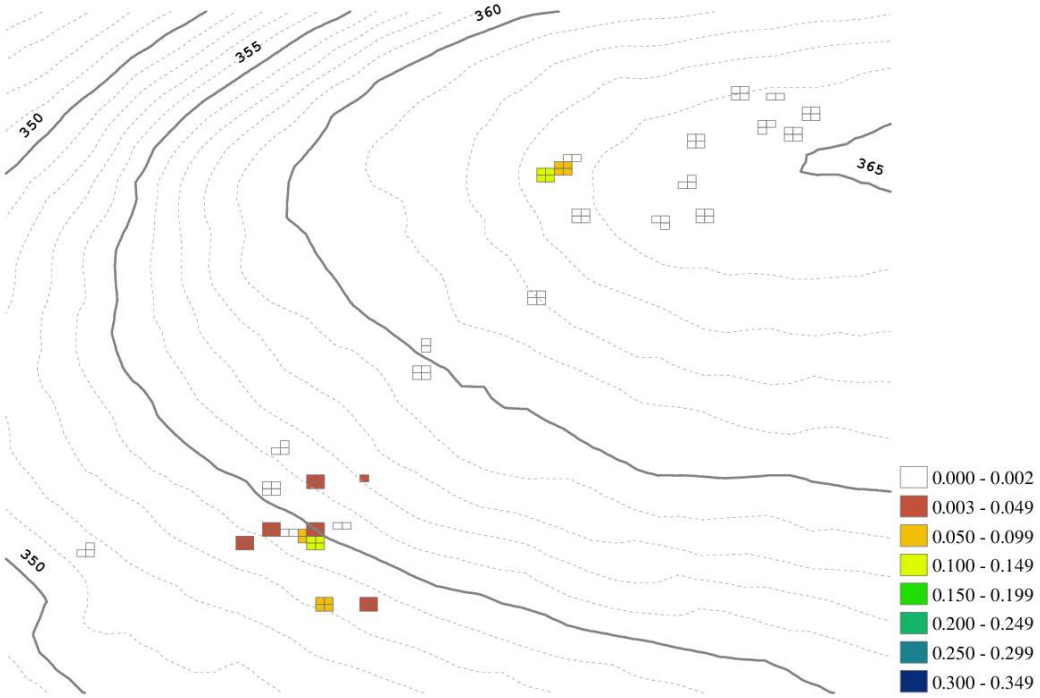

Figure 51. Spatial variation of volumetric water content during the summer of 2010. Volumetric water content measurements are at 30 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

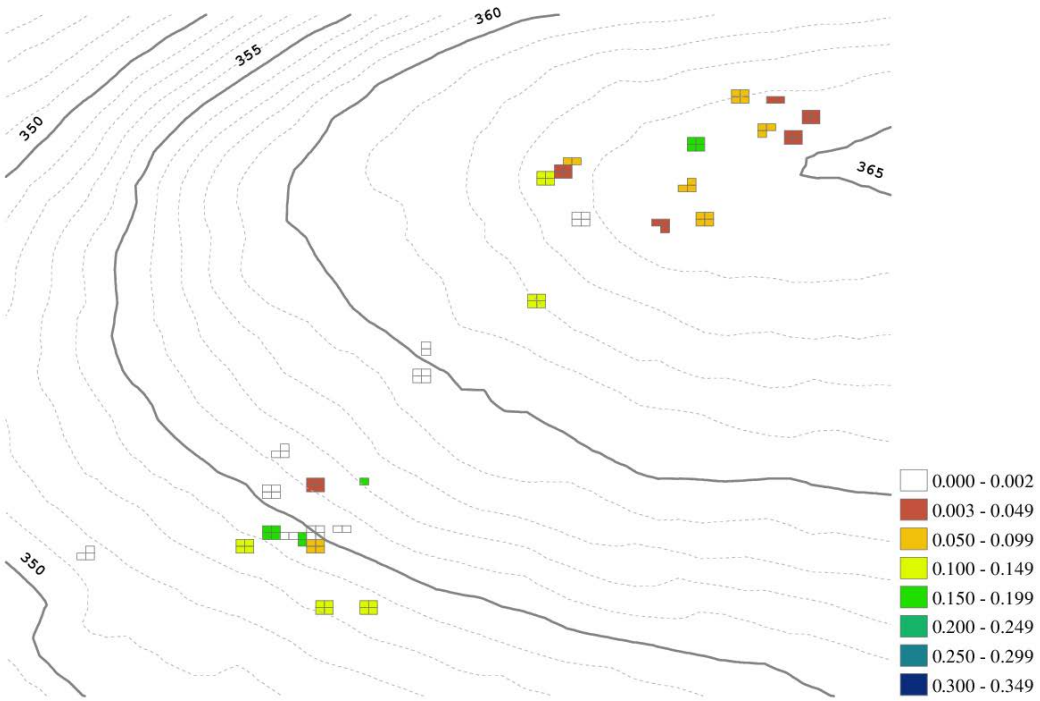

Figure 52. Spatial variation of volumetric water content during the autumn of 2010. Volumetric water content measurements are at 30 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

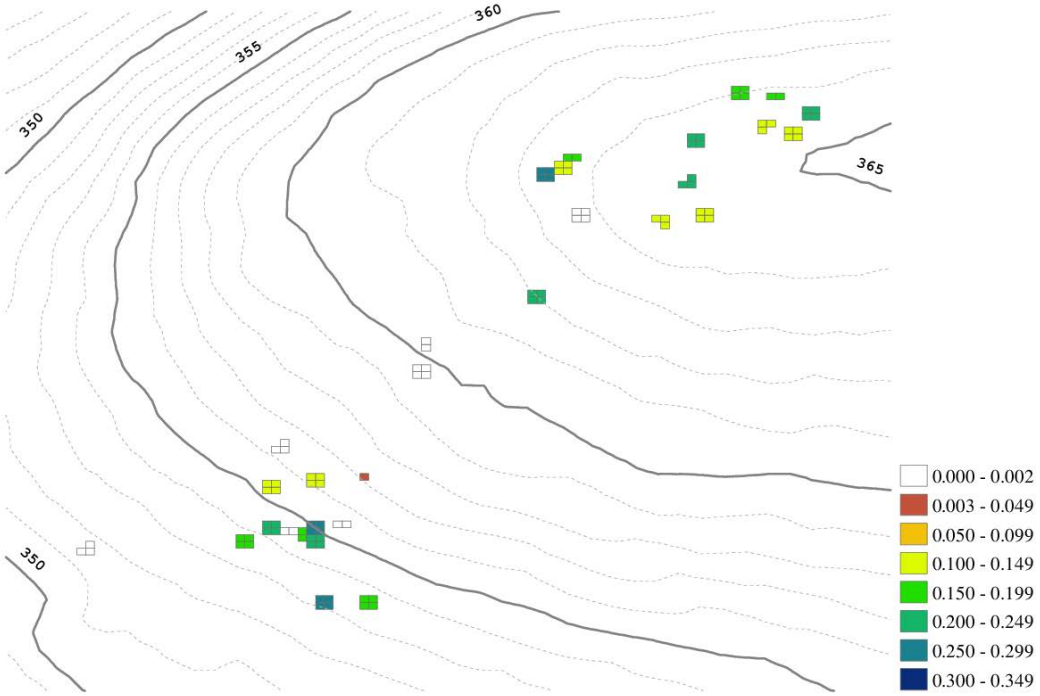

Figure 53. Spatial variation of volumetric water content during the winter of 2011. Volumetric water content measurements are at 30 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

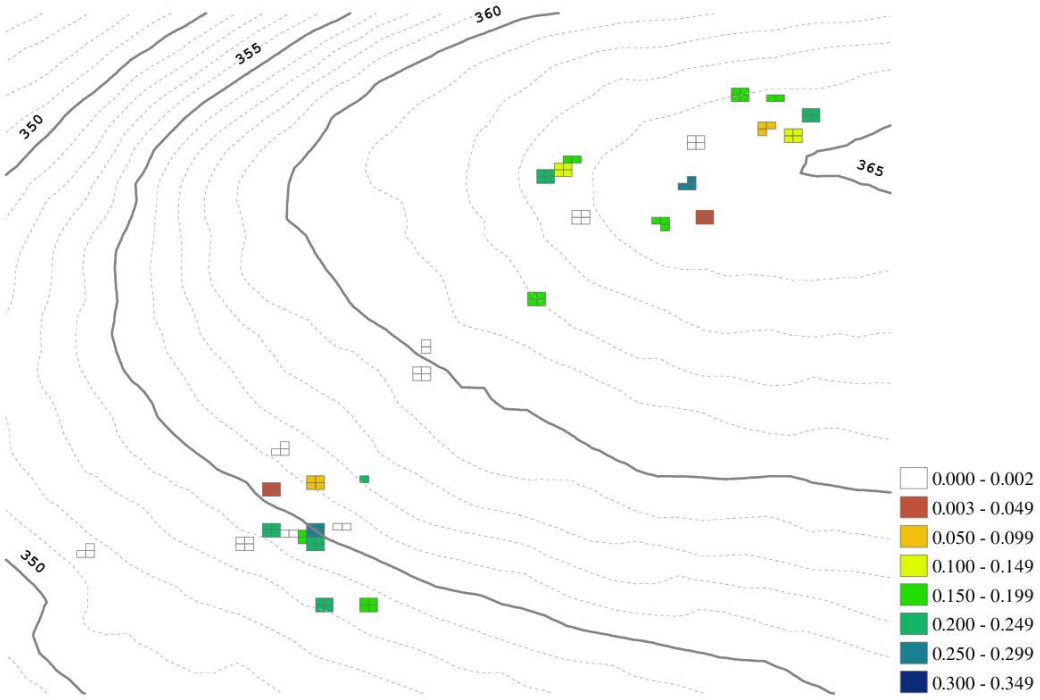

Figure 54. Spatial variation of volumetric water content during the spring of 2011. Volumetric water content measurements are at 30 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

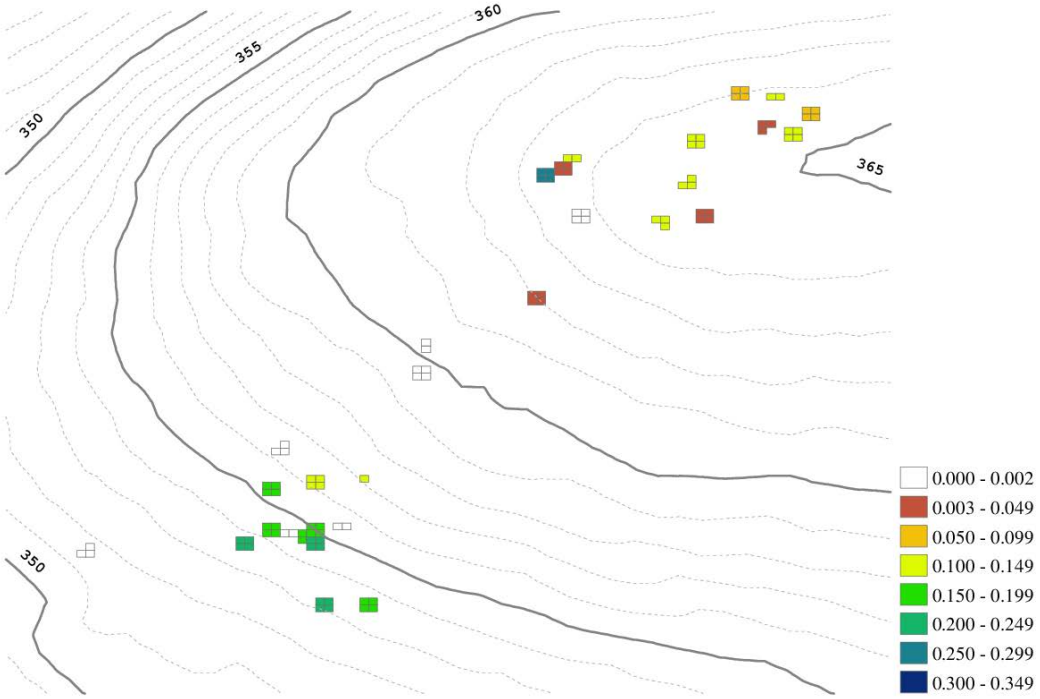

Figure 55. Spatial variation of volumetric water content during the summer of 2011. Volumetric water content measurements are at 30 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

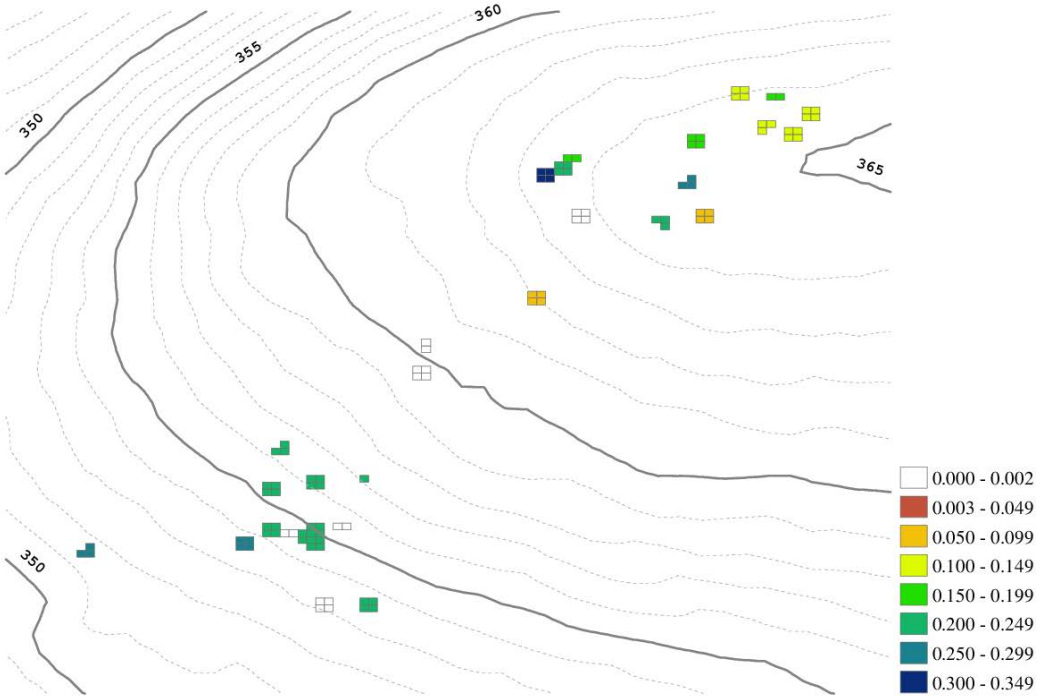

Figure 56. Spatial variation of volumetric water content during the autumn of 2011. Volumetric water content measurements are at 30 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

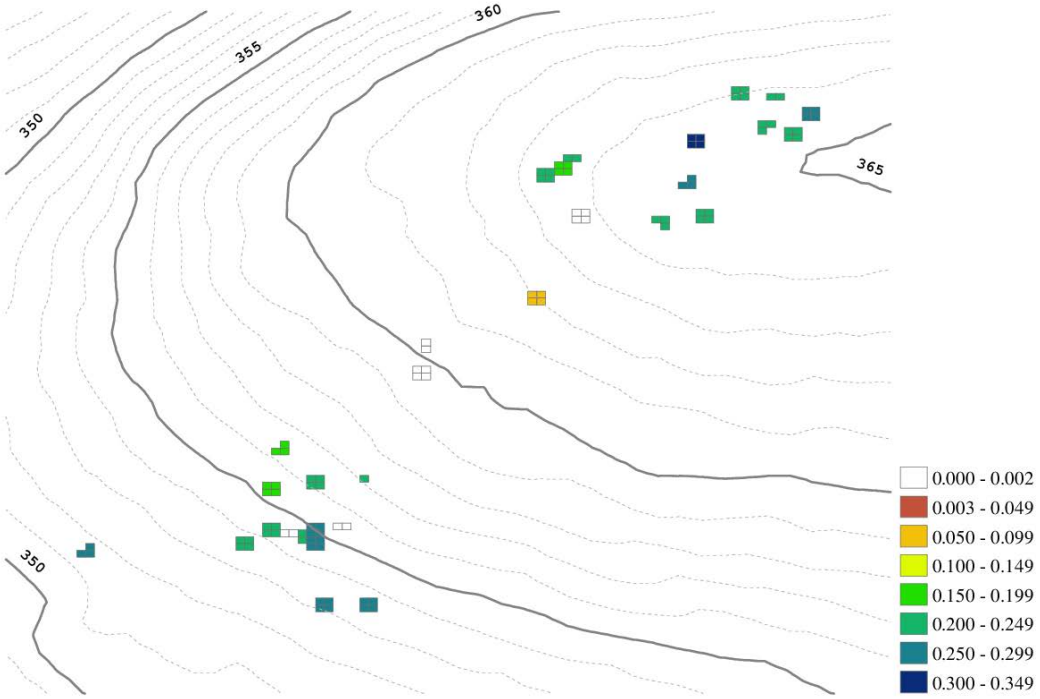

Figure 57. Spatial variation of volumetric water content during the winter of 2012. Volumetric water content measurements are at 30 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

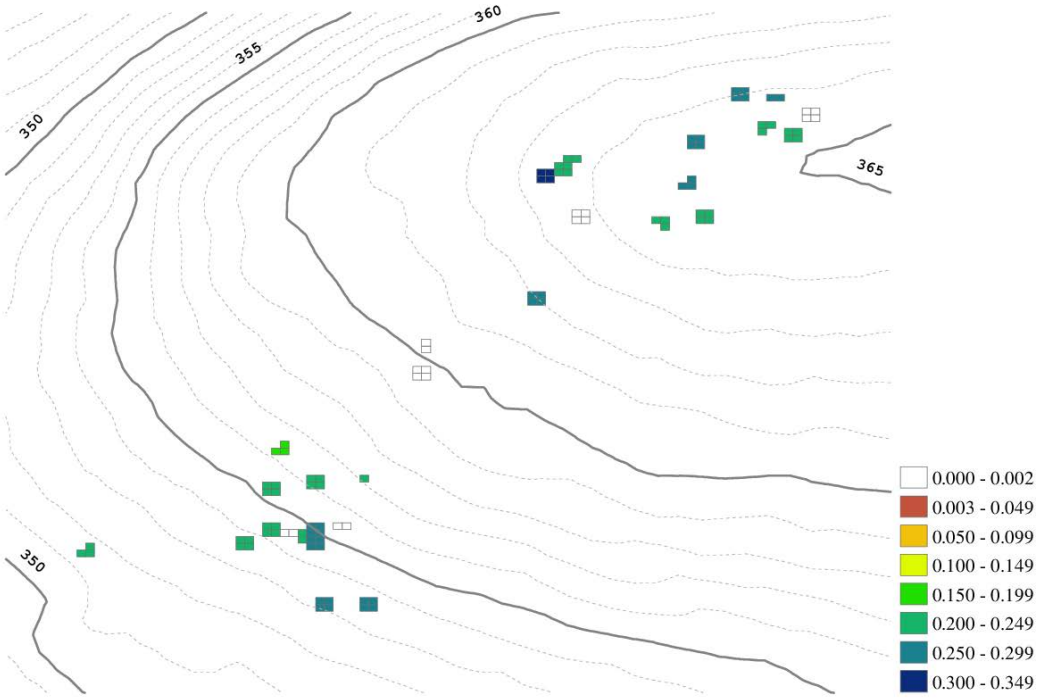

Figure 58. Spatial variation of volumetric water content during the spring of 2012. Volumetric water content measurements are at 30 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

In the summer of 2010, the soil water content in site 2 is low (see Figure 51). Limited data in site 3 presents only two locations with valid measurements, which show moderately dry conditions.

In the autumn of 2010, moderately dry conditions persist at the bottom of the hill in site 2 (see Figure 52). Moderately wet conditions are seen in a few locations near the mid and top of the site 2 hill. Dry conditions remain in one location at the top of the hill in site 2. Site 3 is moderately dry to dry, except in one location in the middle of site 3, which is moderately wet.

In the winter of 2011, some wetting occurs at the bottom of the hill in site 2, while at the top of the hill it remains moderately dry (see Figure 53). In site 3, the northern and western locations are wetter and the eastern and southeastern locations are dryer, with the exception of one location in the west which is moderately dry.

In the spring of 2011, the conditions at the bottom of the hill in site 2 remain wet (see Figure 54). Locations near the top of the hill in site 2 remain dry. In site 3, the spatial distribution of soil moisture remains relatively unchanged from the winter of 2011.

In the summer of 2011, soil moisture conditions in site 2 remain wetter at the bottom of the hill, while conditions throughout site 3 are dry with the exception of one location on the western edge of site 3, which is wet (see Figure 55). This location is an outlier, as it is proximal to another location that is very dry and all other locations within site 3 are moderately dry to dry.

In the autumn of 2011, wet conditions persist throughout site 2, while wet conditions in site 3 are mainly in the western region (see Figure 56). The northeast and southeast of site 3 remain dry.

Conditions in site 2 remain relatively unchanged in the winter of 2012 (see Figure 57). Locations throughout site 3 become moderately wet to wet, with one location in the southwest

that remains moderately dry. In the spring of 2012, wet conditions are present throughout sites 2 and 3 (see Figure 58).

The following figures (Figure 59 to Figure 66) present the variation of volumetric water content at 10 cm during eight seasons over 2010–12.

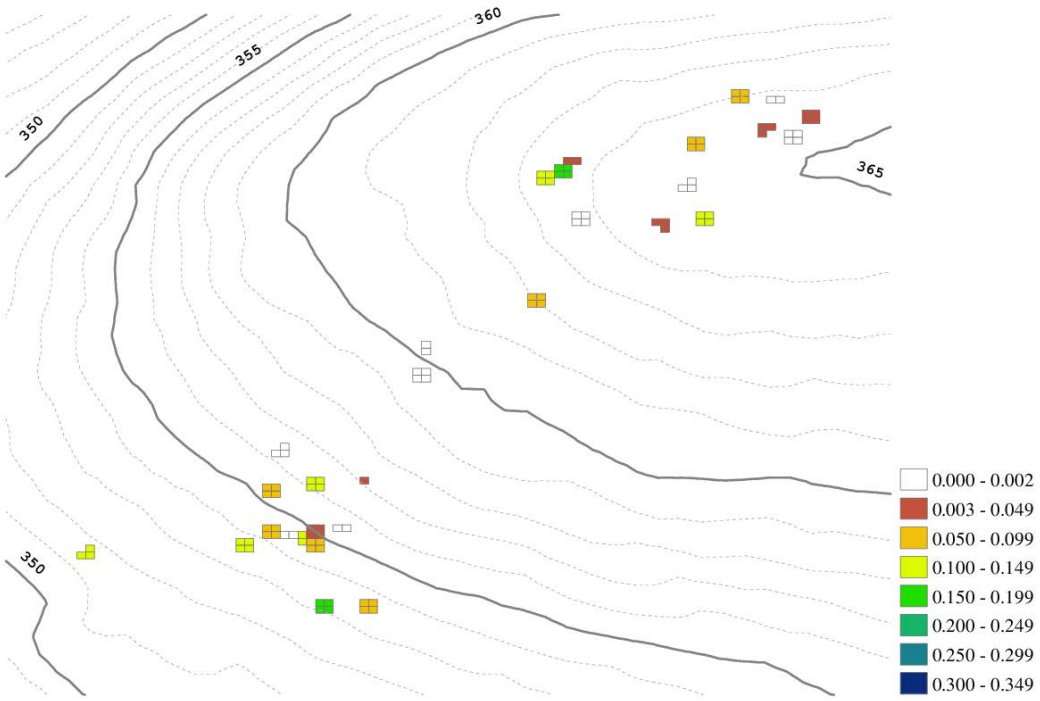

Figure 59. Spatial variation of volumetric water content during the summer of 2010. Volumetric water content measurements are at 10 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

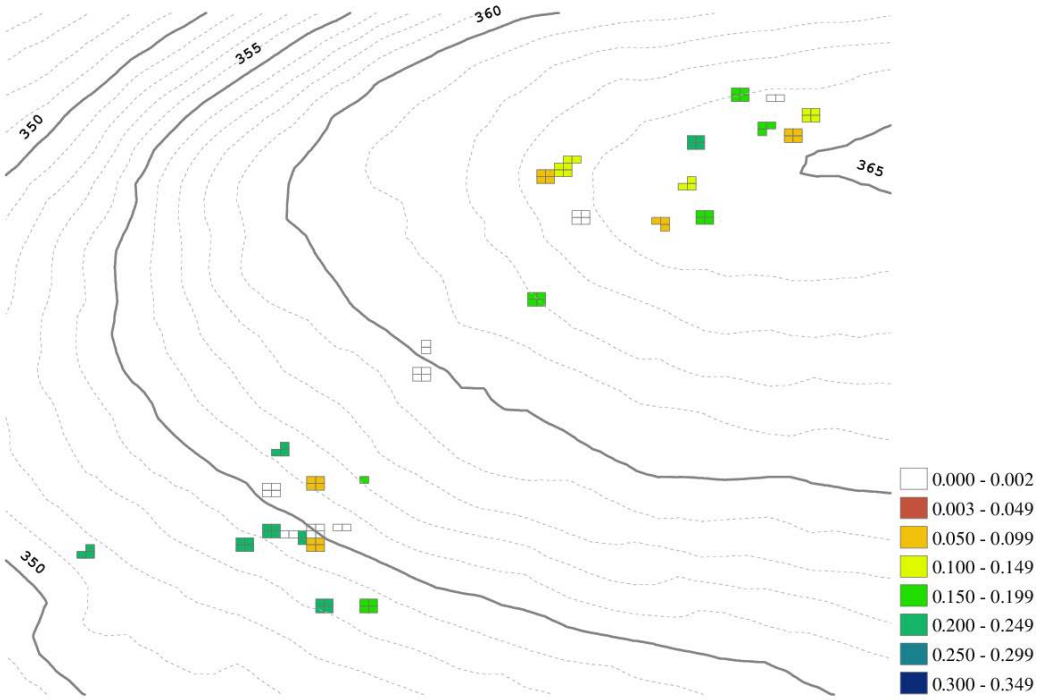

Figure 60. Spatial variation of volumetric water content during the autumn of 2010. Volumetric water content measurements are at 10 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

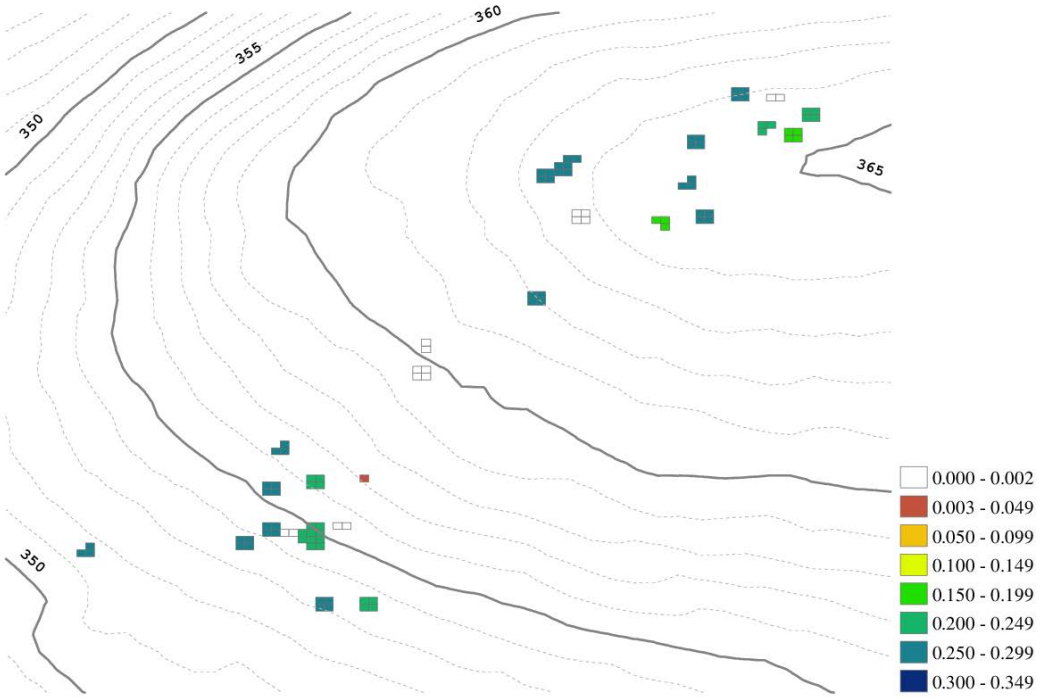

Figure 61. Spatial variation of volumetric water content during the winter of 2011. Volumetric water content measurements are at 10 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

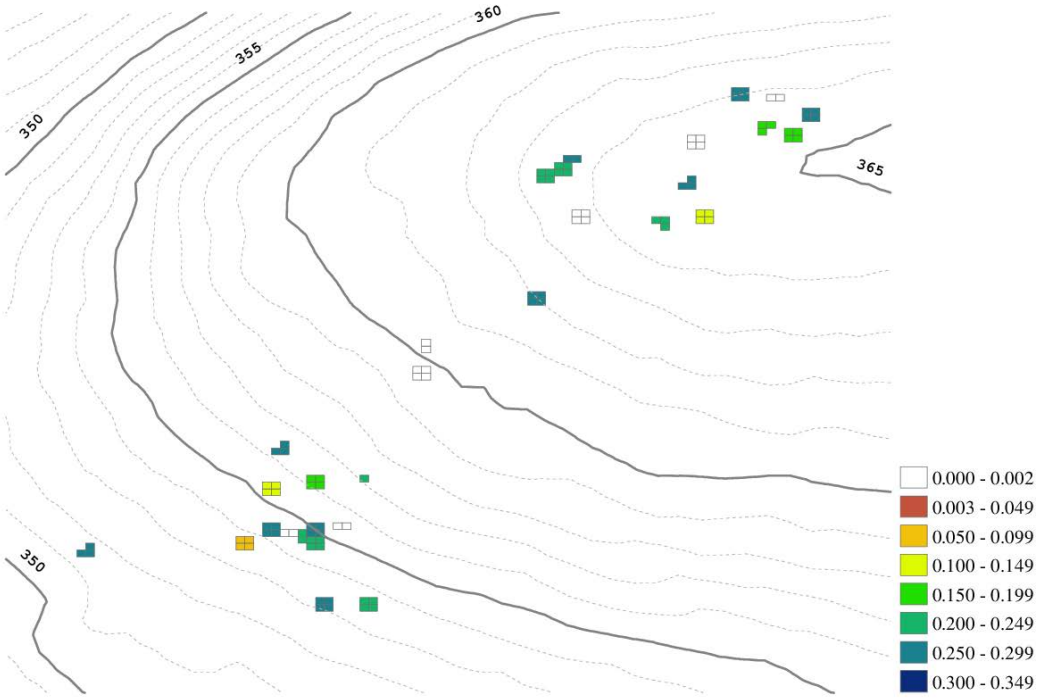

Figure 62. Spatial variation of volumetric water content during the spring of 2011. Volumetric water content measurements are at 10 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

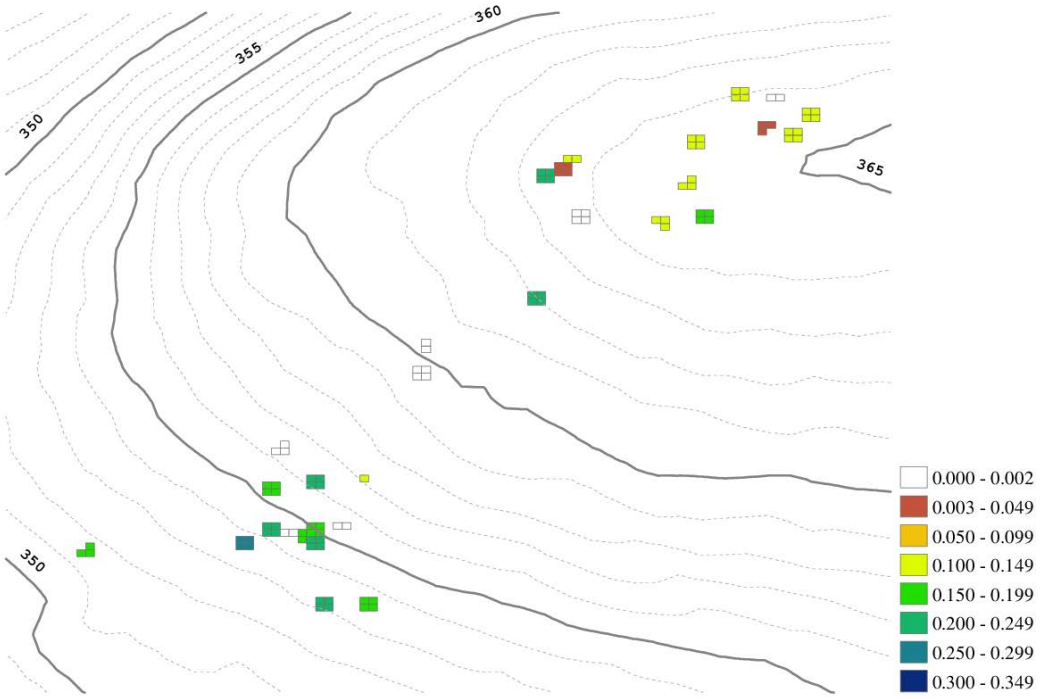

Figure 63. Spatial variation of volumetric water content during the summer of 2011. Volumetric water content measurements are at 10 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

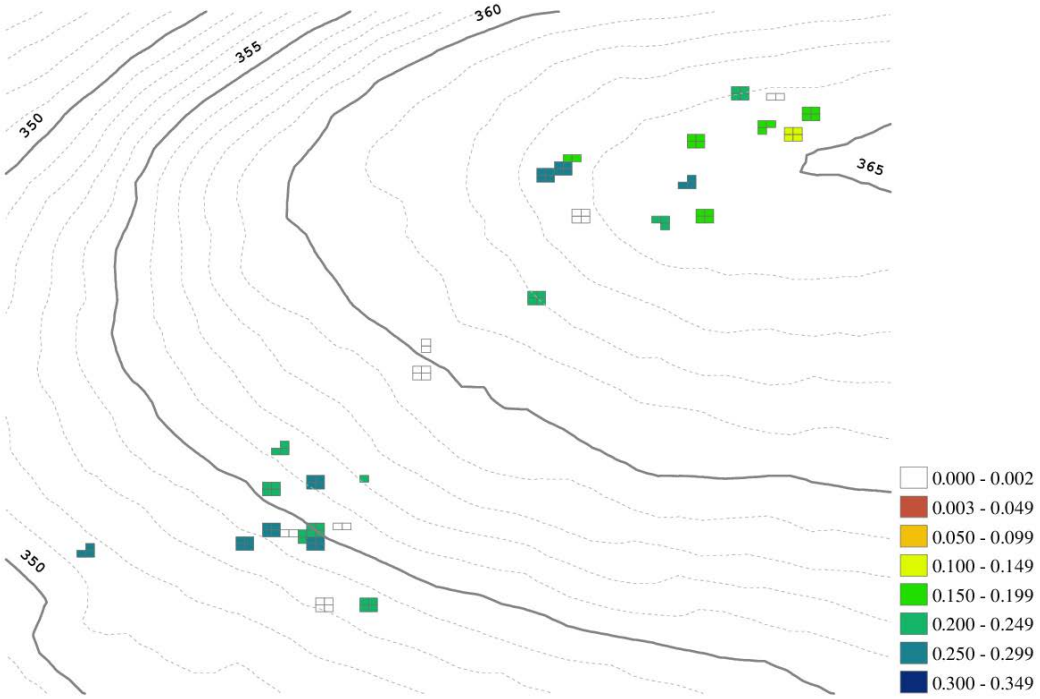

Figure 64. Spatial variation of volumetric water content during the autumn of 2011. Volumetric water content measurements are at 10 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

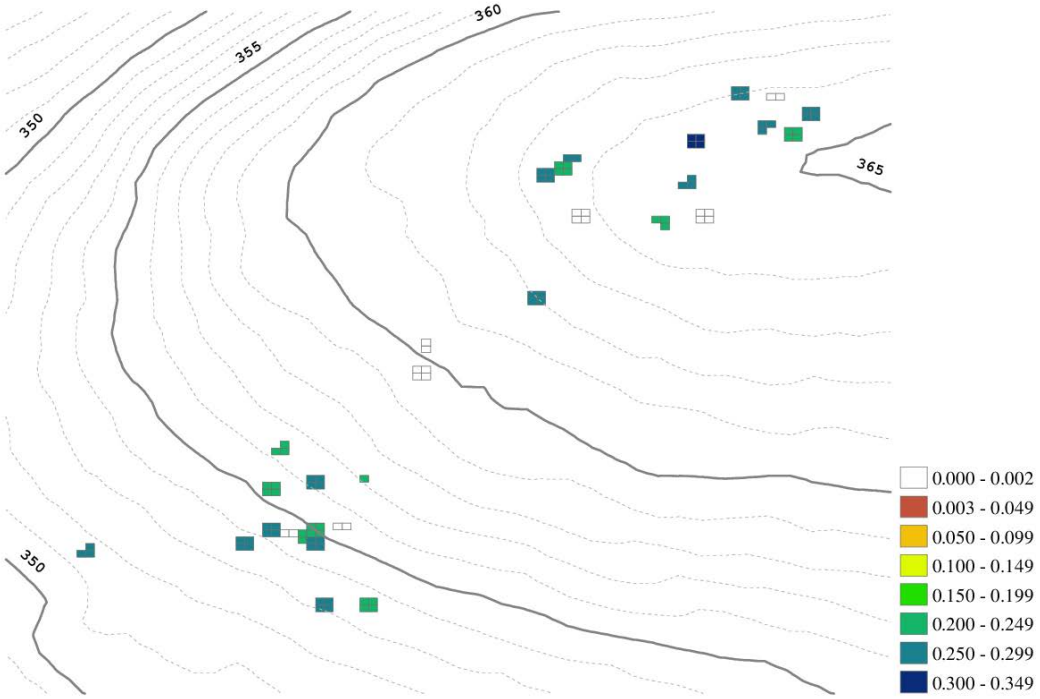

Figure 65. Spatial variation of volumetric water content during the winter of 2012. Volumetric water content measurements are at 10 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

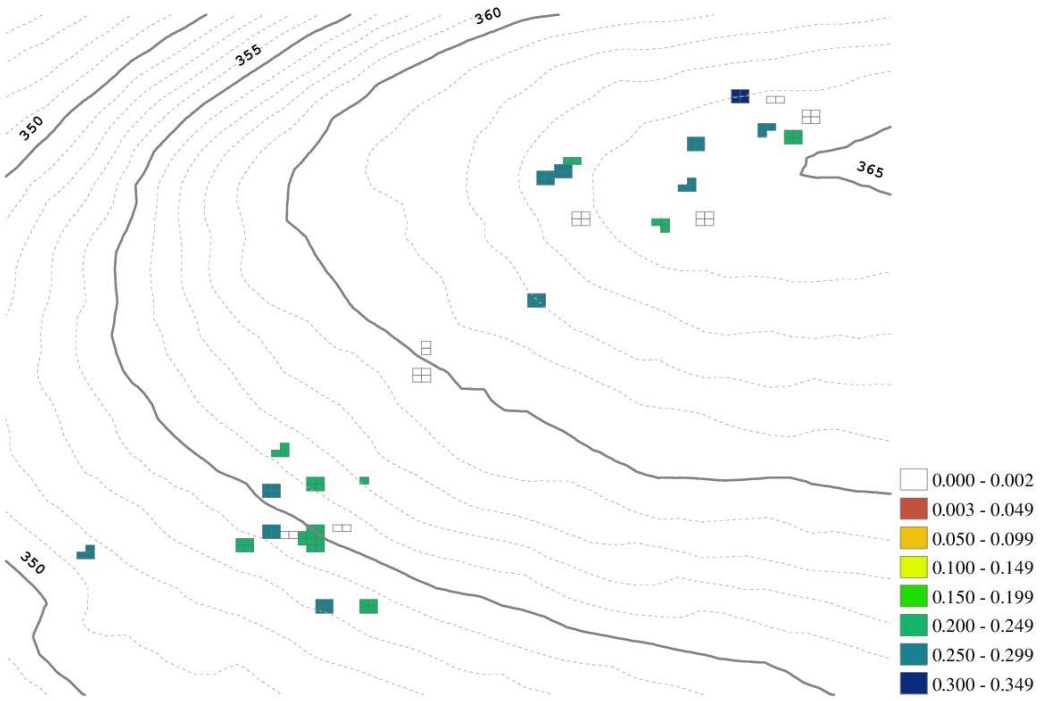

Figure 66. Spatial variation of volumetric water content during the spring of 2012. Volumetric water content measurements are at 10 cm and range from 0.003 to 0.35 $m^3 \cdot m^{-3}$. Color gradation from red to blue indicates regions of dry to wet soil. White grid cells indicate no data or invalid measurements (i.e., 0–0.002 $m^3 \cdot m^{-3}$). Bold contour lines indicate 5 m changes in elevation. Dashed contour lines indicate 1 m changes in elevation.

In the summer of 2010, moisture content throughout sites 2 and 3 is low (Figure 59). Two locations at the bottom of site 2 are the exception: one location represents moderately wet conditions and one location in the western region of site 3 is also moderately wet.

In the autumn of 2010, the bottom of the hill in site 2 and the northeastern area of site 3 became wetter (see Figure 60).

In the winter of 2011, the bottom of the hill in site 2 remained wetter than the top of the hill (see Figure 61). In site 3, the eastern and southeastern regions are dryer.

In the spring of 2011, select locations in site 2 and 3 show signs of drying. This includes the northern regions of site 2 and the southeastern part of site 3 (see Figure 62).

In the summer of 2011, site 2 remains wetter at the bottom of the hill, while areas throughout site 3 are dryer (see Figure 63). Similar to soil moisture at 30 cm (see Figure 55), there is an outlier on the western edge of site 3 which shows wet conditions.

In the autumn of 2011, site 2 is wet and site 3 is wetter than it was in the summer of 2011, with the exception of a moderately dry area in the east (see Figure 64). In the winter of 2012, conditions in site 2 remain as they were in the autumn of 2011; however, wetter conditions are found throughout site 3 (see Figure 65). In the spring of 2012, wet conditions persist throughout sites 2 and 3 (see Figure 66).

6.4.1.2 Duplicate packet analysis

A unique feature of WSN measurements is the presence of duplicated packets of sensor measurements. The cause of duplicate packets is explained in Section 5.2.3.2. The influence of duplicated information on sensor measurements, however, has not yet been investigated. The concern with duplicate packets is the extent of bias they have on the aggregation of data to regular time intervals. For the purposes of analysis, the irregular time series data of received

WSN data packets are often averaged to a regular time series (e.g., daily, weekly, monthly, etc.). Duplicate data can alter the average over these time scales which may affect the results of sensor unit conversion. Critical conditions occur when duplicate packets affect measurements that are made near the boundary threshold values for valid unit conversion (i.e., bias to/from valid range of measurements).

Raw measurements from the 25 sensor nodes were divided into site 2 and site 3 data. Duplicate free sensor measurements were obtained using the algorithm by Navarro et al. (2013) on the original time series measurements. All soil water potential sensor and soil moisture sensor data (original and duplicate-free) in site 2 and site 3, respectively, were averaged over the two-year time period through four time intervals: 15-min, 60-min, daily and seasonally. Seasonal averages were once again based on the USGS water year (see Table 19). Similar to the spatial analysis, only valid measurements were averaged for the soil water potential sensor (i.e., MPS-1) and two soil moisture sensors (i.e., EC-5) and converted to their respective units. To compare the results, the coefficient of determination (R^2) and root mean squared error (RMSE) calculations are given for each pairing of original and duplicate-free calculations. The R^2 calculations are based on a linear regression between the original and duplicate-free data.

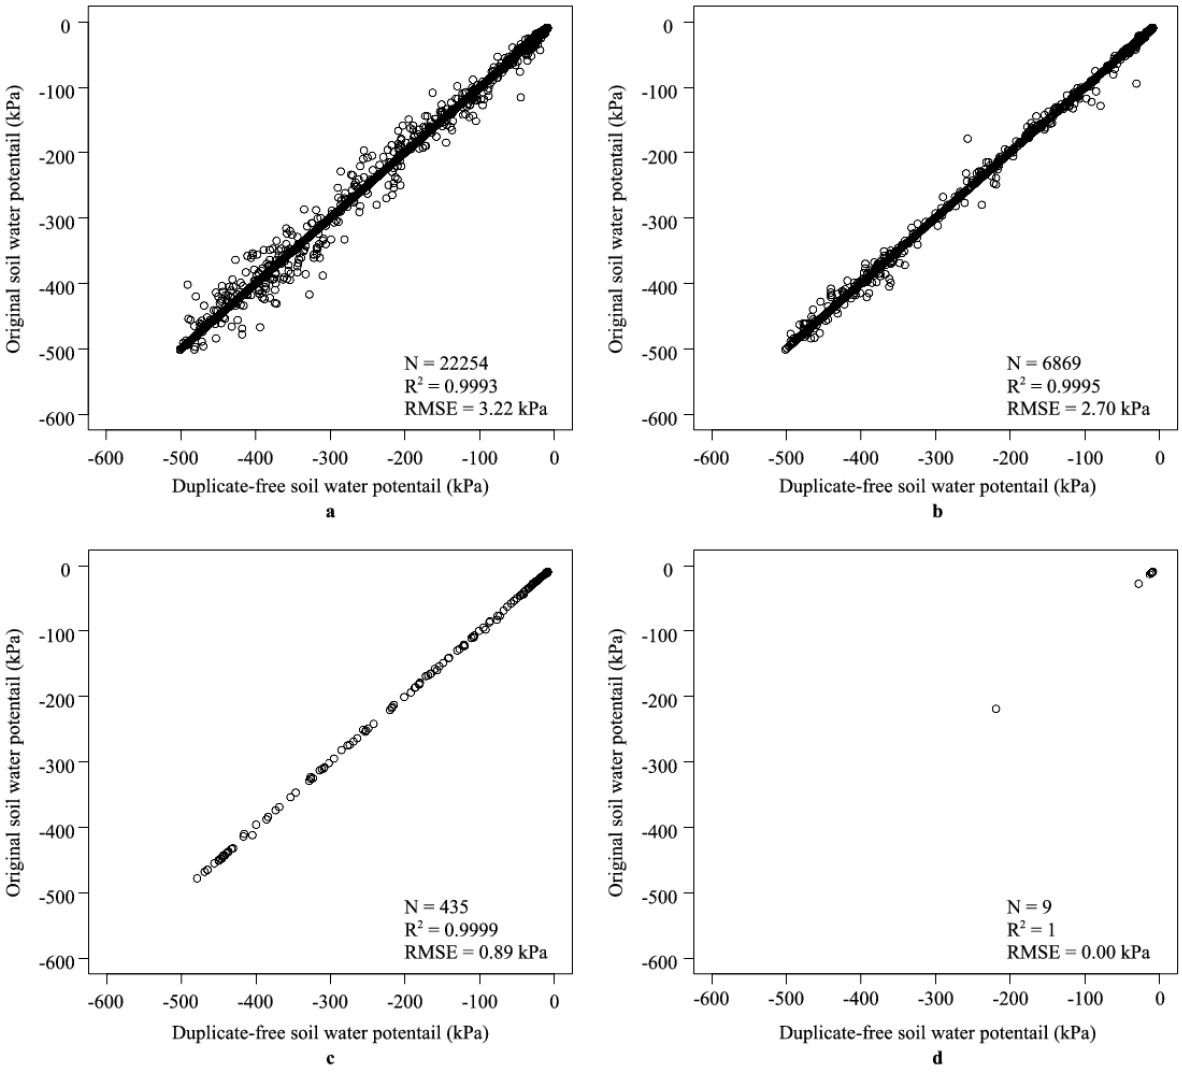

Figure 67. Original and duplicate-free comparison of soil water potential calculations at 30 cm for site 2 averaged at a) 15 min, b) 60 min, c) daily, and d) seasonally. The R^2 and RMSE calculations and the number of points used in the calculation are given in the bottom right corner of each plot.

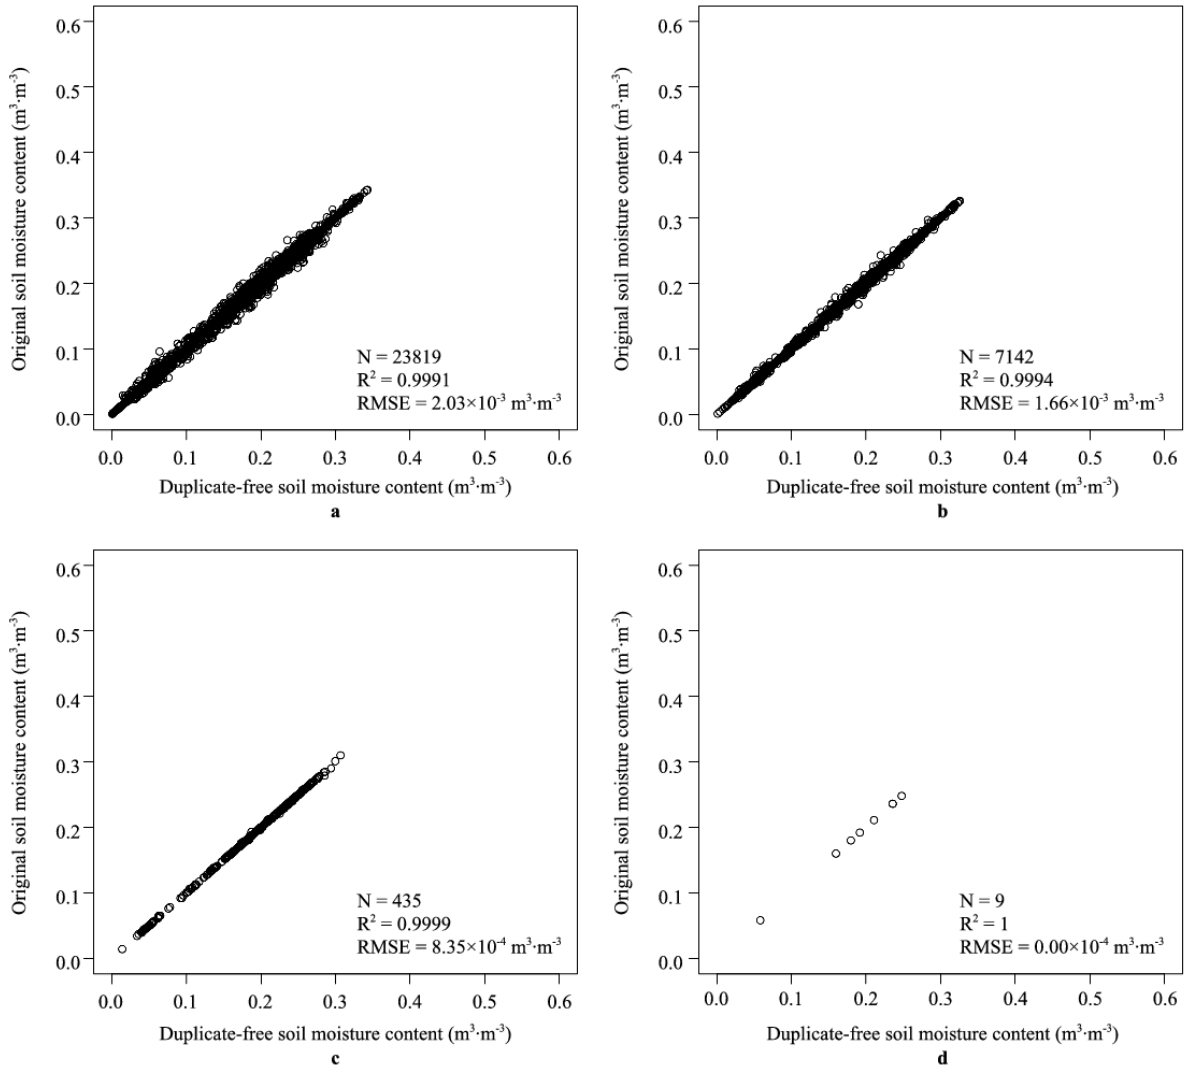

Figure 68. Original and duplicate-free comparison of soil moisture content calculations at 30 cm for site 2 averaged at a) 15 min, b) 60 min, c) daily, and d) seasonally. The R^2 and RMSE calculations and the number of points used in the calculation are given in the bottom right corner of each plot.

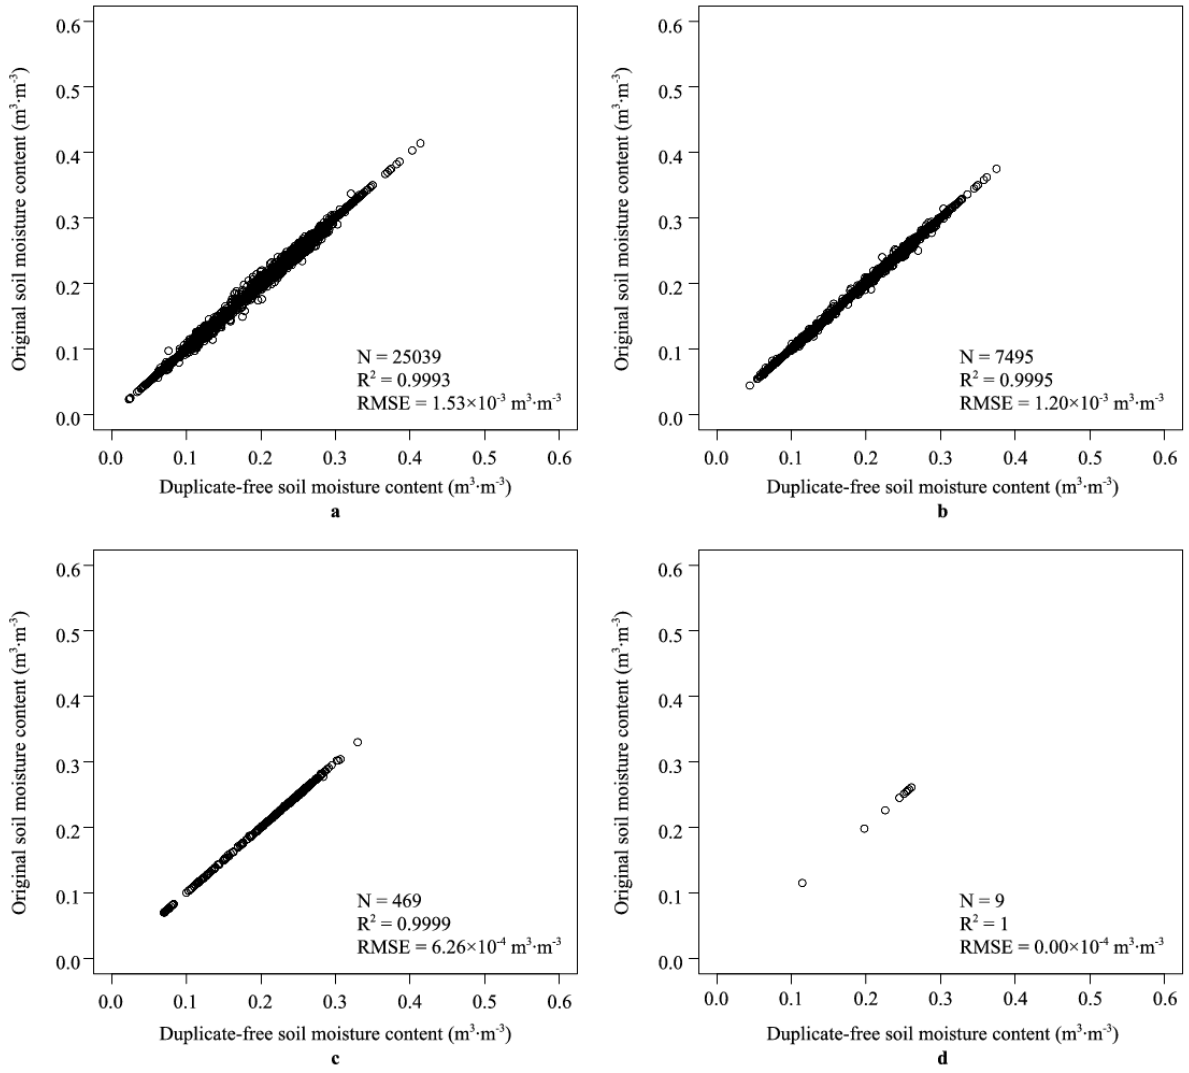

Figure 69. Original and duplicate-free comparison of soil moisture content calculations at 10 cm for site 2 averaged at a) 15 min, b) 60 min, c) daily, and d) seasonally. The R^2 and RMSE calculations and the number of points used in the calculation are given in the bottom right corner of each plot**.**

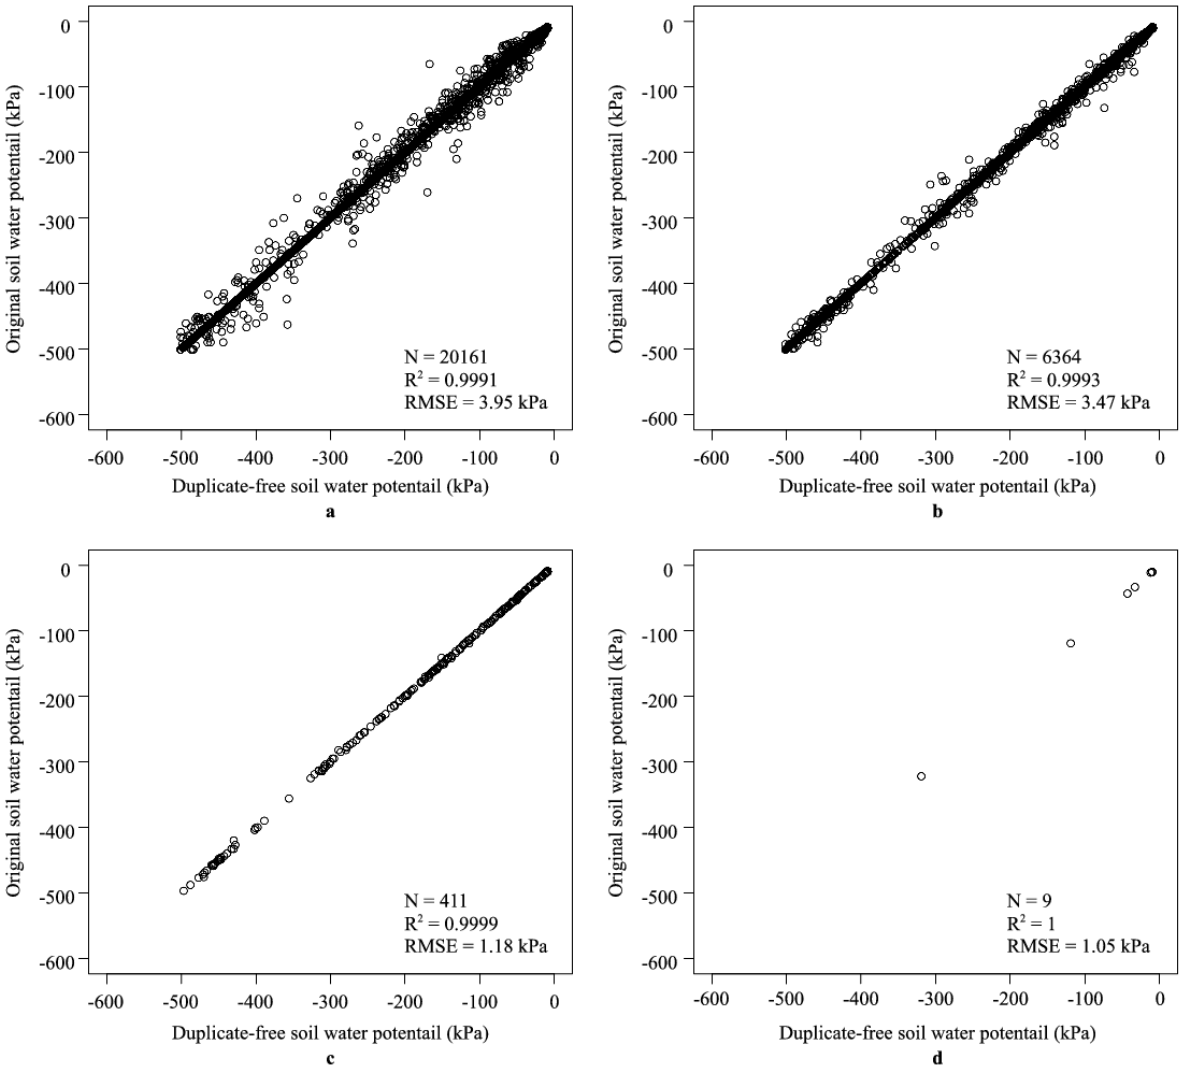

Figure 70. Original and duplicate-free comparison of soil water potential calculations at 30 cm for site 3 averaged at a) 15 min, b) 60 min, c) daily, and d) seasonally. The R^2 and RMSE calculations and the number of points used in the calculation are given in the bottom right corner of each plot.

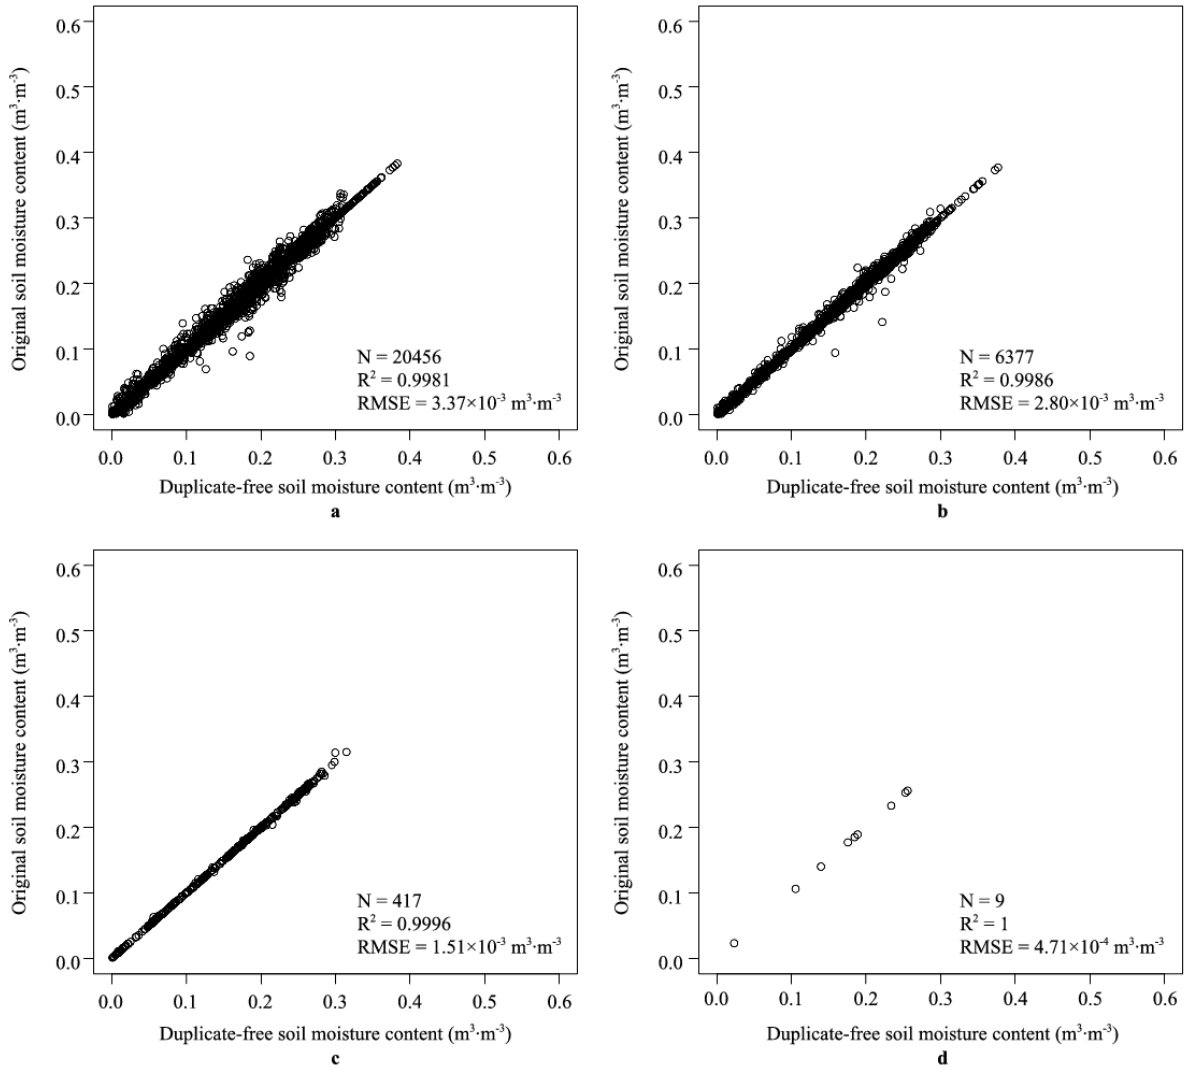

Figure 71. Original and duplicate-free comparison of soil moisture content calculations at 30 cm for site 3 averaged at a) 15 min, b) 60 min, c) daily, and d) seasonally. The R^2 and RMSE calculations and the number of points used in the calculation are given in the bottom right corner of each plot.

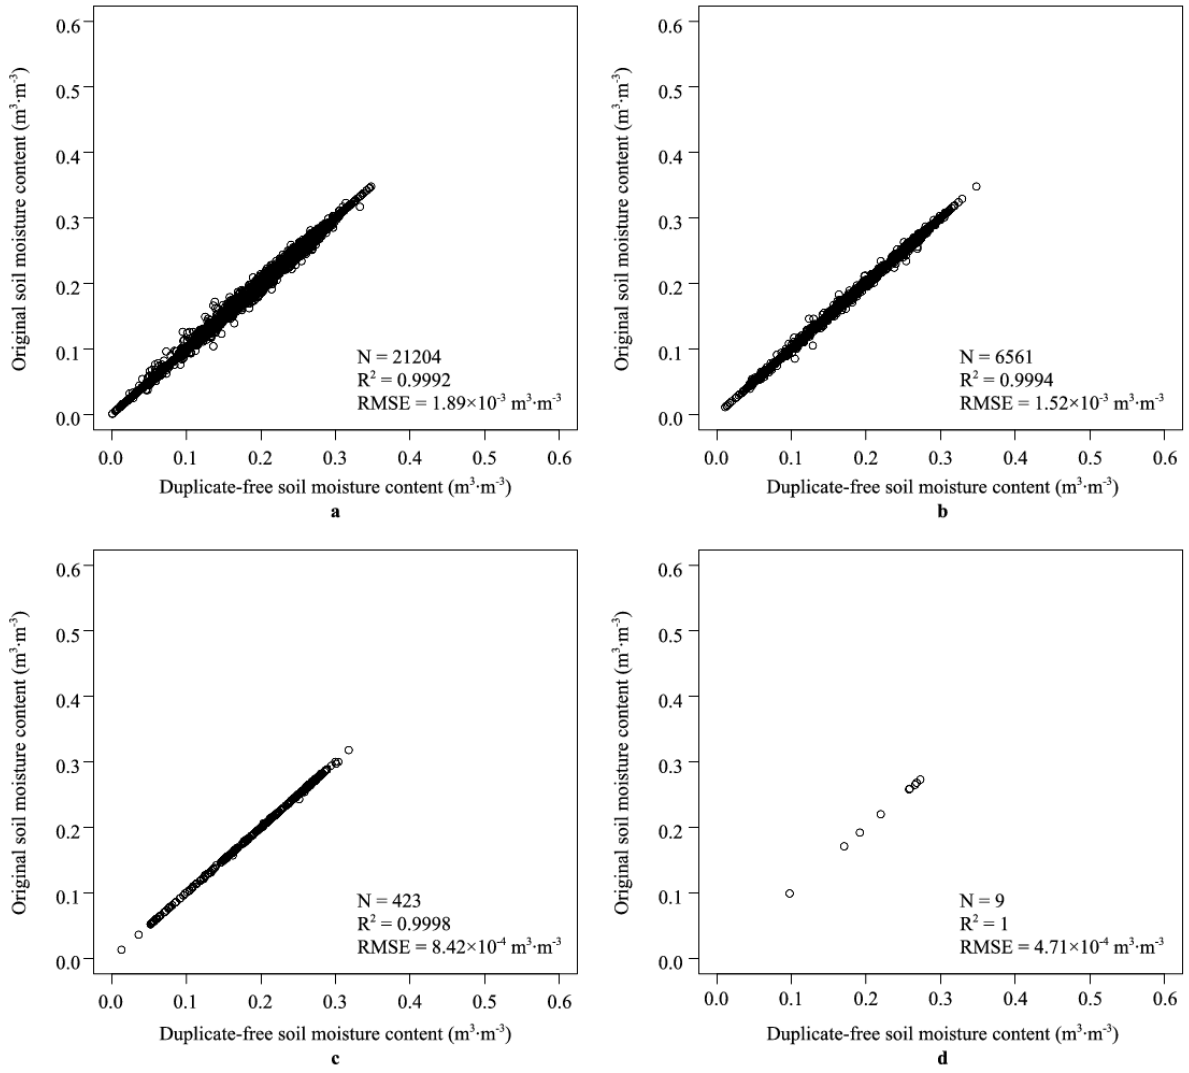

Figure 72. Original and duplicate-free comparison of soil moisture content calculations at 10 cm for site 3 averaged at a) 15 min, b) 60 min, c) daily, and d) seasonally. The R^2 and RMSE calculations and the number of points used in the calculation are given in the bottom right corner of each plot.

During the two-year period, between 3.4% and 4.1% of the sensor measurements made were duplicates. In each figure (Figure 67 to Figure 72), the impact of duplicate packets (in terms of R^2 and RMSE) on the sensor calculations decreases as the period of data aggregation increases (i.e., from minutes to days). For the shortest aggregation period (i.e., 15 min), where the highest error occurs due to duplicates, the amount of error introduced to the sensor calculations is limited as shown by the low RMSE (around 2% of measurements) and high R^2 values (greater than 0.998).

Because only valid measurements are compared, the scatter plots presented in Figure 67 to Figure 72 do not adequately show the impact of duplicate packets on removing valid measurements due to averaging. This typically only affects measurements made during extreme conditions (e.g., very wet or very dry), and therefore under moderate climate conditions (including the climate of the study site) this has a small impact on the overall measurement quality. During the period analyzed, less than 15 calculations in the 15-min and 60-min averaged data sets and no calculations in either the daily or seasonal averaged data sets were removed from the range of validity due to duplicates.

When using WSNs for environmental monitoring it is important to know that duplicate packets received in the data time series can affect sensor calculations. While the magnitude of errors introduced by duplicated data in this example is relatively small, the factors of how duplicates can affect measurement quality are still important and therefore need to be understood. The first factor is the period of aggregation. This factor is dependent upon the sampling interval at which measurements are made. In this example, the sampling interval is 15 min and there are 12 nodes in site 2 and 13 nodes in site 3 used in the data aggregation. The average number of samples used during each of the aggregation periods follows. For the 15-min

aggregation, between three and five measurements were used in each average. For the 60-min aggregation, between 12 and 17 measurements were used in each average. For the daily aggregation, between 185 and 265 measurements were used in each average. For the seasonal aggregation, between 8900 and 13800 measurements were used in each average. The increase in the number of samples used for averages as the aggregation period becomes longer results in a decrease in error caused by duplicates.

The second factor is the sensor type. It is important to understand how the raw voltage measurements are converted to sensor measurement units. This can affect how duplicated data impacts measurement quality, especially under certain environmental conditions. An example of this would be the MPS-1 soil water potential sensor. The conversion from the raw measurement voltage to water potential is given in Equation 6.2. The shape of the exponential (see Figure 41) is more heavily influenced by subtle changes in the measurement voltage under dry conditions (i.e., low measurement voltage or highly negative water potential).

6.4.1.3 Time series analysis of soil moisture and soil water potential

As a means of validation for the WSN soil sensor measurements, soil moisture and soil water potential were also measured in sites 2 and 3 on four EM50 data loggers (Decagon Devices, Inc.). Two EM50 data loggers were located in site 2 and an additional two EM50 data loggers were located in site 3. In site 2, one data logger was placed at the bottom and the top of the hillside. In site 3, the plot of land was relatively flat; since the two data loggers installed here could not be distinguished by location uphill and downhill, as in site 2, they are distinguished by direction (i.e., the southwest and northeast data loggers). Similar to the WSN, soil moisture was measured at 10 and 30 cm and soil water potential measurements were made at 30 cm. The

EM50 data loggers collected measurements at 15-min intervals to correspond with the WSN collection interval.

The time period of analysis begins on 18 April 2012 for site 2 and 4 April 2012 for site 3. These dates correspond to the installation of the MPS-1 sensor at 30 cm on the data loggers. In mid-July 2012, lightning struck a WSN mote which in turn destroyed the nearby EM50 data logger located at the top of the hill in site 2. Therefore, in the time series in site 2, the uphill data logger does not provide measurements following July 2012. The comparison period ends on 24 September 2012.

Duplicate packets were removed from the WSN data using the same algorithm that was used in the duplicate packet analysis (Section 6.4.1.2). Similar to the duplicate packet analysis, WSN data from sensor nodes in site 2 and 3 were aggregated into 15-min time series and processed using only valid data measurements.

Daily rainfall totals were provided by the Weather Underground personal weather station (PWS) KPAPITTS16. This weather station is operated by the Dorseyville Middle School, located approximately three miles away from the study region. Daily data is available online from the weather station for free public use.

The time series from the data loggers and the WSN are plotted together for comparison of the two time series. The following three figures (Figure 73 to Figure 75) show the comparison of the site 2 soil sensor measurements compared to the same variables measured on the data loggers.

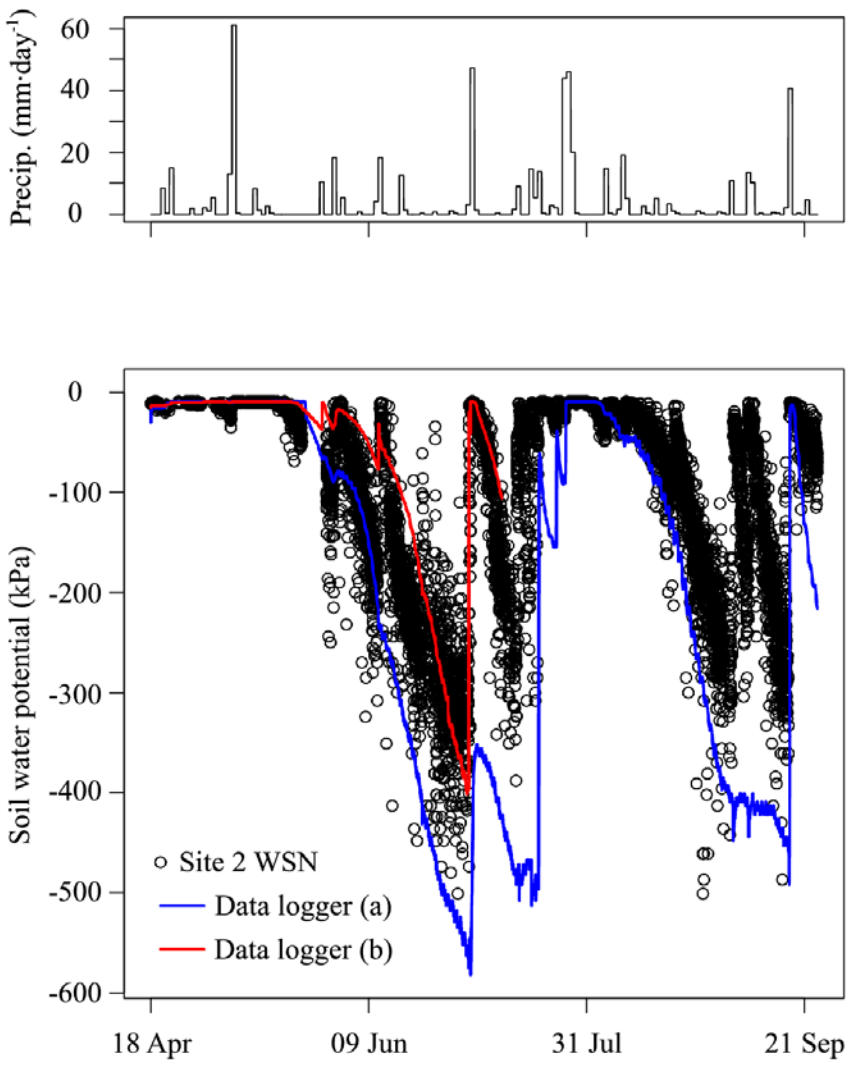

Figure 73. Daily rainfall and site 2 soil water potential measurements at 30 cm from the WSN (open circles) and data loggers located at the bottom of the hill (blue) and top of the hill (red).

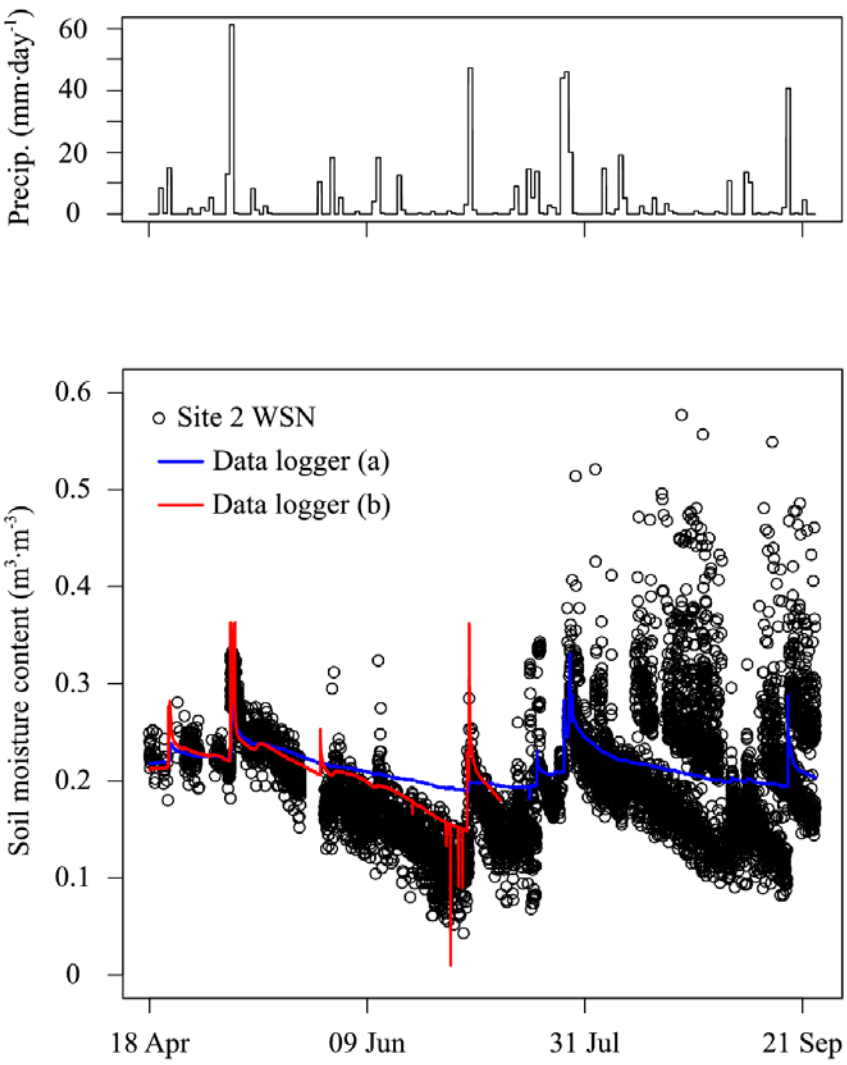

Figure 74. Daily rainfall and site 2 soil moisture content measurements at 30 cm from the WSN (open circles) and data loggers located at the bottom of the hill (blue) and top of the hill (red).

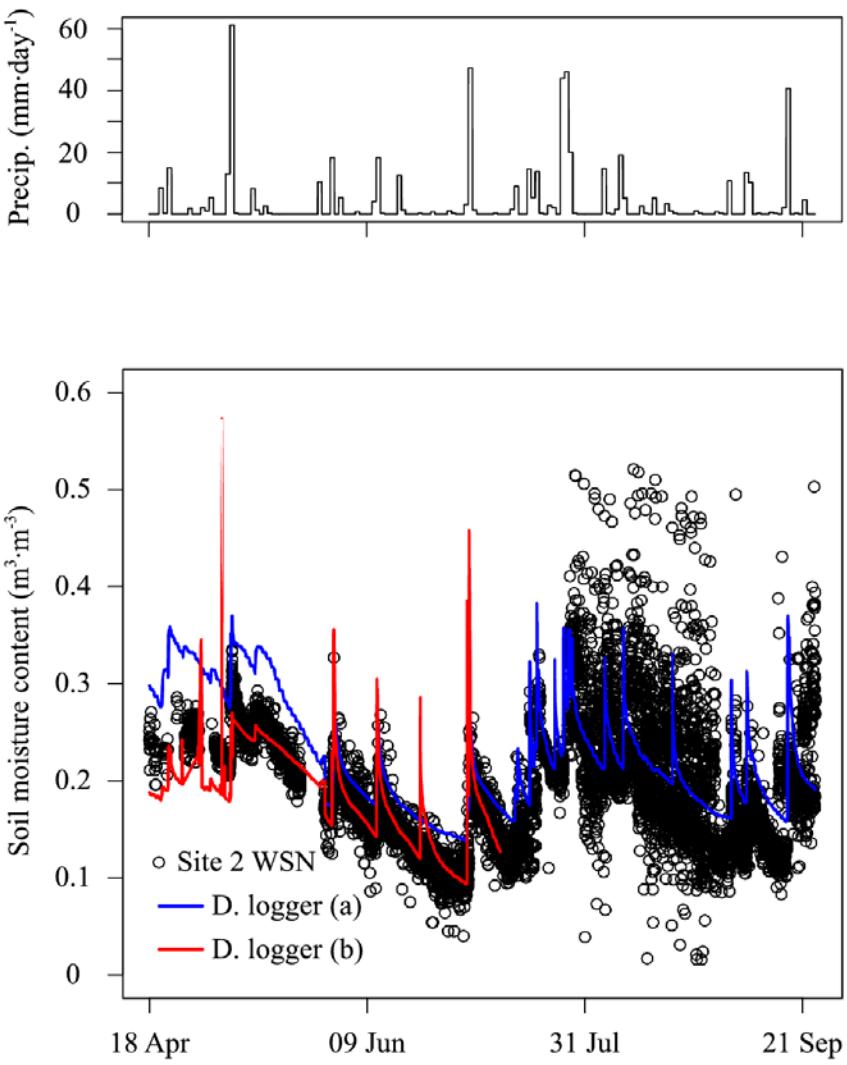

Figure 75. Daily rainfall and site 2 soil moisture content measurements at 10 cm from the WSN (open circles) and data loggers (solid lines) located at the bottom of the hill (blue) and top of the hill (red).

The results in site 2 show a close comparison between measurements made by the WSN and data loggers. Soil water potential measurements from the data logger at the bottom of the hill are similar to the drier WSN measurements, while the data logger at the top of the hill follows the mean water potential of the WSN measurements (Figure 73). Comparing the soil moisture content at 10 and 30 cm it can be seen that both correlate with both data loggers and precipitation time series. The soil moisture content at 10 cm is more responsive to rainfall in both the WSN measurements and data logger measurements (see Figure 75). As seen in Figure 75, following the lightning strike, which occurred between the 18 and 25 of July 2012, there is a considerable amount of noise introduced into the WSN data. It is unclear as to whether this increase in noise in the soil moisture measurements is due to residual static electricity in the soil or if it is damage incurred to the soil sensors as a result of the lightning strike. It is clear that this noise begins at the time of the lightning strike occurrence, as it is present in site 2 and not in site 3. The 30 cm soil moisture content in site 2 (see Figure 74) also shows some change in the water content measurements following the lightning strike; however, the effect is not as significant as that seen at 10 cm. The soil water potential measurements (which are at 30 cm) in site 2 do not appear to be effected.

The following three figures (Figure 76 to Figure 78) show the comparison of the site 3 soil sensor measurements compared to the same variables measured in the data loggers.

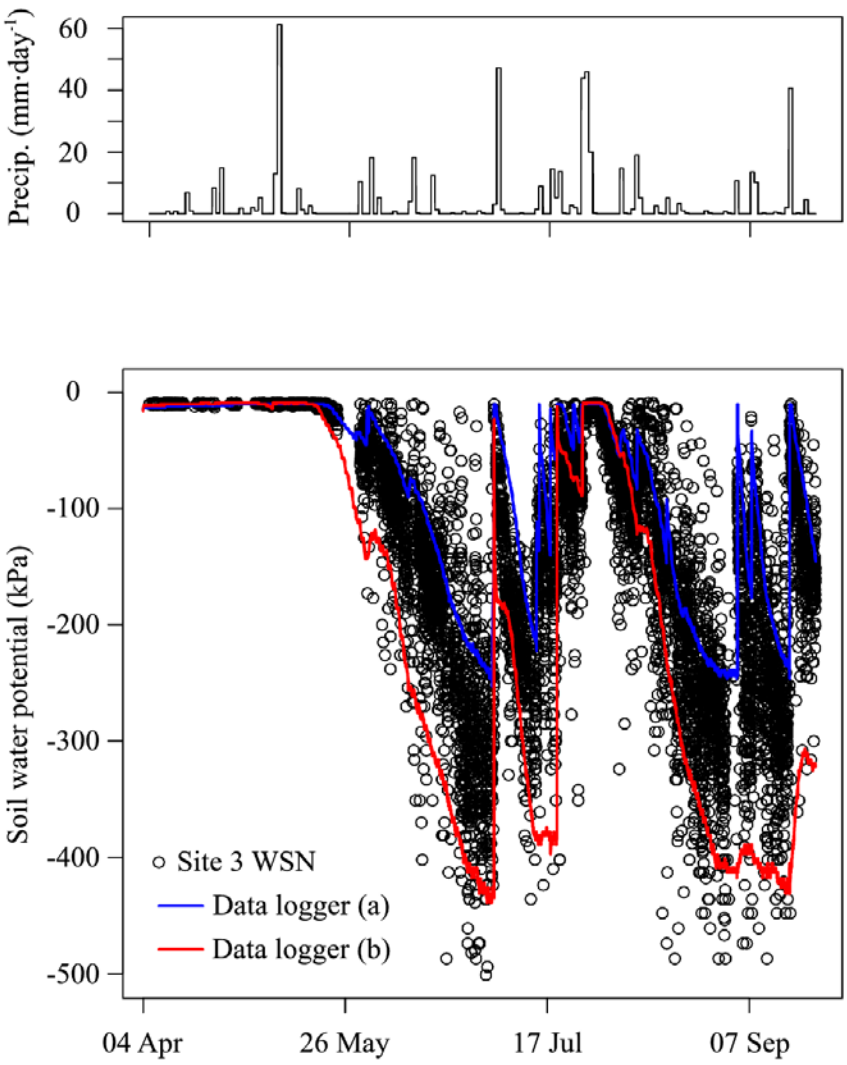

Figure 76. Daily rainfall and site 3 soil water potential measurements at 30 cm from the WSN (open circles) and data loggers (solid lines) located in the south west (blue) and north east (red).

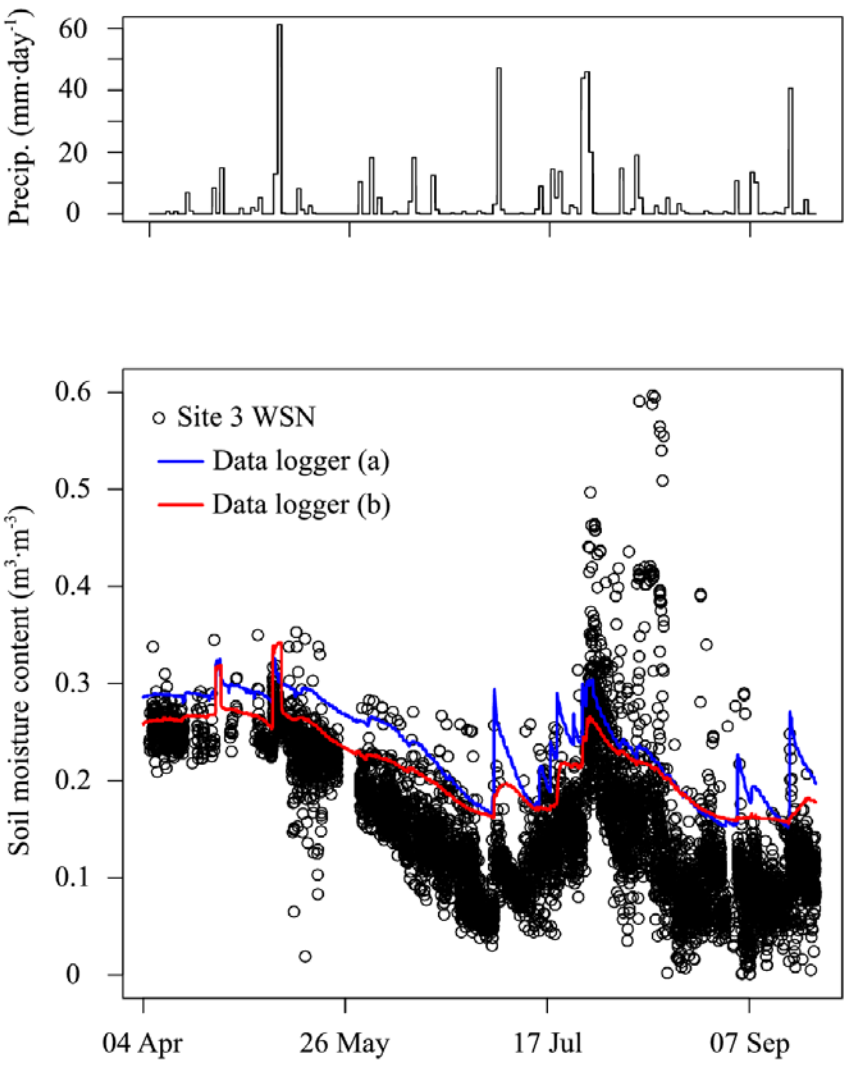

Figure 77. Daily rainfall and site 3 soil moisture content measurements at 30 cm from the WSN (open circles) and data loggers (solid lines) located in the south west (red) and north east (blue).

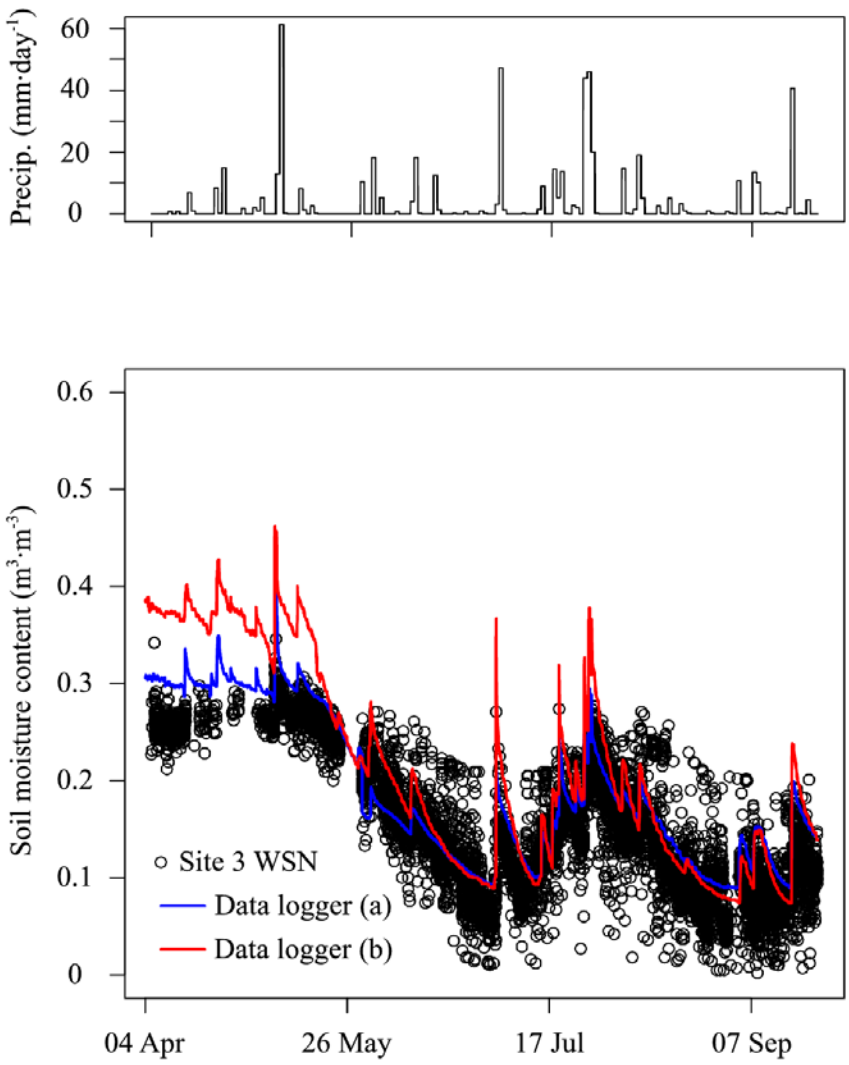

Figure 78. Daily rainfall and site 3 soil moisture content at 10 cm from the WSN (open circles) and data loggers (solid lines) located in the south west (blue) and north east (red).

The results of the site 3 comparison between the WSN and data loggers, similar to site 2, show an acceptable correlation between the measurements. The 30 cm soil water potential measurements of the southwest and northeast data loggers present high and low bounds of the WSN measurements, respectively (see Figure 76). The 30 cm soil moisture content measurements are consistently drier from the WSN compared to both data loggers (see Figure 77). Despite the drier measurements, the WSN response to rainfall events is closely represented by both measurement methods. The 10 cm soil moisture content is initially drier when measured by the WSN; however, it becomes consistent with data logger data after roughly the first month of measurements (see Figure 78). The WSN measurements at 10 cm do not appear to be as sensitive to the rainfall events as the two data loggers; however, the overall trend is well represented.

6.4.2 Sap flow analysis

The other environmental variable being monitored at the ASWP test bed is sap flow. Eleven nodes (five nodes in site 2 and six nodes in site 3) were outfitted with TDM sap flow sensors in 2012. Sap flow monitoring was also conducted during the previous two years. In 2010, the 12 V battery status was not monitored in the WSN data, which made distinguishing valid and invalid sap flow measurements difficult. In 2011, persistent network outages resulted in limited data available for analysis. In 2012, both of the problems faced in the past deployments were resolved and therefore 2012 data is the focus of this analysis.

The sap flow analysis individually examines the measurement results collected from sites 2 and 3. The sap flow measurements from both sites are seasonally separated such that daily trends and seasonal changes can be observed. Unlike data logger sap flow measurements, the

WSN required specially made circuitry for the sap flow sensors to be integrated with the data acquisition board (i.e., MDA300). This analysis includes a brief description of the WSN sap flow circuit and some of the implications of its use.

6.4.2.1 Seasonal sap flow measurements

There is a predominant difference in the tree size and type distinguishing sites 2 and 3. With the assistance of Dr. Tonsor (Biological Science Department, University of Pittsburgh) the tree species within the two sites were identified. Site 2 consists mainly of maple and ash trees ranging in size from saplings to over 40 cm diameters. All of the wireless motes monitoring sap flow in site 2 are attached to silver maple trees (*Acer saccharinum*). Ash trees were specifically avoided due to the emerald ash borer, a beetle responsible for the destruction of many ash trees in the area. Trees in site 3 are mostly old growth oaks, some in excess of 70 cm in diameter. Sap flow sensors in site 3 are attached to red oak (*Quercus rubra*) and white oak (*Quercus alba*) trees. Sapwood area for each of the monitored trees in sites 2 and 3 was estimated based on three coring samples taken from each tree.

The 11 nodes that were outfitted with sap flow sensors in 2012, their respective installation date, the tree type, the direction the sap flow sensor is oriented on the tree, the tree diameter at breast height (DBH), the elevation of the sensor above the ground, and the percent sapwood area (SWA) are given in Table 21. The locations of these nodes with respect to their individual sites are given in Figure 31.

Table 21. The 2012 monitoring start date, tree type, sensor orientation, breast-height tree diameter, node elevation above the ground, and percent sapwood area of the 11 WSN nodes and their respective trees.

Node ID	Start date	Tree type	Orien.	DBH (cm)	Elev. (cm)	SWA (%)
2045	25 May	Silver maple	S	34	133	77
2055	24 May	Silver maple	S	31	136	70
2065	07 Aug	Silver maple	W	32	150	77
2095	07 Aug	Silver maple	S	41	141	82
2115	07 Aug	Silver maple	W	24	136	83
3005	13 Jun	Red oak	SE	62	124	13
3045	13 Jun	Red oak	N	44	130	17
3055	13 Jun	Red oak	S	50	159	15
3065	25 Aug	Red oak	N	39	145	12
3075	08 Aug	Red oak	S	53	144	17
3125	21 Jun	White oak	E	57	135	14

The time series data for each of the 11 nodes (from the start date indicated in Table 21 to 8 October 2012) were processed to remove the duplicate packets using the algorithm defined in Navarro et al. (2013). The irregular time series data from each mote was aggregated to regular 15-min intervals. For each node, between 45 and 138 individual days of measurements were examined. Each time series was filtered to remove days from the analysis that did not meet specific criteria. The criteria for valid sap flow are: 1) measurements that were taken during low mote battery power (i.e., less than 2.6 V) or low 12 V battery power (i.e., less than 6 V) were discarded; 2) measurements were counted within three 6-hr time windows: a) 03:00–09:00, b) 09:00–15:00, and c) 15:00–21:00 and days that did not have at least three measurements taken within each time window were discarded. The purpose of these three time windows was to ensure that enough data was present throughout an entire day to formulate the sap flow curve.

The valid days of measurements for each of the 11 nodes were separated into three seasons: spring (May and June), summer (July and August), and autumn (September and

October). The total number of valid days for each node and the respective number of days within each season are given in Table 22. Based on the filtering results and the time periods analyzed, there are fewer days presented in the spring compared to the summer and autumn (approximately 22–32%). There is almost a two-month difference in the number of valid measurements received in site 2 compared to site 3. This is most likely due to reduced communication success of motes in site 3 compared to site 2.

For each season, the daily sap flow calculations of each mote were averaged. The mote daily averages are given in Figure 79 and Figure 80 for sites 2 and 3, respectively. The daily average of all the motes within each season is also given. During the time sap flow measurements were made, understory photos were taken in sites 2 and 3 to depict changes in the tree canopy. The photos representing the respective seasons are presented beside the seasonal sap flow calculations.

Table 22. The total number of valid days of sap flow measurements for the five nodes in site 2 and six nodes in site 3 and the respective number of sap flow measurement days within the three analyzed seasons in 2012.

Node ID	Spring (May–Jun)	Summer (Jul–Aug)	Autumn (Sep–Oct)	Total
2045	17	29	23	69
2055	14	30	20	64
2065	0	9	21	30
2095	0	16	18	34
2115	0	13	22	35
Site 2 totals	31	97	104	232
3005	7	14	5	26
3045	6	24	9	39
3055	4	19	13	36
3065	0	3	16	19
3075	0	5	19	24
3125	1	17	17	35
Site 3 totals	18	82	79	179

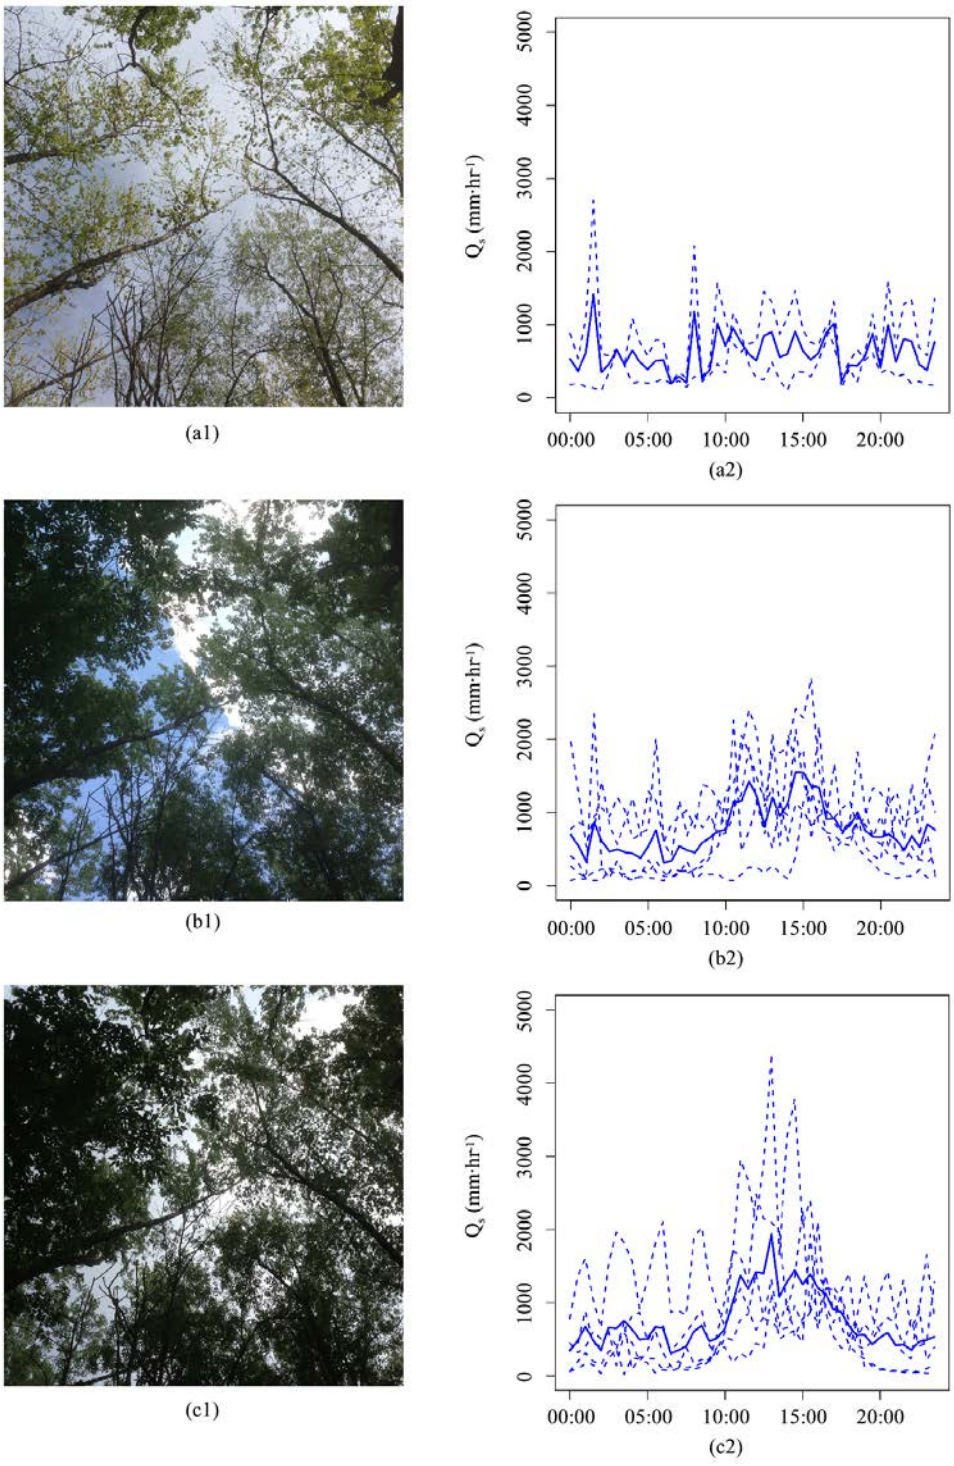

Figure 79. Site 2 understory photos during spring (a1), summer (b1), and autumn (c1) 2012. Site 2 daily average (dashed) and seasonal daily average (solid) WSN sap flow measurements during the spring (a2), summer (b2), and autumn (c2) 2012.

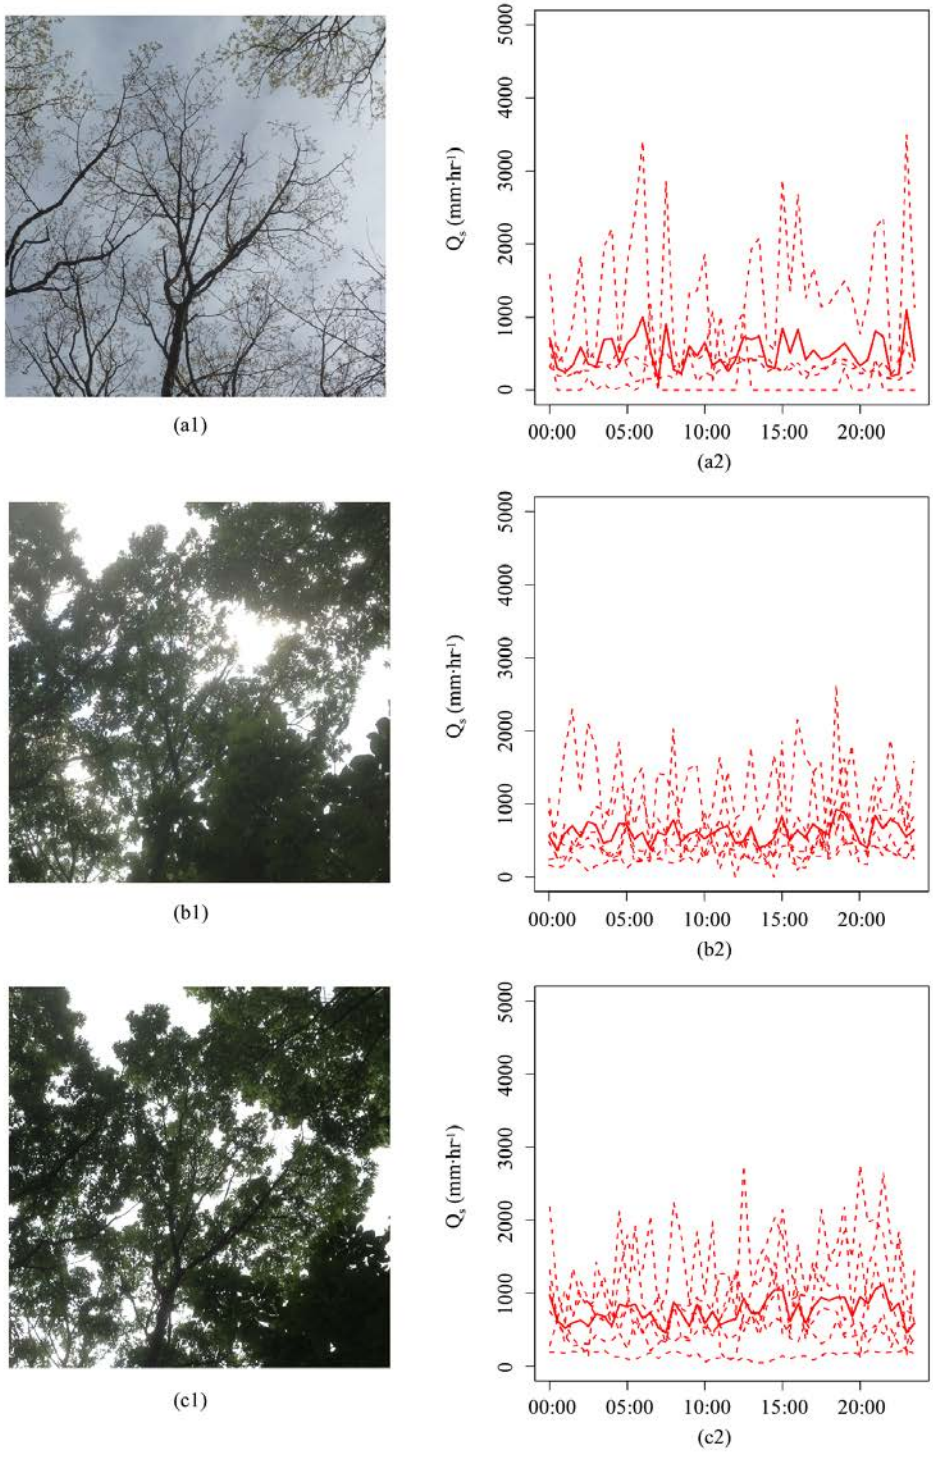

Figure 80. Site 3 understory photos during spring (a1), summer (b1), and autumn (c1) 2012. Site 3 daily average (dashed) and seasonal daily average (solid) WSN sap flow measurements during spring (a2), summer (b2), and autumn (c2) 2012.

Daily trends in sap flow are shown in the calculation results for site 2. This is especially the case during the summer and autumn seasons (Figure 79 b2 and c2). In addition, the magnitude of the average sap flow measurements increases over the seasons. This is shown by the increase in the peak of the average seasonal sap flow from around 1000 $mm \cdot hr^{-1}$ during the spring to approximately 1500 $mm \cdot hr^{-1}$ in the summer and 2000 $mm \cdot hr^{-1}$ in the autumn. The increases in sap flow magnitude correspond with the increases in the canopy foliage (Figure 79 a1, b1, and c1).

There are no discernible daily trends in the sap flow calculation results for site 3 during any of the three months (Figure 80 a2, b2, and c2). There are a few possible causes for the absence of a trend. First, the tree sizes in site 3 are generally larger than the trees in site 2 (see Table 21). Larger trees may be more susceptible to circumferential variations of sap flow. Due to the single location of sap flow measurements on each tree, it is possible that poor locations were selected for sap flow monitoring. As noted in Section 2.4, circumferential sap flow variations have been reported as high as 50% depending on the location of the sap flow sensor. It should be noted that the orientation of the sap flow sensors in site 3 represent three of the four major directions (Table 21); therefore the orientation may not be a major cause for the static sap flow measurements seen throughout all the sap flow results in site 3.

Another cause may be the relative sapwood area of the trees in site 3. Table 21 shows estimates of the percentage of sapwood area (with respect to the total cross-sectional area of the tree). These estimates show that the sapwood area of the trees in site 3 is considerably less than the sapwood area of the trees in site 2. The sap flow sensor probe length is designed to be 3.81 cm in length (Section 2.2.1). Due to the small representative sapwood area of the trees in site 3, it is possible that the sensor probe length is positioned within both the tree's sapwood and

heartwood. This would cause only a fraction of the sensor probe length to be measuring sap flow while the remaining sensor probe length (within the heartwood) would be measuring no sap flow. The relative percentage of the sap flow probe length within the heartwood, as opposed to sapwood, would have to be large for the results to show little to no change over daily/seasonal scales. This is partially explained by the sapwood estimates presented in Table 21. Future sap flow monitoring in this site may need to consider a different probe design. James et al. (2002) propose a thermal dissipation sensor with a 1 cm effective measurement element. This may reduce errors associated with measuring across sapwood and heartwood with longer probe designs.

6.4.2.2 Wireless sap flow circuit

The sap flow sensors used in this study are based on single-end thermocouple temperature sensors. The voltage responses to the Type E thermocouple, used in this sensor design, range from 0–3 mV over temperature changes from 0–50 °C (Figure 1b). The MDA300 data acquisition board, discussed in Section 5.1.2, has a 0.6 mV sensitivity to input voltages. The low voltage response of the thermocouple probes and the sensitivity of the MDA300 data acquisition board allows for as fine as 10 degree temperature changes to be monitored. This is inadequate for sap flow when maximum temperature differences experienced throughout the day are on the order of 10 degrees. Therefore, in order to improve the resolution of temperatures that can be monitored by the MDA300 data acquisition board, a signal amplifier is required.

TDM sap flow monitoring requires a secondary power source for the constant heater probe. The wireless mote's AA battery does not provide the voltage high enough or capacity large enough to support the power requirements of the constant heater probe. To maintain the validity between the sap flow sensors deployed in the field with Granier's empirical sap flow

equation (Equation 2.1), the output power of the constant heater probe needs to match the same output power used to develop the TDM equation (i.e., 0.2 W, as discussed in Section 2.2.1.2). A voltage regulator is therefore required to adjust the output power of the heater probes in the field. The ability to adjust the output voltage in the field is critical in that the heater probe resistance can alter during field installation (e.g., heater coil shortages).

Based on these two design needs, a circuit board was developed to improve wireless sap flow monitoring deployments. Early prototypes of this circuit design attempted to address each problem separately. The current circuit board design integrates both the signal amplifier and voltage regulator together on the same circuit board (Kuo et al., 2010).

A single 12 V battery input (Figure 81a) is partitioned into two variable voltage outputs. The first voltage output (Figure 81b) provides between 2.5–3.75 V which corresponds to the range of input voltages for powering the MICAz mote (as well as the Iris mote). The second voltage output (Figure 81c) ranges from less than 1 V to approximately the current voltage setting of the first voltage output. The design of the two output voltages limits the maximum of the second output to the current setting of the first. In this way, only one voltage regulator is required for the circuit design. The second voltage output allows for adjusting the power of the constant heat probe. The 12 V battery is responsible for powering both the voltage regulator (Figure 81j) and the two signal amplifiers (Figure 81i).

Due to the necessity of measuring two temperature probes simultaneously, the circuit board design requires two separate signal amplifiers. Individual thermocouple inputs are provided: one for the constant heater probe (Figure 81e) and one for the reference temperature probe (Figure 81g). The two thermocouple inputs are amplified by their own respective signal

amplifier (Figure 81i). The single-end outputs of the amplified thermocouple voltage responses are accessible at individual output ports (Figure 81f and h).

The circuit board design is relatively low power (less than 1 W) and has a current draw estimated between 30 and 75 mA depending on the load (i.e., powering the heater, the mote, or both). Lab and field tests using the circuit design (Figure 81) presented measurements with unexpected levels of noise. An investigation on the causes for this noise was conducted.

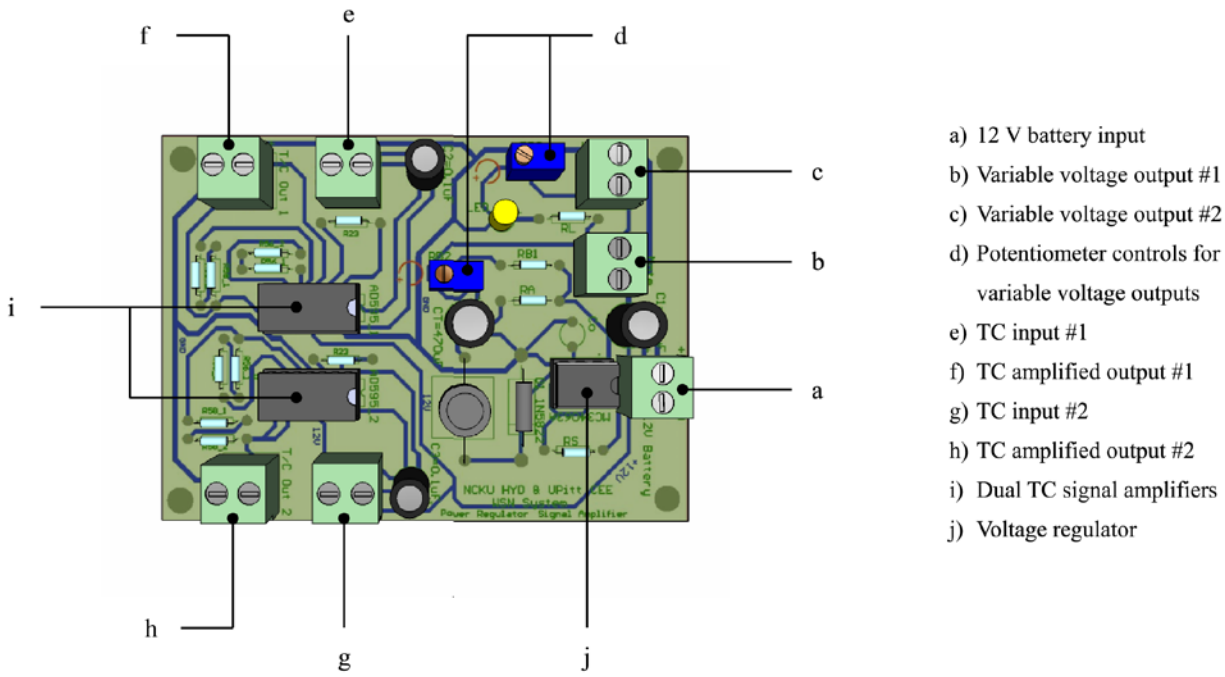

Figure 81. WSN sap flow circuit design highlighting the major components (a–j).

6.4.2.3 Noise analysis of wireless sap flow measurements

To begin this analysis, some levels of noise were also found in data logger measurements at times of peak sap flow occurrence. These measurements were analyzed to see if the noise levels in the sap flow measurements were significant. Two periods of sap flow measurements were chosen: 16–23 September 2007 (Figure 82) and 22–25 August 2008 (Figure 8). Both of these sets of sap flow measurements were made on CR1000 data loggers (Campbell Scientific) using TDM sap flow sensors. The 2007 measurements (three sensors) were made at 1-min intervals and the 2008 measurements (one sensor) were made at 1-s intervals.

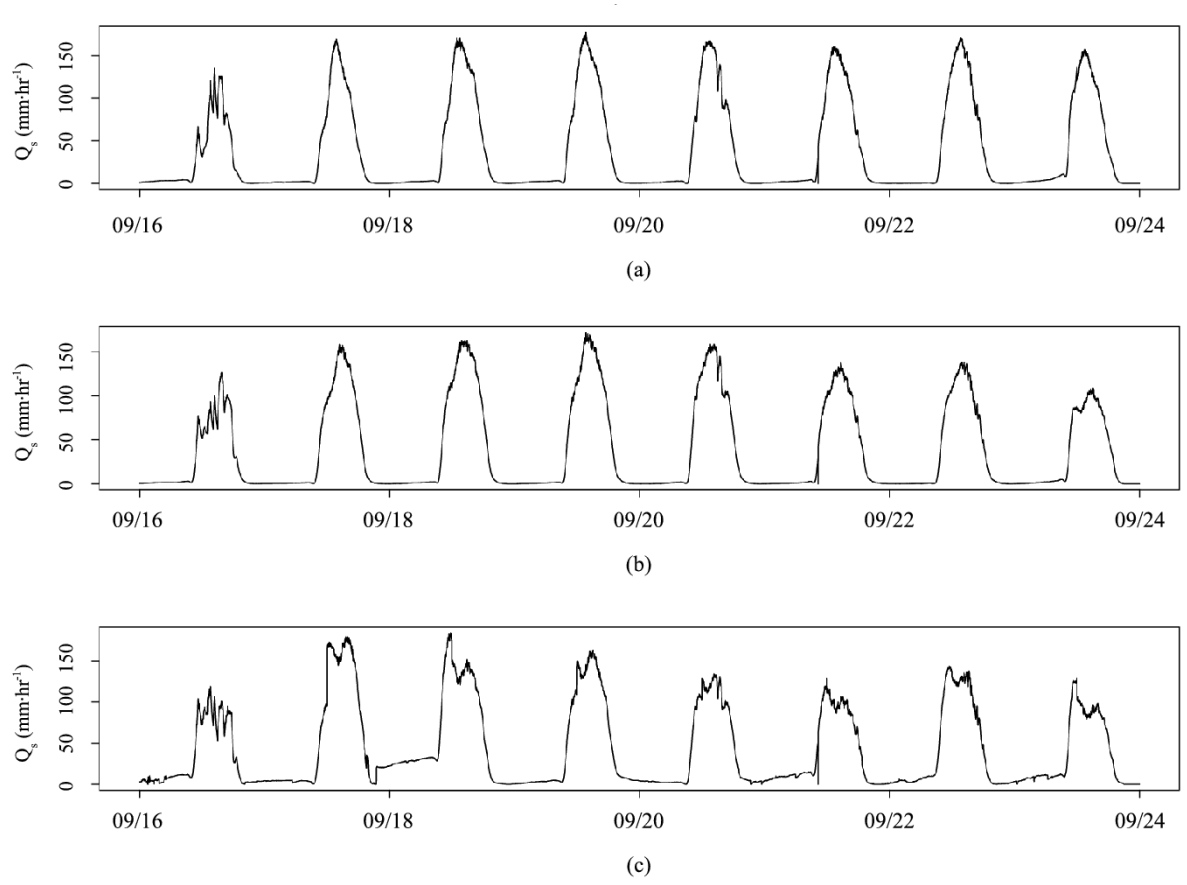

Figure 82. Sap flow measurements from three thermal dissipation style sap flow sensors from 16–24 September 2007.

A power spectrum analysis was performed on the three sap flow sensor measurements from 2007 and one sap flow sensor's measurements from 2008. To calculate the scaled periodogram of the sap flow measurements, the following equation was used:

Equation 6.3

$$P(Q_s) = \left| \frac{2 \cdot fft(Q_s)}{N(Q_s)} \right|^2$$

where: Q_s = sap flow time series

P = scaled periodogram

N = number of data points

fft = fast Fourier Transform

The scaled periodograms for the three sap flow sensors shown in Figure 82 are presented in Figure 83.

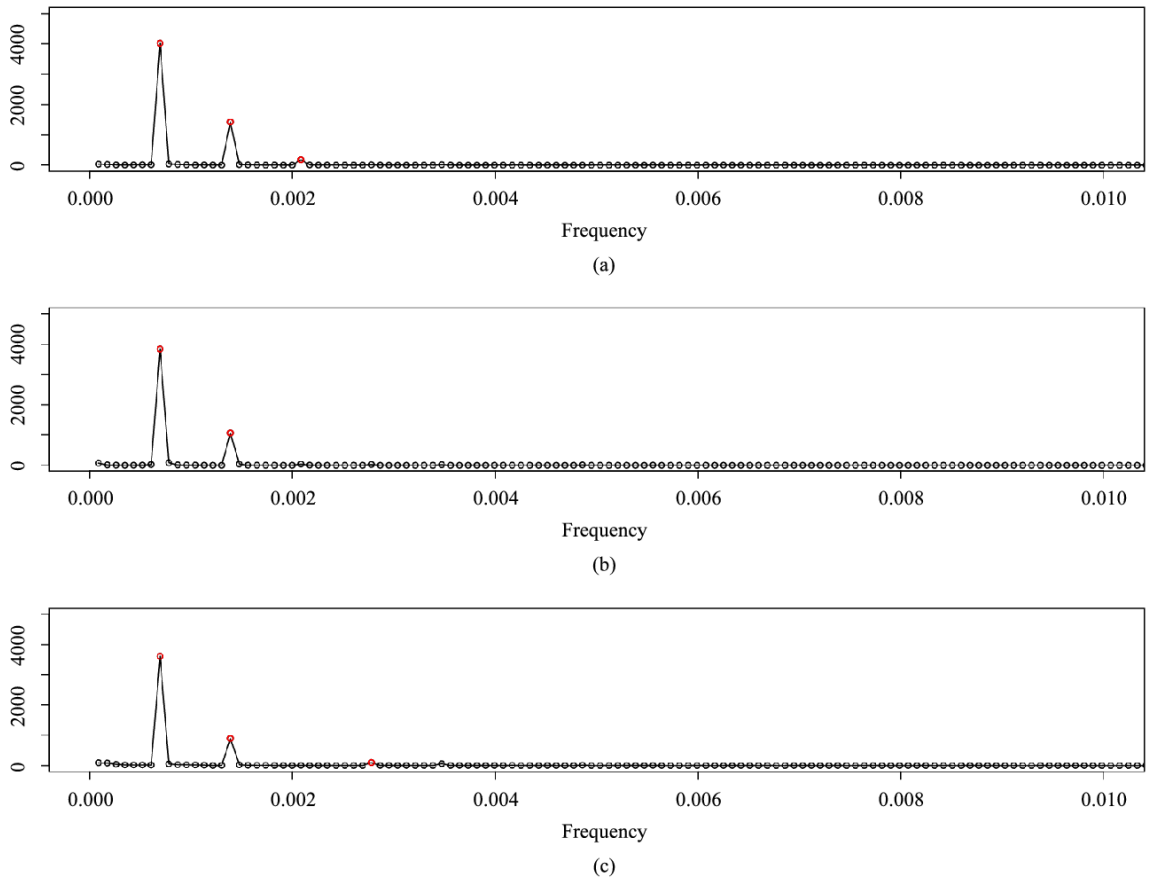

Figure 83. Periodograms of sap flow sensor measurements from 16–24 September 2007.

The dominate frequencies, for each of the three sap flow sensor measurements, are highlighted in red. The dominate periods (i.e., frequency $^{-1}$) for the results in Figure 83a are 24 hr (0.00069 min $^{-1}$), 12 hr (0.00139 min $^{-1}$), and 8 hr (0.00208 min $^{-1}$). The dominate periods for the results in Figure 83b are 24 hr and 12 hr. The dominate periods for the results in Figure 83c are 24 hr, 12 hr, and 6 hr (0.00278 min $^{-1}$).

The scaled periodogram for the 2008 measurements (not shown) provides similar results to the 2007 data, with dominate periods of 24 hr and 12 hr. The results from both the 2007 and 2008 power spectrum analysis show that high frequency noise in peak sap flow measurements is not a dominate frequency and therefore can be ignored. After determining that high frequency variability in sap flow measurements is noise, a series of tests were completed to understand the sources of noise present in the wireless measurements. Potential sources for the noise included the data acquisition board, the lab-made sap flow probes, and/or the wireless sap flow circuit.

The data acquisition board (i.e., the MDA300) was decidedly not the source of the noise. If it had been the source, measurements made by other environmental sensors would have also exhibited similar noise patterns in their data (e.g., soil moisture and soil water potential sensors). However, this was not the case and therefore the remaining two potential sources were investigated.

The first test conducted measured ambient temperature from lab-made sap flow probes, which were connected to a mote through the sap flow circuit board (Section 6.4.2.2) over a period of 24 hr. For comparison, professional grade thermocouples (Omega Engineering, Inc.) were also tested, in place of the sap flow probes, to determine if the lab-made sap flow probe construction was the cause for the noise in the measurements. The results of these two tests are presented in Figure 84.

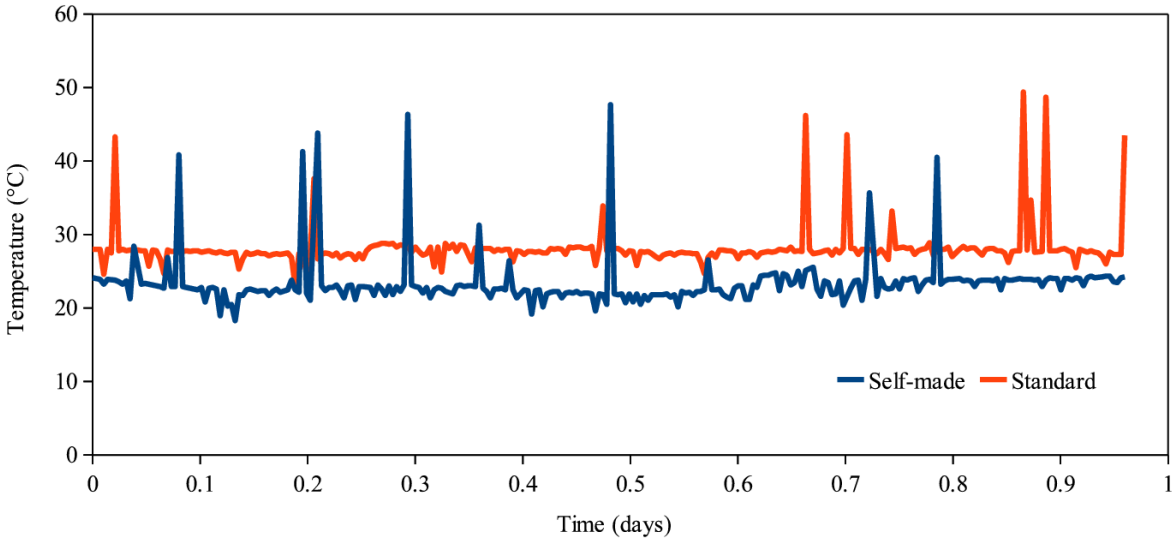

Figure 84. Variability in lab-made and professional grade (standard) thermocouples measuring ambient temperature by a wireless mote through the sap flow circuit board.

The difference in the magnitude of the two temperature series can be attributed to the fact that they were tested on different days. While the lab-made thermocouple has slightly higher variability (± 3.48 °C) than the professional grade (standard) thermocouple (± 2.95 °C), both measurements have considerable fluctuations resulting in temperature peaks greater than 20 °C above the mean. These results suggest that the quality of the thermocouples in the lab-made sap flow sensors is not the cause for noise in sap flow measurements.

The other possible source for the noise was the sap flow circuit board. An examination of the signal amplifier (i.e., AD595a) data sheet revealed that the amplifier is calibrated with an accuracy of 3 °C, which happened to be the variability shown in both sensors within Figure 84. It was theorized that this amplifier was the potential source for noise. An alternative amplifier

(i.e., AD595c) was available and was calibrated with an accuracy of 1 °C. To determine whether noise could be reduced by using the AD595c, as opposed to the AD595a, tests were conducted. Five days worth of ambient temperature measurements were collected from lab-made sap flow sensor probes using both the AD595a and AD595c. Temperature anomalies from both tests are presented in Figure 85.

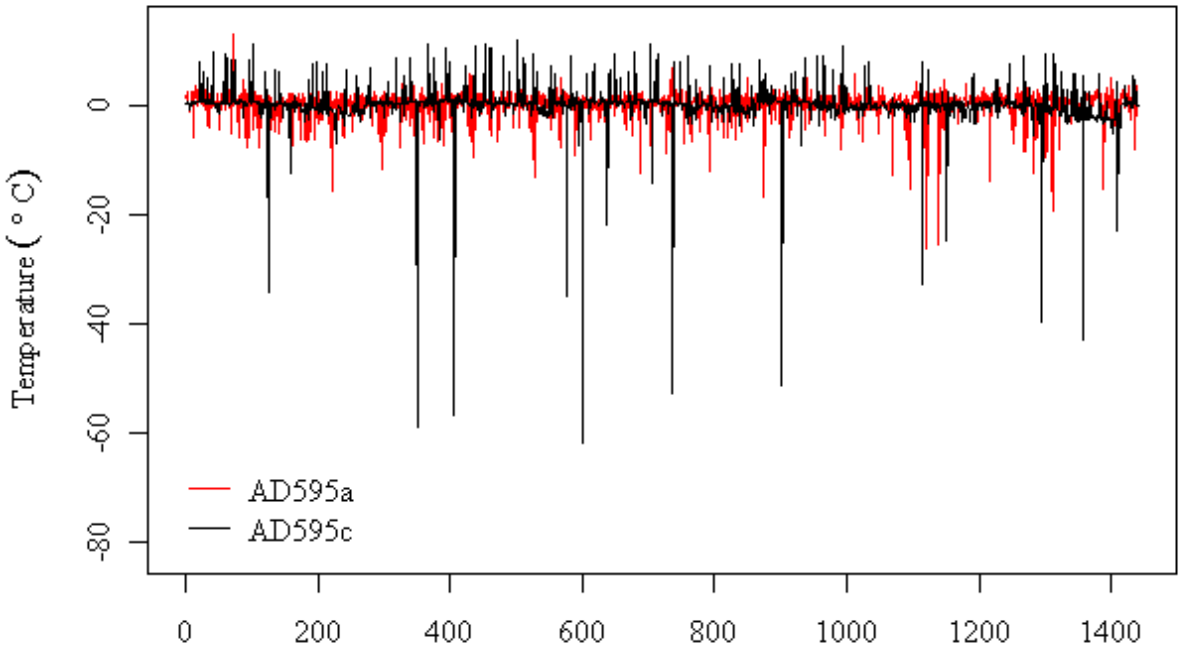

Figure 85. Five days of ambient temperature anomalies made at 5-min intervals from lab-made thermocouples by wireless motes through the AD595a amplifier (red) and the AD595c amplifier (black).

The results from these tests show that measurements made using the AD595c have reduced amounts of high frequency noise compared to the measurements made using the AD595a. However, irregular high amplitude noise still occurs in the measurements made by the

AD595c. These high amplitude fluctuations are larger than those measured using the AD595a (standard deviation of temperature measurements are 29.6 °C and 48.0 °C for the AD595a and AD595c, respectively). These results suggest that the AD595a, while exhibiting more noise, is more appropriate for wireless sap flow monitoring due to its relatively lower noise amplitude.

6.5 DISCUSSION

6.5.1 Compensating for sap flow measurement noise

Due to the presence of noise in sap flow sensor measurements, introduced by the sap flow circuit board, data preprocessing is necessary. There are a number of different procedures for noise removal and reduction. Outliers (i.e., noise) can be eliminated using standard statistical analysis based on probability density or distribution estimation (Staszewski, 2002). There are several options for data filtering and/or smoothing, such as the Wiener filter and the Savitzky-Golay and Nadaraya-Watson methods for smoothing. In addition, there is the orthogonal wavelet analysis, which can be used for denoising signals in datasets.

While these preprocessing methodologies aim to improve noisy datasets, there are shortcomings to each of these methods. For example, an assumption made using the Wiener filter is that the signal and the noise are stationary linear stochastic processes. This may not be a reasonable assumption given that temperature fluctuations measured within the tree are influenced by atmospheric temperature changes, which is non-stationary. Also, the Nadaraya-Watson kernel regression smoothing method tends to smooth out maxima and minima within datasets, which can affect the magnitude and time of peak occurrence in sap flow calculations.

The Savitzky-Golay polynomial regression-smoothing filter, which does preserve the maxima and minima of dataset measurements, is only appropriate when the sampling rate is faster than the rate of change in the data. Given the low success rate of packets in the network and tendency for network outages (as seen in Section 5.2.3.2), this method may not be suitable. The orthogonal wavelet analysis, which acts as a high-pass filter, requires a user-defined threshold value. The decision on this threshold value may significantly alter results.

6.5.2 Network costs

For the purposes of comparison, monitoring networks are often broken down into costs per monitored area. Based on the geographic study of the WSN test bed, the monitored areas of sites 2 and 3 can be estimated. The estimated areas of sites 2 and 3 are 632 m^2 and 642 m^2, respectively. Each node in the field and its components (e.g., hardware, power, enclosures, mounting, and sensors) along with the computer (gateway) and base station are included in the estimation for the total cost of the 2012 WSN, which is approximately $25,000 (see Table 15). Based on this estimation, each node on average costs $575 ($350 per relay node and $800 per sensor node). Given the combined areas of sites 2 and 3, the cost per monitored area is approximately \$19.50 m^{-2}.

6.5.3 Natural phenomenon

As discussed in Section 6.4.1.3, a lightning strike hit a section of site 2 in July of 2012. Upon finding the damage on a routine maintenance visit to the site on 25 July 2012, it was discovered that node 2115 was completely destroyed. Pieces from the node enclosure were found up to 5 m away from where the node was originally attached to a tree. Additionally, it was found that the wires, which had connected the soil sensors to the mote, were burnt. A nearby data logger, which was roughly 5 m away from node 2115, was also destroyed. The lid enclosure was blown open and several batteries were found scattered throughout the area. It appears that the electricity from the lightning travelled to the data logger via the soil sensors in the ground. The data logger was attached to a soil sensor, which was buried within 1 m of the soil sensors from node 2115. The lightning damage was evident by the blackening that had occurred around the data logger port that was attached to this soil sensor.

It is unclear as to what the lightning struck directly. It may have hit the ground and travelled through the sensor wires, it may have hit a node enclosure on a tree and travelled into the ground, or it could have struck the tree directly. The findings at the site were inconclusive; however, there are scorched fallen trees within site 2, which appear to be evidence of previous lightning strikes. The damage to both node 2115 and the data logger can be seen in Figure 86.

(a) (b)

Figure 86. (a) Remnants of node 2115 enclosure. (b) EM50 data logger with damage to port 1 (P1) and scorching to enclosure door.

Initially the damage was thought to be restricted to the 2115 node sensors and the uphill data logger sensors; however, it was later found that nodes 2055, 2065, and 2095 all incurred permanent damage to their soil moisture sensors. These three nodes happen to fall within a specific radius of node 2115 (see Figure 87). The soil sensors of nodes outside this radius were still operable; however, as shown in Figure 75 the soil moisture readings in the WSN at 10 cm show an increased level in noise following the lightning strike. It is possible that the effects at 10 cm were residual, since September data suggests that the noise levels begin to reduce.

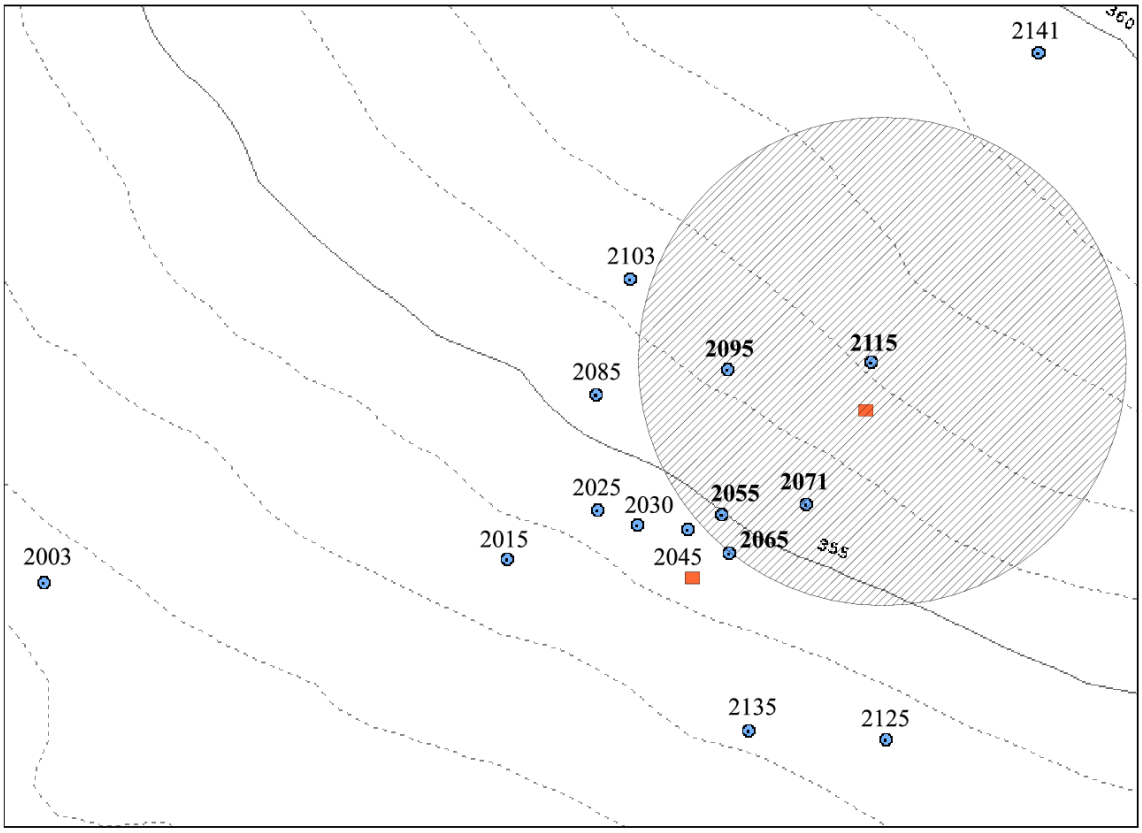

Figure 87. An assumed radius of impact from the lightning based on damaged sensors.

7.0 CONCLUSION

The most cost-effective method for sap flow monitoring is a combination of pulsed and thermal dissipation methods. However, based on the power characteristics of both the pulsed and thermal dissipation methods, only the thermal dissipation method was practical for WSN applications. The heat pulse method is specifically designed to be low power (i.e., it does not require the constant power that the thermal dissipation method uses); however, it was not used in the ASWP test bed because it requires a high sampling rate. Many samples need to be taken in a short window of time after the heat pulse. The way the current system works, every time a sample is taken it must be transmitted. Many samples require numerous transmissions, which would drain mote battery power and cause network congestion. Logistically, the way that the sampling is required for the heat pulse method is too complex for the current commercial WSN systems. Until there is a way to sample multiple times and not transmit after every sample, this is not a realistic means of monitoring sap flow.

It was also found that utilizing common parts from wholesale vendors saves money for large productions. The lab-made sensors are more cost effective than buying commercial grade sensors and purchasing wholesale increases this savings. Compared to the Dynamax TDP30 sap flow sensor, both the HRM and TDM lab-made sensors provide an acceptable alternative to commercial sensors at a fraction of the price. This is particularly true for large deployments, such as the ASWP test bed.

Based on the investigation of the power characteristics of the wireless motes, the sampling frequency for the environmental sensors was determined to be 15 min. This sampling rate was chosen because it was a compromise between mote battery life and the temporal resolution of sap flow data. Based on the low percentage of successful packet transmissions in outdoor networks, which was between 38 and 48% for the pilot test bed and between 45 and 54% for the ASWP test bed, it was found that the 15 min sampling rate was not realized in the data measurements. For comparison, the percentage of valid sap flow data collected from the number of days sap flow was monitored in 2012 ranged from 22 to 56%. Examining the sap flow data by site, the average percent of valid days in site 2 is at 51% where as in site 3 the average percent of valid days is at 33 %. It can be seen that the transmission success rate of mote packets is comparable to the number of successful days of sap flow measurements. Therefore, the transmission rate of the motes in the field is still adequate for sap flow monitoring purposes. The success rate of the network, however, remains an area of concern.

Network stability was found to be a function of the quality of links between nodes in the deployment area. The motes' radio power as well as environmental conditions (including vegetation and radio interference) influence the link quality for transmissions. A mote's radio is powered by the mote battery pack. The transmission distance is dependent on the radio's antenna power and the height of the antenna above the ground. The average distance between motes at the ASWP test bed is approximately 38 m. However, transmission distances have been recorded in excess of 100 m. This is due to the relatively high power level (0 dBm) in the antenna. The antenna also includes a 4.9 dBi gain.

Vegetation, which was anticipated to be problematic, was in fact not an issue. The ASWP test bed has nodes located in open areas and on vegetated hill sides. The dense

vegetation was initially assumed to pose a problem with transmission distances. However, the field data shows that long transmission distances (e.g. approximate 100 m transmission distance between nodes 2115 and 2141 to node 1091) were achieved in a highly vegetated area.

While radio interference was found in the frequency channel the motes transmit through (i.e., channel 12), an investigation found radio interference in all 16 available channels. Background noise can fluctuate in magnitude over daily, weekly, and monthly time periods. Therefore, the noise sampling conducted in the spring and summer of 2012 (at an 18 min sampling rate per channel) does not offer a comprehensive view of the background noise. It has been hypothesized that the levels of noise in channel 12 may influence link quality. Further investigation needs to be conducted to determine if there is a better channel to operate on in the ASWP test bed. Until further studies can take place, channel 12 has been deemed an acceptable channel to operate on in this WSN.

Health statistics data, received by each node in the network, allows the link quality between nodes to be investigated. Metrics for analyzing the communication performance include the received signal strength indicator (RSSI) and transmission path cost. Path cost, which is a general descriptor of the number of transmission links a node is away from the base station, tends to increase the farther away from the base station a node is located. Therefore, drops in the network-averaged path cost tend to correlate with partial network outages. RSSI, which is a descriptor of how well a mote can communicate with its parent, does not exhibit the natural gradation based on node locations with respect to the base station. However, as exhibited in the ASWP test bed, RSSI is shown to increase during low network path cost. This trend suggests that low connectivity between nodes and their parents are present in the farther reaches of the network.

Asymmetry in the link quality between nodes was identified as the cause for duplicate messages received by the base station. Approximately 4% of network data collected was found to be duplicates. While this duplicated data can present bias when studying network performance (i.e., statistics health data) or environmental variables (i.e., sensor data), an analysis on the impacts of duplicates was not found to be significant on soil sensor measurements made at the ASWP test bed.

Route utilization found at the pilot test bed, which could affect up to 50% of a node's battery life, was accounted for in the ASWP test bed by increasing the number of batteries in the nodes in site 1, which are close to the base station. These site 1 nodes are the ones expected to have high route utilization.

It was found that the use of INDAMS, the web-based management system, together with the battery monitoring of network motes, was effective in improving network maintenance. Improving the efficiency of network maintenance in turn helped improve network performance. In addition, it was found that the gateway was an important factor in the success of the WSN test bed. The shortcomings of the Sluggo gateway were addressed in the development of the new Linux gateway, which hosts INDAMS. The Linux gateway, which provides the same collection and network protocol (XServe), also stores the network's health statistics data, which was not collected by the Sluggo gateway. The benefit of having the gateway operate on a PC or laptop is the increase in performance, reliability, and functionality.

Environmental monitoring was successful in terms of capturing both spatial distributions and temporal changes in hydrological sensor measurements at the ASWP test bed. Based on the comparisons between data logger measurements (i.e., high precision) and WSN measurements, the natural trends in the hydrology are captured and are closely correlated between both data

collection methods. The WSN measurements are noisy; however they provide both spatial and temporal views of the data. The data logger measurements are high precision, but they only provide temporal data and show limited spatial distributions of the data. When studying hydrological processes, it is better to have numerous noisy data than to have a few high precision measurements (Li, 2012).

As seen in Chapter 3.0 corrections to attenuated sensor measurements collected during low mote battery power can be used to extend mote operational time. This correction algorithm was not used in the sensor measurements in Chapter 6.0 for two reasons. The first reason is that in field deployments low mote battery power is accompanied by low link quality. As discussed in Beutel et al. (2010), as motes' batteries deplete they lose communication with their neighbors, which is referred to as link failures. This limits the amount of data collected by nodes at low battery power levels, which in turn limits battery life savings by using the sensor attenuation correction algorithm. The second reason the correction is not used is because it is not understood how it works with the EC-5 soil moisture sensors. Therefore, for consistency between the EC-5 and MPS-1 results the correction was not used.

In comparison to a data logger network, WSNs offer a more affordable means of capturing the spatial distributions and temporal changes of environmental variables in a particular area. Startup costs of the ASWP test bed were approximately $25,000. As a point of comparison, the EM50 data loggers used in this study cost $440. The nodes in the WSN cost $350 without sensors. In this case, there is a savings of $90. Each data logger can support up to five sensors, where as the WSN nodes can support up to seven sensors. In addition, the WSN allows for real-time remote monitoring of the variables being studied at the test bed. The data loggers used in this study do not allow for real-time remote monitoring. Also, given the

relatively low cost per additional mote for the network, the WSN can be easily expanded upon. However, the drawback to expanding the network size is that as the network gets larger, it becomes more complex in terms of maintenance. For example, transmission bottlenecking (i.e., communication congestion) and packet transmission losses (i.e., lower success rates of packets being received) become a greater concern. Given the low success rate of mote packet transmissions at the ASWP test bed, which are similar to the success rate of motes at the pilot test bed, it is unlikely that transmission success (which is currently around 50%) would improve in a larger deployment.

WSNs for environmental monitoring are feasible, but the practically of using them depends on many different variables. They are not practical for sensors that require ancillary power sources (i.e., something in addition to battery packs) with higher maintenance than the wireless motes. For example, the sap flow sensors in the ASWP test bed were powered by a 12 V battery. The thermal dissipation sap flow sensors require constant power at around 3 V. In order to have a long-term study of sap flow, a large battery is needed. Until battery life is increased and the size is less cumbersome, it is not practical for outdoor deployments. Batteries are heavy (the size and weight vary by capacity, but the ones currently being used in this study weigh 8.25 kg each) and it can challenging to replace them in a field setting. Maintenance costs can be considerable when performing upkeep for the sap flow sensors. The only benefit of monitoring sap flow in a WSN is that numerous lab-made sap flow sensors can be deployed over a larger area, giving a spatial and temporal view of the plant water usage in the test bed. However, sap flow monitoring through a WSN is not practical in its current state in that maintenance costs are high and the goal in deploying a WSN is to have a low-cost, low-maintenance network.

8.0 FUTURE WORK

After the two-year WSN test bed study period concluded, numerous areas for further investigation were determined based on concerns with the WSN gateway, mote power, WSN communication, and environmental sensors.

One of the main issues with maintaining gateway functionality involves mitigating power outage effects. It is recommended that the gateway be supported by an uninterrupted power supply (UPS). Most power outages that affect the ASWP Nature Center, where the gateway is located, last for relatively short durations (i.e., in the order of minutes to hours). Therefore, a UPS would be able to mitigate most of these outages.

While most of the motes in the ASWP test bed are located in a forested area, there is still a possibility (especially during winter months) of incorporating a method of energy harvesting (e.g., solar panels) at the individual mote scale. There are affordable options for small solar panels (i.e., can fit on mote enclosure lids), which provide a voltage that could either support exhausted wireless motes (during daylight hours) or recharge wireless mote batteries to prolong their life.

As mentioned previously, radio frequency interference may be a concern for WSN communication. Additional studies of the 16 available frequencies needs to be conducted in order to reflect the variability over both daily and seasonal time scales. Another concern with WSN communication is the low success rate of network packet transmissions. Current

developments in WSN technology have provided more intelligent and creative approaches for saving mote battery power and improving transmission success (e.g., data compression and smart sampling techniques). Methods for implementing these improvements can be considered for future phases of the ASWP test bed.

As discussed in Chapter 6.0 the failure of successful sap flow measurements in site 3 may be due to the current sap flow sensor design. It is hypothesized that the probe length is too long for the old growth oak tress being studied in site 3. Experimenting with the probe length in future phases of the ASWP test bed is recommended.

A final recommendation involves the implementation of the low mote battery power attenuation correction algorithm. Currently, the EC-5 soil moisture sensor response to low mote battery power is not well understood. Therefore, it is recommended that the EC-5 sensor be investigated under low battery power conditions. An examination of the potential for the attenuation algorithm in outdoor WSN deployments is also recommended.

BIBLIOGRAPHY

Akyildiz, I.F., W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: a survey," *Comput. Netw.*, vol. 38, no. 4, pp. 393–422, 2002.

Anastasia, G., M. Contib, M. D. Francescoa, and A. Passarella, "Energy conservation in wireless sensor networks: A survey," *Ad Hoc Netw.*, vol. 7, no. 3, pp. 537–568, 2009.

Atmel Corp. (2011, June). 8-bit Atmel Microcontroller with 128KBytes In-System Programmable Flash, ATmega128(L) datasheet. [Online]. Available: http://atmel.com/dyn/resources/prod_documents/doc2467.pdf

Barrenetxea, G., F. Ingelrest, G. Schaefer, M. Vetterli, "The hitchhiker's guide to successful wireless sensor network deployments," in *Proc. 6th ACM Conference on Embedded Network Sensor Systems*, pp. 43–56, ACM Press, New York, NY, USA, 2008.

Barrett, D.J., T.J. Hatton, J.E. Ash, and M.C. Ball, "Evaluation of the heat pulse velocity technique for measurement of sap flow in rainforest and eucalypt forest species of south-eastern Australia," *Plant Cell Environ.*, vol. 18, pp. 463–469, 1995.

Beutel, J., K. Römer, M. Ringwald, and M. Woehrle, "Deployment techniques for sensor networks," in *Sensor Networks: Where Theory Meets Practice*, G. Ferrari, Ed., pp. 219–248, Springer, Berlin, Germany, 2010.

Braun, P., J. Schmid, "Sap flow measurements in grapevines (*Vitis vinifera* L.) 2. Granier measurements," *Plant Soil*, vol. 215, pp. 47–55, 1999.

Burgess, S.S.O., M.A. Adams, N.C. Turner, C.K. Ong, "The redistribution of soil water by tree root systems," *Oecologia*, vol. 115, pp. 306–311, 1998.

Burgess, S.S.O., M.A. Adams, N.C. Turner, C.R. Beverly, C.K. Ong, A.H. Khan, and T.M. Bleby, "An improved heat pulse method to measure low and reverse rates of sap flow in woody plants," *Tree Physiol.*, vol. 21, pp. 589–598, 2001.

Burgess, S.S.O., T.M. Bleby, "Redistribution of soil water by lateral roots mediated by stem tissues," *J. Exp. Bot.*, vol. 57, pp. 3283–3291, 2006.

Burgess, S.S.O., T.E. Dawson, "Using branch and basal trunk sap flow measurements to estimate whole-plant water capacitance: A caution," *Plant Soil*, vol. 305, pp. 5–13, 2008.

Burgess, S.S.O., M.L. Krantz, N.E. Turner, R. Cardell-Oliver, and T.E. Dawson, "Harnessing wireless sensor technologies to advance forest ecology and agricultural research," *Agricultural and Forest Meteorology*, vol. 150, no. 1, pp. 30–37, 2010.

Čermák, J., J. Kučera, N. Nadezhdina, "Sap flow measurements with some thermodynamic methods, flow integration within trees and scaling up from sample trees to entire forest stands," *Trees*, vol. 18, pp. 529–546, 2004.

Chaczko, Z., A. Kale, C. Chiu, "Intelligent health care—a motion analysis system for health practitioners," in *Proc. 6th International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP '10)*, pp. 303–308, Brisbane, Australia, December 2010.

Chen, X., G.R. Miller, *Manufacture of Heater Probes*. Unpublished report, ESPM and Berkeley Atmospheric Science Center Biometeorology Lab, Ecosystem Science Division, College of Natural Resources, University of California, Berkeley: Berkeley, CA, USA, 2012a.

Chen, X., G.R. Miller, *Manufacture of Thermocouple Probes-containing two TC's*. Unpublished report, ESPM and Berkeley Atmospheric Science Center Biometeorology Lab, Ecosystem Science Division, College of Natural Resources, University of California, Berkeley: Berkeley, CA, USA, 2012b.

Chipcon, A.S., "SmartRF CC2420 preliminary datasheet," 2004. [Online]. Available: http://inst.eecs.berkeley.edu/~cs150/Documents/CC2420.pdf

Clearwater, M., F. Meinzer, J. Andrade, G. Goldstein, N. Holbrook, "Potential errors in measurement of nonuniform sap flow using heat dissipation probes," *Tree Physiol.*, vol. 19, pp. 681–687, 1999.

Crossbow Technology, Inc. *MoteView Users Manual (Doc.#7430-0008-05) Rev. A.*, San Jose, CA, USA, 2007a.

Crossbow Technology, Inc., *MPR-MIB Wireless Module Users Manual (Doc.#7430-0021-08) Rev. A.*, San Jose, CA, USA, 2007b.

Crossbow Technology, Inc., *XMesh Users Manual (Doc.#7430-0108-01) Rev. D.*, San Jose, CA, USA, 2007c.

Crossbow Technology, Inc., *XServe Users Manual (Doc.#7430-0111-01) Rev. E.*, San Jose, CA, USA, 2007d.

Crossbow Technology, Inc., *MTS/MDA Sensor Board Users Manual (Doc.#7430-0020-04) Rev. B.*, San Jose, CA, USA, 2007e.

Crossbow Technology, Inc., *MICAz Wireless Measurement System (Doc.#6020-0060-04) Rev. A.,* San Jose, CA, USA, 2007f.

Davis, T.W., C.-M. Kuo, X. Liang, and P.-S. Yu, "Sap flow sensors: construction, calibration, and comparison," *Sensors*, vol. 12, pp. 954–971, 2012a.

Davis, T.W., X. Liang, C.-M. Kuo, Y. Liang, "Analysis of power characteristics for sap flow, soil moisture, and soil water potential sensors in wireless sensor networking systems," *IEEE Sensors Journal*, vol. 12, no. 6, 2012b.

Davis, T.W., X. Liang, M. Navarro, D. Bhatnagar, and Y. Liang, "An experimental study of WSN power efficiency: MICAz networks with XMesh," *International Journal of Distributed Sensor Networks*, vol. 2012, article ID 358238, 14 pp., 2012c.

De Couto, D. S. J. (2004). *High-throughput routing for multi-hop wireless networks* (Doctoral dissertation, Massachusetts Institute of Technology). Retrieved from http://pdos.csail.mit.edu/papers/grid:decouto-phd/thesis.pdf

Edwards, W., P. Becker, J. Èermák, "A unified nomenclature for sap flow measurements," *Tree Phiol.*, vol. 17, pp. 65–67, 1996.

Gay, D., P. Levis, R.V. Behren, M. Welsh, E. Brewer, D. Culler, "The nesC language: a holistic approach to networked embedded systems," *ACM SIGPLAN Notices*, vol. 38, no. 5, pp. 1–11, 2003.

Granier, A., "Une nouvelle méthod pour la mesure du flux de sève brute dans le tronc des arbres," [A new method of sap flow measurement in tree stems], *Ann. Sci. Forest*, vol. 24, pp. 193–200, 1985. [Note: Original paper in French. Translated to English by T.W. Davis.]

Granier, A., "Evaluation of transpiration in a Douglas fir stand by means of sap flow measurements," *Tree Physiol.*, vol. 14, pp. 179–190, 1987a.

Granier, A., "Mesure du flux de sève brute dans le tronc du Douglas par une nouvelle method thermique," [Measurement of sap flow in Douglas-fir tree trunks by means of a new thermal method], *Ann. Sci. Forest.*, vol. 44, pp. 1–14, 1987b. [Note: Original paper in French. Translated to English by T.W. Davis.]

Hamilton, M., E.A. Graham, P.W. Rundel, M.F. Allen, W. Kaiser, M. L. Hansen, and D. L. Estrin, "New approaches in embedded networked sensing for terrestrial ecological observatories," *Environmental Engineering Science*, vol. 24, pp. 192–204, 2007.

Hart, J.K., K. Martinez, R. Ong, A. Riddoch, K.C. Rose and P. Padhy, "A wireless multi-sensor subglacial probe: design and preliminary results," *Journal of Glaciology*, vol. 52, no. 178, pp. 389-397, 2006.

Huber, B. and E. Schmidt, "Eine Kompensationsmethode zur thermoelektrischen Messung langsamer Saftströme," *Ber. Deutsch. Bot. Ges.*, vol. 55, pp. 514–529, 1937.

IEEE 802.15.4 Standard-2003, "Part 15.4: wireless medium access control (MAC) and physical layer (PHY) specifications for low-rate wireless personal area networks (LR-WPANs). IEEE-SA standards board," 2003.

Ingelrest, F., G. Barrenetxea, G. Schaefer, M. Vetterli, O. Couach, M. Parlange, "SensorScope: application-specific sensor network for environmental monitoring," *ACM Transactions on Sensor Networks*, vol. 6, no. 2, article 17, 2010.

James, S.A., M.J. Clearwater, F.C. Meinzer, G. Goldstein, "Heat dissipation sensors of variable length for the measurement of sap flow in trees with deep sapwood," *Tree Physiol.*, vol. 22, pp. 277–283, 2002.

Jiménez, M.S., N. Nadezhdina, J. Čermák, M. Morales, "Radial variation in sap flow in five laurel forest tree species in Tenerife, Canary Islands," *Tree Physiol.*, vol. 20, pp. 1149–1156, 2000.

Juang, P., H. Oki, Y. Wang, M. Martonosi, L.-S. Peh, and D. Rubenstein, "Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with ZebraNet," in *Proc. of the 10th Intl. Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '02)*, pp. 96–107, October 2002.

Kerkez, B., S.D. Glaser, R.C. Bales, and M.W. Matthews, "Design and performance of a wireless sensor network for catchment-scale snow and soil moisture measurements," *Water Resour. Res.*, vol. 48, W09515, doi:10.1029/2011WR011214, 2012.

Krämer, M. and A. Geraldy, "Energy measurements for MicaZ node," Fachgespräch Drahtlose Sensornetze, GI/ITG KuVS, 6 pp., 2006.

Kuo, C., T.W. Davis, C. Tseng, C. Cheng, X. Liang, and P. Yu, Wireless sap flow measurement system. Poster presented at the American Geophysical Union, Fall Meeting, San Francisco, California, December 2010. Abstract retrieved from http://adsabs.harvard.edu/abs/2010AGUFM.H31F1056K.

Lee, H.J. (2007). "RssiSample." *Geeknet, Inc*. [Online]. Available: http://tinyos.cvs.sourceforge.net/tinyos/tinyos-2.x-contrib/stanford-sing/apps/RssiSample/

Li, S.-G. 2012. The Department of Civil and Environmental Engineering Graduate Seminar Series. *Data Intensive Computing for Complex Groundwater Systems and Groundwater-Dependent Ecosystems*. October 19. [Lecture] Pittsburgh: University of Pittsburgh.

Li, X., Y. Deng and L. Ding, "Study on precision agriculture monitoring framework based on WSN," in *Proc. 2nd Intl. Conference on Anti-counterfeiting, Security and Identification (ASID '08)*, pp. 182–185, August 2008.

Liu, Y., Y. He, M. Li et al., "Does wireless sensor network scale? A measurement study on GreenOrbs," in *Proc. 2011 IEEE International Conference on Computer Communications (INFOCOM '11)*, pp. 873–881, April 2011.

Lu, P., “A direct method for estimating the average sap flux density using a modified Granier measuring system,” *Aust. J. Plant Physiol.*, vol. 24, pp. 701–705, 1997.

Lu, P., L. Urban, and P. Zhao, “Granier's thermal dissipation probe (TDP) method for measuring sap flow in trees: theory and practice,” *Acta Bot. Sin.*, vol. 46, no. 6, pp. 631–646, 2004.

Mainwaring, A., D. Culler, J. Polastre, R. Szewczyk, and J. Anderson, “Wireless sensor networks for habitat monitoring,” in *Proc. of the 1st ACM Intl. Workshop on Wireless Sensor Networks and Applications (WSNA ‘02)*, pp. 88–97, September 2002.

Marshall, D.C., “Measurement of sap flow in conifers by heat transport,” *Plant Physiol.*, vol. 33, no. 6, pp. 385–396, 1958.

Martinelli, M., L. Ioriatti, F. Viani, M. Benedetti and A. Massa, “A WSN-based solution for precision farm purposes,” in *Proc. 2009 IEEE Intl. Geoscience and Remote Sensing Symposium (IGARSS ‘09)*, pp. 469–473, July 2009.

Martinez, K., J.K. Hart and R. Ong, “Environmental sensor networks,” *IEEE Computer*, vol. 37, no. 8, pp. 50–56, 2004.

Musăloiu-E., R., A. Terzis, K. Szlavecz, A. Szalay, J. Cogan and J. Gray, “Life under your feet: a wireless soil ecology sensor network,” in *Proc. 3rd Workshop on Embedded Networked Sensors (EmNets ‘06)*, pp. 51–55, May 2006.

Navarro, M., D. Bhatnagar, Y. Liang, “An integrated network and data management system for heterogeneous WSNs,” in *Proc. 8th IEEE International Conference on Mobile Ad-hoc and Sensor Systems (MASS ’11)*, Valencia, Spain, October 2011.

Navarro, M., T.W. Davis, Y. Liang, and X. Liang. (2013). *ASWP: a long-term WSN deployment for environmental monitoring*. Manuscript submitted for publication.

Panchard, J. (2008). *Wireless sensor networks for marginal farming in India* (Master's thesis, Ecole polytechnique fédérale de Lausanne, Lausanne, Switzerland). Retrieved from http://commonsense.epfl.ch/Resources/thesis.pdf.

Pearsall, K., L. Williams, A. McElrone, *The Measurement of Volumetric Water Use in Grapevine with Dual Heat Pulse Sensors*, University of California at Davis: Davis, CA, USA, 2011.

Pennystone Project. (2012). Allegheny County - key soil properties. Retrieved from http://www.pennystone.com/soils/allegheny.php

Ruiz-Garcia, L., P. Barreiro, and J. Robla, “Performance of ZigBee-based wireless sensor nodes for real-time monitoring of fruit logistics,” *J. Food Eng.*, vol. 87, pp. 405–415, 2008.

Rundel, P.W., E.A. Graham, M.F. Allen, J.C. Fisher, and T.C. Harmon, “Environmental sensor networks in ecological research,” *New Phytologist*, vol. 182, pp. 589–607, 2009.

Sakurtani, T., "A heat balance method for measuring water flux in the stem of intact plants," *J. Agric. Meteorol.*, vol. 37, pp. 9–17, 1981.

Selavo, L., A. Wood, Q. Cao, T. Sookoor, H. Liu, A. Srinivasan, Y. Wu, W. Kang, J. Stankovic, D. Young, and J. Porter, "LUSTER: wireless sensor network for environmental research," in *Proc. 5th Intl. Conf. Embedded Networked Sensor Systems (SenSys '07)*, pp. 103–116, November 2007.

Stanjo, F. D. Cvrcek, M. Lewis, "Steel, cast iron and concrete: security engineering for real world wireless sensor networks," in *Proc. 6th Applied Cryptography and Network Security Conference (ACNS '08)*, vol. 5037 of *Series Lecture Notes in Computer Science*, pp. 460–478, Springer, June 2008.

Staszewski, W.J., "Intelligent signal processing for damage dectection in composite materials, " *Composites Science and Technology*, vol. 62, pp. 941–950, 2002.

Swanson, R.H., *An Instrument for Detecting Sap Movement in Woody Plants*, (Station Paper 68), USDA Forest Service, Rocky Mountain Forest and Range Experiment Station: Fort Collins, CO, USA, 1962, pp. 1–16.

Swanson, R.H., "Significant historical developments in thermal methods for measuring sap flow in trees," *Agr. Forest Meteorol.*, vol. 72, pp. 113–132, 1994.

Szewczyk, R., E. Osterweil, J. Polastre, M. Hamilton, A. Mainwaring, and D. Estrin, "Habitat monitoring with sensor networks," *Communications of the ACM — Wireless sensor networks*, vol. 47, issue 6, pp. 34–40, 2004a.

Szewczyk, R., J. Polastre, A. Mainwaring, and D. Culler, "Lessons from a sensor network expedition," in *Proc. 1st European Workshop on Sensor Networks (EWSN '04)*, pp. 307–322, January 2004b.

Teo, A., G. Singh, J.C. McEachen, "Evaluation of the XMesh routing protocol in wireless sensor networks," in *Proc. 49th Midwest Symposium on Circuits and Systems (MWSCAS '06)*, pp. 113–117, IEEE, San Juan, Puerto Rico, August 2006.

"*TinyOS*", 2010, http://tinyos.net/.

Tolle, G., J. Polastre, R. Szewczyk, D. Culler, N. Turner, K. Tu, S. Burgess, T. Dawson, P. Buonadonna, D. Gay, and W. Hong, "A macroscope in the redwoods," in *Proc. 3rd Intl. Conf. Embedded Networked Sensor Systems (SenSys '05)*, pp. 51–63, November 2005.

Trubilowicz, J., K. Cai, and M. Weiler, "Viability of motes for hydrological measurement," *Water Resour. Res.*, vol. 45, W00D22, 6 pp., 2009.

Werner-Allen, G., J. Johnson, M. Ruiz, J. Lees, and M. Welsh, "Monitoring volcanic eruptions with a wireless sensor network," in *Proc. 2nd European Workshop on Wireless Sensor Networks (EWSN '05)*, pp. 108–120, 2005.

Werner-Allen, G., K. Lorincz, M. Welsh, O. Marcillo, J. Johnson, M. Ruiz, and J. Lees, "Deploying a wireless sensor network on an active volcano," *IEEE Internet Comput.*, vol. 10, no. 2, pp. 18–25, 2006.

Wilson, K.B., P.J. Hanson, P.J. Mulholland, D.D. Baldocchi, S.D. Wullschleger, "A comparison of methods for determining forest evapotranspiration and its components: Sap-flow, soil water budget, eddy covariance and catchment water balance," *Agric. For. Meteorol.*, vol. 106, pp. 153–168, 2001.

Wullschleger, S., K. Childs, A. King, P. Hanson, "A model of heat transfer in sapwood and implications for sap flux density measurements using thermal dissipation probes," *Tree Physiol.*, vol. 31, pp. 669–679, 2011.

Xu, N., S. Rangwala, K.K. Chintalapudi et al., "A wireless sensor network for structural monitoring," in *Proc. 2nd International Conference on Embedded Networked Sensor Systems (SenSys '04)*, pp. 13–24, ACM Press, New York, NY, USA, November 2004.

Ye, W., F. Silva, A. DeSchon, and S. Bhatt, "Architecture of a satellite-based sensor network for environmental observation, " in *NASA Earth Science Technology Conference* (2008), 6 pages.

Ye, W., J. Heidemann, and D. Estrin, "Medium access control with coordinated adaptive sleeping for wireless sensor networks," *IEEE/ACM Trans. Netw.*, vol. 12, pp. 493–506, 2004.