Title
A cyber-infrastructure for the measurement and estimation of large-scale hydrologic processes

Author
Kerkez, Branko

Publication Date
2012

Peer reviewed|Thesis/dissertation

A cyber-infrastructure for the measurement and estimation of large-scale hydrologic processes

By

Branko Kerkez

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Civil and Environmental Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Steven D. Glaser, Chair
Professor Alexandre M. Bayen
Professor Kristofer J. Pister

Fall 2012

A cyber-infrastructure for the measurement and estimation of large-scale hydrologic processes

Abstract

A cyber-infrastructure for the measurement and estimation of large-scale hydrologic processes

by

Branko Kerkez

Doctor of Philosophy in Civil and Environmental Engineering

University of California, Berkeley

Professor Steven D. Glaser, Chair

Water shortages, particularly evident in the state of California, emphasize the need for a better hydrologic understanding, and improved water management techniques. The majority of the state's water originates in the Sierra Nevada as snow, melting throughout the year to meet the needs of various stakeholders. Current measurement techniques are unable to resolve variability of the snowpack at the basin scale, and snowmelt processes are not well captured by existing hydrologic models. A system-level solution is introduced to facilitate scientific understanding and water management decisions in basins of the Sierra Nevada. The core of this thesis focuses on expanding current sensing methods at the hydrologic catchment (1-2km^2) scale to develop an improved understanding of the mountain water balance. It is shown that Wireless Sensor Networks (WSNs) offer an ultra low-power, cost-effective solution to instrument catchment-scale regions. An explicit WSN deployment strategy is derived, leveraging in-situ network statistics of the *Packet Delivery Ratio*, and *Received Signal Strength Indicator* to optimize network performance. A hydrologic variability analysis is carried out on the dataset collected by the network, showing temporal stability in the variability of snowdepth at the catchment scale, and validating the use of a stratified sampling approach based on the even instrumentation of physiographic variables. To derive estimates of Snow Water Equivalent (SWE) at the larger basin scale, an estimation framework based on Gaussian Processes, and an optimal sampling strategy, which maximizes *Mutual Information*, is described. The sensor placement method outperforms a number of other sampling designs, reducing estimation error by up to 100mm. A computationally tractable, *Hybrid System* model of snow dynamics is introduced and shown to accurately reflect the various stages of snowmelt when compared to observations in the Sierra Nevada. The need for improved sensing and estimation procedures is highlighted by an analysis of the effects of improved SWE estimates on basin-scale streamflow forecasting methods. The use of historically reconstructed SWE data shows that streamflow forecasting error could be reduced by nearly 10% through improved SWE estimates.

To my family.

Contents

List of Figures and Tables

Acknowledgements

This thesis is the direct culmination of countless conversations with fellow colleagues, mentors and friends. My experience at Berkeley has been everything I ever hoped it would be and I would do it all again if I could. First and foremost, I would like to thank my adviser, Dr. Steven Glaser, for his continuous support, guidance, and encouragement. I have tremendous respect for his vision to incorporate holistic, systemic thinking into Civil and Environmental Engineering, and it is my sincere belief that his impact on CEE will become lasting and tremendously respected. Observing Steve manage various cross-disciplinary and often orthogonal projects encouraged me to branch out of my comfort zone to explore domains and tools which would have surely escaped me otherwise. I want to thank Steve for giving me the freedom to pursue what may have at the time appeared as whimsical, high-risk topics. Working in the Glaser lab allowed me to reinvent my pursuits, and ultimately discover research topics that I am truly passionate about. Above all, our many espresso-fueled conversations have resulted in what is surely to be a lasting friendship.

I would like to thank Professor Roger Bales for exposing me to a broad suite of hydrologic research topics. His passion for his discipline and his ability to lead the deployment of unprecedented hydrologic measurement campaigns were truly inspiring. Working with Roger has taught me to think big, and I am very gretaful for his support of our systems-based research goals.

I want to thank Professor Kris Pister for hosting me in his EECS research group. Learning from his expertise in Wireless Sensor Networks, and collaborating with members of his research groups has significantly strengthened my understanding of WSN topics, and has vastly improved the quality of our WSN deployments. Above all, I have been inspired by his dream-big mentality and his encouragement of independent thinking. I would also like to thank him for teaching me the value of counting micro amps.

I want to thank Dr. Alex Bayen for his one-on-one conversations, which began even before I was admitted into the CEE program. I have tremendous respect for his drive, and his frankness has been extremely valuable in guiding my research and career plans. I would also like to thank Professor John Dracup. I have very much enjoyed our numerous conversations and collaborations on hydrologic systems. I want to also thank Raja Sengupta for our many conversations and ideas regarding system-level philosophies.

I would like to thank Fabien Chraim for his friendship and collaboration. I want to thank Stephen Welch for his collaboration on our optimal sensor placement strategies. I hope the algorithms resulting from our collaborations result in similarly positive impacts on the development of acoustic instrument pickups. I want to thank Matt Meadows for his continuous support of our field-operations, as well as his patience in the early years of our WSN deployments. I would like to thank Edouard Bulteau for working with me on the impact of SWE data on real-time water resource operations. I would also like to thank my friend Jerry Jariyasunant for his collaboration on our work on real-time transportation applications.

I want to thank Tomas Watteyne for his patience and many lessons on embedded programming. Our collaboration on OpenWSN has been instrumental to my betterment as a programmer and researcher. I am continuously impressed by his organizational skill and passion for his work. I want to thank the National Science Foundation for supporting my work through so many channels, spanning REUs, graduate fellowships, and research grants.

I also owe much to my dear friends - old and new. Thank you so much for all the years of support and fun. I want to thank my friend Tom Icard, for our many road trips and musical adventures. I hope to continue our musical collaborations and finally get that stage gig. Our Ragtime ensemble would also not be complete without Greg McLaskey, a great friend and lab mate.

I want to sincerely thank my family. My parents, Milenko and Smiljana Kerkez, have left behind their own aspirations to enable my own. Thank you for believing in me, and allowing me to pursue my aspirations without any external pressures. You are the best parents and anyone could ask for. I want to thank my brother Vladimir for his support and patience. I am very proud of the person you have become, and I will look back at our past five years as one of the best of my life. Last but not least, I want to thank my beautiful wife Lindsey for her relentless support. Needless to say, I would not have been able to do any of this without you.

Chapter 1: Introduction

Traditional water systems were built to serve the "analog economy," where various ownership and management boundaries are loosely interconnected to meet the needs of an oblivious water consuming public. Inefficiencies are inherent in such a system, since little or no coordination exists between stakeholders. Though significant strides have been made in Civil Engineering to study water networks, as well as the hydrologic systems that feed them, the tools necessary to tie this domain knowledge into a real-time, information-driven system are still lacking. This is especially evident in the state of California, where, by most metrics, the current water infrastructure is struggling to meet the requirements imposed by a growing population and a thriving agricultural sector [*WEF,* 2011; *ACWA,* 2011; *Sierra Club,* 2011]. Crippling droughts and devastating floods are one motivating example that informed decision-making will require new management techniques to ensure equitable and reliable distribution of water supplies. As the following sections will show, an improvement to current water supply forecasting, massive in scope and spanning various areas of science and engineering, has the potential significantly improve water resource management.

As with all substantial infrastructure problems, a solution to large-scale water resource management cannot be addressed through an individual treatment of sub-domain problems, but must rather be tackled through a holistic and systemic approach. This thesis addresses one aspect of this system - the theoretical and practical underpinnings of a system-level architecture for deriving real-time water balance estimates in large-scale water basins. As with most complex systems, detailed spatial and temporal data, scientific domain knowledge, and robust estimation procedures are required to meet the needs of such a framework. This framework will be applied directly to water supply forecasting in California's Sierra Nevada, showcasing the direct benefits of improved sensing, estimation, and modeling to short-term and long-term water supply forecasting.

1.1 Water supply forecasting in California

The current population of California is estimated at 36 million, and is expected to grow to 60 million by 2050 [*California Water Plan*, 2009]. The state boasts one of the most productive agricultural sectors in the world, and contains a plethora of forest terrain and environmental habitats. Aside from weathering historic drought seasons, flood risks are growing due to aging levees, extreme weather scenarios, and climate change. California's water infrastructure composes a broad, multi-use system that balances the water supply to meet goals of flood-control, irrigation, recreation, hydropower generation, and other uses. Water demands in the state are met by a combination of federal, state, and local management projects. The majority of the supply is dedicated to meeting agricultural irrigation needs, with a smaller fraction going to urban use. Historically, 10-30% of the state's energy is also generated via hydropower. The major hydrologic parameter associated with hydropower production is snowfall [*Aspen*, 2005].

1.2 Current management practices

It is estimated that seasonal snow cover is the primary source of water for over 60 million people in the western United States, and that melting snow is responsible for 80% or more of soil moisture and stream flow in semi-arid mountain basins [*ACWA*, 2011; *Daly et al*., 2001; *Marks et al.,* 2001]. Globally, seasonal snowpack and glaciers provide water to more than one sixth of the world's population and may be at risk in a changing climate [*Barnett et al*., 2005]. A closer look at the problem traces the origins of California's water supply to the Sierra Nevada, a mountain range spanning a 600 km north-south transect of the state. Precipitation is deposited in the wintertime as snow, which acts as storage, and melts throughout the year to provide a steady supply of water through a network of streams and rivers. Significant state and federal resources are devoted annually to the monitoring of snow related phenomena in the Western US.

Currently, the main source of snow data is provided by a network of manual and automated measurements managed by the US Department of Agriculture, the National Weather Service, and the National Resource Conservation Services (Figure 1.1) [*California Cooperative Snow Surveys*, 2012; *Rice and Bales*, 2010; *SNOTEL,* 2012]. The manual measurements are known as *snowcourses,* pre-selected locations where observations of snow-depth and snow water equivalent (SWE) are made by human surveyors on a monthly basis during the winter period. Snowcourses were selected to be representative of particular water-producing regions. These locations were chosen beginning in the 1930s through empirical means, with the main criteria of general accessibility and limited public disturbance [*SNOTEL,* 2012]. Automated measurements are also made by a number of permanently located *snow pillows*, large liquid filled pillows which convert the pressure exerted by accumulated snow to a representative SWE estimate.

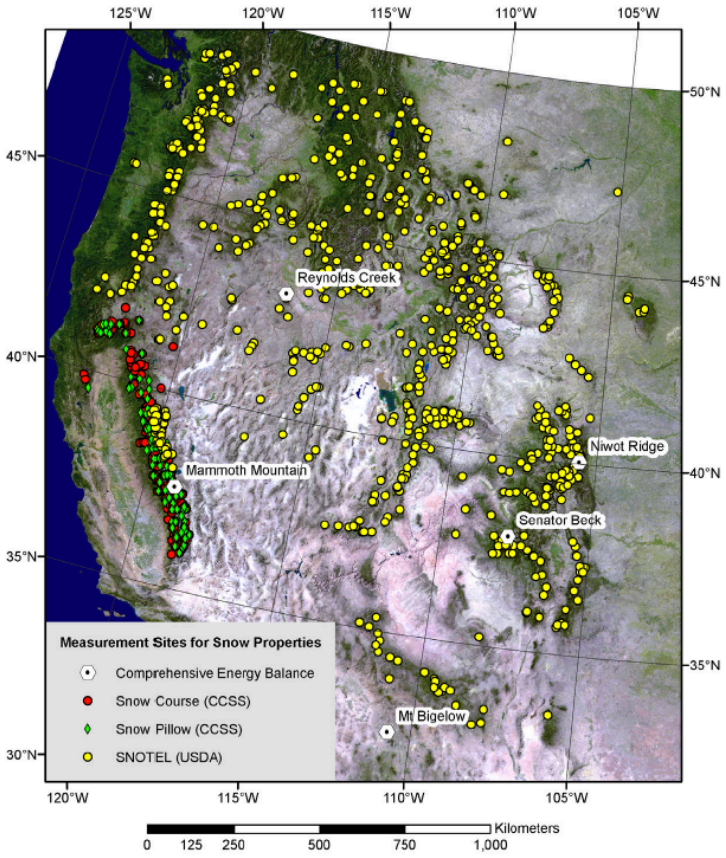

Figure 1.1 Location of snow course, snow pillow, and SNOTEL sites across the Western US. Data source: *Bales et al.* (2006).

Presently, operational water management decisions, and specifically those relating to control of reservoirs and dams, rely on two predictive models. The *California Department of Water Resources* (DWR) maintains an annual statistical forecast, leveraging snow course historical data, and correlations between previously observed snowdepth readings and stream runoffs [*California Cooperative Snow Surveys*, 2012]. This predictive method relies heavily on multivariate regression techniques to predict the total volume of the April-July streamflow based on the April 1st snowdepth surveys [*Garen,* 1992]. Recent warming trends are challenging the stationarity assumption of this statistical approach, suggesting that historical correlations may no longer be the best proxy for predicting future behavior [*Bales et al*., 2006].

The second source of stream flow predictions is provided by the National Weather Service's River Forecast Centers, which employ a physically-based process model to provide ensemble forecasts of river levels and reservoir inflows [*Hartman et al.,* 2002]. The DWR and NWS predictions are used by various agencies and decision makers to allocate water resources and to dictate long-, and short-term reservoir releases.

1.3 Limitations of current practices

Two primary challenges arise in present water management in the American West: limited understanding of mountain hydrologic processes, and sparse data collection. While the

study of hydrologic phenomena is a relatively mature field, the hydrologic processes pertaining to water and energy fluxes in mountainous regions, where basins are largely dominated by snowmelt, are not well captured by current hydrologic models [*Bales et al.,* 2006]. Statistical forecast models tend to perform well when predictions are made within an acceptable window around the historical mean. A continuation of the current warming trend could lead to significant deviation from this historical mean, and has the potential induce a significant amount of error into statistical forecasting methods. Aside from impacting model forecasts, further warming will shift the rain-snow transition zone to higher elevations, thus accelerating the snow melting process. Rapid melting of the snowpack could reduce the amount of water available during spring and summer months, and has the potential to reduce hydropower production between 8% to 50% [*Aspen,* 2005]. An improved hydrologic understanding hinges directly upon a broader scientific initiative to measure and understand the spatiotemporal coupling of snowmelt driven processes.

While snow distribution is known to vary significantly spatially, even along short transects [*Meromy et al*, 2011; *Meadows et al.,* 2010], current monitoring strategies still only provide very sparse spatiotemporal data. It has been shown that the combination of snow-course and snow-pillow infrastructure is unable to resolve SWE variability at a basin scale [*Bales et al.,* 2006; *Molotch et al.,* 2005a], which adversely affects even the most advanced hydrologic models. These measurement locations were selected in part due to their ease of accessibility (flat terrain, for example) and thus suffer from a lack of the ability to capture the range of possible terrain seen in most basins. In some cases, point readings were found to be 200% greater than the their respective basin mean [*Molotch et al.,* 2005b]. Dense, large-scale distributed measurements are thus required to capture the spatiotemporal variability of the snowpack.

Developing an improved scientific understanding of mountain hydrologic processes, and thus improving real-time SWE estimates and supply forecasts is a two-part effort:

1. Expanding and improving ground-based snow measurement methods to more accurately capture hydrologic variability.
2. Upscaling and distributing these ground measurements to improve basin-scale estimates of water quantities.

These two problems are equally important, inherently coupled, and must be solved in tandem.

Scaling measurement campaigns to the size of hydrologic basins imposes unprecedented demands on sensing hardware, which must be met if the deployment scheme is to be effective and sustainable. This, in part, explains the sparse set of currently available measurements, since sensor network deployments at this scale are subject to significant constraints in energy, accessibility, sensor coverage area, and cost. Given the size and remote location of mountain basins, it is often challenging, and inconvenient, to physically collect cached data, readily detect faulty hardware, and alter equipment parameters. Additionally, when the utility of sensor data extends beyond scientific purposes, as it does in the case of water management, sensor data must be available in real-time. Many off-the-shelf data-logging components do not lend themselves to real-time data acquisition at this scale, as wiring runs are limited by cost, signal noise and the need to maintain acceptable levels of current. Transmission along wires is further complicated by heavy snow loads, and the proclivity of rodents to chew on cables. This renders a wired monitoring solution practically unfeasible. The use Wireless Sensor Networks (WSNs) offers a viable solution, by enabling distributed sensing of large regions in real-time and at extremely low power consumption.

1.4 Thesis contributions

This thesis lays out a system-level approach to facilitate the sensing, estimation, and modeling of large-scale environmental phenomena, focusing specifically on snowmelt-driven hydrologic processes in the Sierra Nevada of California. My approach can be summed up through a set of core contributions, and will emphasize the need to treat small-scale (catchment) processes first, before expanding to the study of larger-scale (basin) phenomena:

- **Validation of wireless sensor network technology for real-time measurements of hydrologic processes**
  - The successful design, deployment, and operation of an embedded sensor network (57 wireless nodes and over 300 sensors) to provide real-time measurements of catchment-scale snow depth and soil moisture in the Sierra Nevada of California.
  - The development a three-phase procedure to optimize the WSN deployment by focusing on the need to quantitatively evaluate in-situ network performance.
  - Introduction of the use of the *packet delivery ratio* and *received signal strength indicator,* as valuable a deployment metric for environmental WSNs. Using these metrics, we show that for our site conditions, a conservative 50-m node-to-node spacing ensures low-power, reliable, and robust network communications.
- **Improved scientific understanding of catchment-scale hydrologic processes**
  - Demonstration that a stratified sensor placement approach can effectively capture the catchment-wide mean and variability of snow depth and soil moisture, when compared to three comprehensive gridded synoptic and LIDAR surveys.
  - The spatiotemporal analysis of snow depth variability, revealing that canopy cover is the major explanatory variable of snow depth, and that under-canopy measurements persistently show higher variability compared those in open terrain.
  - An analysis of soil moisture patterns, showing lower variability at deeper soil depth, and a correlation between mean soil moisture and variability for shallow soils.
- **Development of a sampling design and sensor placement framework for basin-scale hydrologic estimation**
  - The introduction of the first comprehensive, quantitatively driven sensor placement (sampling design) strategy for basin-scale SWE estimation. The procedure uses *Gaussian Process* models to derive estimates of basin-wide SWE, and selects the optimal sensing locations by coupling rank-based clustering with *Mutual Information*-based optimization.
  - Use of an 11-year data set, reconstructed by carrying an energy balance on satellite data in the American River basin, to show that that our sampling design provided the best placement strategy, improving RMSE between 0 and 100 mm compared to a number of other placement approaches.
  - A comparison of rank-based clustering to geographically based clustering (sub-basin delineation) to validate the existence of stationary covariance structures within the overall SWE dataset.
- **Development of a hybrid systems snowmelt model**
  - The presentation of a direct application of *Hybrid Systems* for the modeling of the seasonal snowmelt cycle, showing that through the *Hybrid Systems* framework it is possible to significantly reduce the complexity offered by existing models.

  - Demonstration that currently existing heuristics can be embedded into a coherent mathematical framework to enable powerful analytical techniques while preserving physical intuition about the system.
  - Modeling snowmelt as a three state hybrid automaton, representing the sub-freezing, sub-saturated, and fully saturated physical states, and showing that the model accurately reproduces melt patterns, by simulating over actual data sets collected in the Sierra Nevada.
- **Study of the effects of improved SWE estimates on basin-scale streamflow forecasting**
  - The sensitivity analysis of streamflow forecasting methods in the Sierra Nevada, showing that a given error in SWE estimates can manifest itself in a two-fold error in streamflow forecasting.
  - Demonstration on the effects of improved SWE data on current regression methods, showing that better SWE data can improve streamflow forecasting by nearly 10% compared to presently used methods.
  - The investigation of the effects of improved SWE data on a physical process model of the basin water balance, showing that the use of better SWE estimates can improve the precision of model outputs, while reducing error variance by 50%.

1.5 Organization

The core of this thesis focuses on improving and expanding current sensing methods, while developing the tools necessary to utilize the resulting data for large-scale environmental modeling. The viability of wireless sensor networks as a distributed real-time, catchment-scale measurement platform will be investigated in Chapter Two. Larger-scale hydrologic systems will be addressed in Chapter Three, where a *Gaussian Process* estimation algorithm, along with an *Information Theoretic* sensor placement framework will be employed to derive better basin-scale SWE estimates. To facilitate streamflow forecasting, a physically-based snowmelt model will then be introduced in Chapter Four. Finally, to validate the need for an expanded sensing infrastructure, the benefits of improved SWE estimates on present forecasting methods will be investigated in Chapter Five. Concluding remarks will be offered on the performance of the discussed methods, and future work will be examined in the context of tying these methods into a formal systems framework that incorporates downstream operations and man-made infrastructure.

Chapter 2: Wireless Sensor Networks for Catchment-scale Snow and Soil Moisture Measurements

The Earth's *Critical Zone* encompasses the complex interaction of near-surface biological and chemical processes, involving ecosystems, soil, water, and the atmosphere [*Levia et al*., 2011]. Dynamics of hydrologic systems in large-scale mountain basins (1000 km^2) are thus inherently guided much smaller-scale *Critical Zone* processes at the catchment level (1-2km^2). The primary objective of this chapter is to describe new instrumentation techniques that can be employed to carry out long-term, real-time studies of these environments, and to investigate the spatial variability of snowdepth and soil moisture at the catchment scale.

2.1 Introduction

A comprehensive understanding of the mountain water cycle demands a deeper insight into the spatiotemporal coupling between various hydrologic processes at the headwater-catchment, and larger-basin scales. This is particularly true in the mixed-conifer zone of California's Sierra Nevada, a productive ecosystem situated in the mid-elevation rain-snow transition zone, where rain falls at elevations below 1500 m, and snow accumulates above 2200 m [*Bales et al.,* 2006]. It has been noted that this transition zone is sensitive to annual temperature fluctuations, both in terms of the quantities of accumulated snow, as well as melt timing [*Christensen et al.,* 2008]. A complete understanding of the links between snow cover and soil moisture on the forest water cycle is still lacking.

One of the major challenges in hydrologic science relates to characterizing and understanding spatial variability [*Bogena et al.,* 2010]. Spatial measurement of hydrologic processes at the catchment, and basin scales is subject to significant constraints in energy, accessibility, sensor coverage area, and cost. Many off-the-shelf data-logging components do not lend themselves to real time data acquisition when sensors are spatially distributed at the km scale, as wiring runs are limited by prohibitive cost, signal noise and the need to maintain acceptable levels of current. Transmission along wires is further complicated by heavy snow loads, and damage by wildlife. Recent advances in sensing technology, particularly in the area of Wireless Sensor Networks (WSNs), now enable environmental monitoring in real time, and at unprecedented spatial and temporal scales. While access to real time data reduces site visits and inconvenience experienced during scientific research campaigns, a major benefit of distributed real time data relates to decision-making support and operational hydrology. *Rice and Bales* [2010] analyzed the performance of a prototype WSN and concluded that the technology can be effectively used to capture spatially representative measurement of snow depth with relatively few sensors.

For most spatially distributed sensor deployments it is often challenging, and inconvenient, to physically collect cached data, readily detect faulty hardware, and alter equipment parameters. While there are many commercial wireless systems that can transmit data between two points, such hardware generally use relatively high-powered radios and requires substantial energy for transmission. Such point-to-point transmission systems operate

independently, and can interfere with each other if multiple links are deployed within a region. WSN technology offers an alternative for cost-effective instrumentation of extended regions with limited accessibility, while permitting real time data access to distributed sensors through a centralized communication architecture. Given the nascent nature of low-power wireless, in particular for remote field applications, off-the-shelf performance cannot be guaranteed, and much work remains to be conducted to quantify the real-world performance of WSNs for hydrologic monitoring.

The research reported in this chapter aims to develop techniques for efficient, scalable, and robust WSN deployments for monitoring hydrologic phenomena, while using methods that can readily be extended to many environmental monitoring applications. We address four specific questions: First, what soil moisture and snow depth variability patterns are revealed by strategically distributed sensors in a densely instrumented mountain catchment, and how do these sensor readings compare to gridded surveys and LIDAR data? Second, what is the performance of a large-scale, low-cost WSN built from off-the-shelf hardware, when exposed to the harsh conditions in the snow-covered Sierra Nevada? Third, which metrics can be used to quantify WSN performance and to evaluate the design of the wireless monitoring system? Fourth, what is an efficient approach to designing and deploying WSNs for long-term environmental monitoring campaigns, while maintaining robust and reliable network performance?

2.2 WSN design principles

The core component of a WSN is known as a *mote* - a tiny, ultra-low power radio operated by a micro-controller, and featuring analog and digital interfaces to which sensors can be connected. A mote exchanges information wirelessly with neighboring motes in a distributed network, which can relay this information to their neighbors, until the data reaches a central hub. Software optimization of power control (also known as duty cycling) allows some devices to operate for several years on a pair of AA batteries [*Karl and Willig,* 2005; *Dust,* 2006]; embedded software can be optimized to permit motes to remain in an extremely low-power state the majority of the time (powering the most necessary components, such as the clock and basic micro-controller features), waking up only periodically to transmit or receive data [*Dust,* 2011; *IEEE,* 2009]. The recent mass adoption of these devices, particularly for industrial applications, has made them an extremely cost-effective alternative to placing wires [*Emerson,* 2010; *Honeywell,* 2010; *ISA,* 2009; *Song et al.,* 2008].

Depending on the implementation and protocols of the *network stack* (the general hardware and software architecture), motes can aggregate into a number of network topologies. The two most common are the *star*-, and the *mesh*-topology (Figure 2.1). In a *star* configuration, each mote exchanges information only with a central base station. In the majority of such implementations, the base station is programmed to keep its radio on continuously to listen for incoming transmissions, which results in high energy usage. The span of a *star*-network is limited by the distance of a single link between a mote and its network manager. The network manager is thus often placed in a central location. While it is possible to increase transmission distance using more powerful radios, star networks are still limited by the lack of radio path redundancy and the need to communicate with one central location. If the region being instrumented cannot be covered by a single *star*-network, the *ZigBee Protocol* specifies a hybrid topology, which enables *router nodes* to be added to the network to extend coverage [*ZigBee Alliance*, 2008; *Bogena et al.,* 2010]. These nodes serve as range extenders, typically using more radio resources to listen for incoming transmissions from nearby nodes, and then forwarding this

data to the base station or other nearby router nodes. Given their specialized function, and higher power consumption, most networks use limited router nodes, which limits path redundancy to the base station.

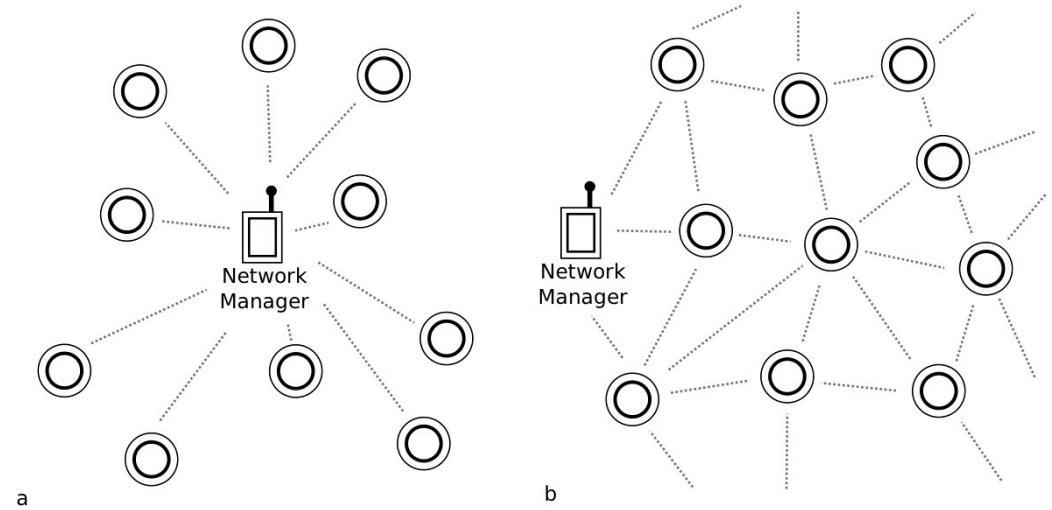

Figure 2.1 WSN network topologies. (a) WSN *star-topology*: all motes have a direct link to the network manager; the span of this network is limited by the distance between a mote and the manager. (b) WSN *mesh-topology*: motes exchange information with neighbors to create a redundant *multi-hop* network.

In a *mesh*-based topology, motes communicate with multiple neighbors to create an internally redundant *multi-hop* network. Multiple paths between network nodes allow data to be transmitted even if certain network links fail [*Karl and Willig,* 2005]. Although requiring a relatively larger software overhead, mesh networks can adapt their topology to reflect varying attenuation of radio signals caused by changes in the environment, or changing needs of the operator. The *mesh*-topology thus allows the network to span much larger areas, compared to a *star* layout, as out-of-range motes can exchange data by transmitting, or *hopping*, it via any neighboring motes in the mesh.

Data transmitted within WSNs is quantized into formatted data blocks called *packets* [*Tanenbaum,* 2003]. A *successful transmission* is achieved if the transmitting mote receives an acknowledgment from the receiving mote, notifying it that the transmitted data packet has been received correctly. This acknowledgment redundancy, while not built into all available WSN platforms, ensures that packets reach their intended destination, even if retransmission is required.

Since remote wireless networks operate on battery power, overall power consumption becomes a critical constraint. As radio transmissions consume more than 95% of a mote's of power use [*Dust,* 2006] (not including sensors), it is important to design WSNs that will minimize retransmissions. By optimizing use of radio resources, and ensuring low power use, motes can be added to an existing data logging infrastructure to provide real time wireless data access, without significantly increasing overall power consumption. Extreme weather conditions, especially cold weather, can take significant tolls on hardware and batteries [*Hasler et al.,* 2008; *Mainwaring et al.,* 2002]. Successful design and operation of such a field-deployed WSN hinges upon a rational set of metrics that can be used to quantify performance.

Communication in WSNs is challenged by multipath radio propagation and narrow-band interference [*Karl and Willig,* 2005; *Watteyne et al.,* 2010] where topography can result in radio signals bouncing off ambient surfaces, causing phase-shifted copies of the same signal to arrive at the receiver antenna. This can lead to destructive interference, which cancels out the original signal, effectively eliminating the ability of certain mote pairs to communicate. The adverse effects of multipath propagation on radio communications are a function of the deployment environment, and it has been shown that even slight changes in node-to-node distance (on the order of centimeters) can have significant impact on this behavior [*Watteyne et al.,* 2009].

The IEEE802.15.4 standard divides the 2.4 GHz band into 16 transmission channels (or, narrow-band sub-frequencies) that can be used to mitigate communication challenges posed by external interference and multipath propagation [*IEEE,* 2009]. Improper channel selection can cause data to be lost during transmission, leading to the need to retransmit data, and manifesting itself in much higher energy requirements on the network [*Watteyne et al.,* 2009]. A common approach to address this issue is to conduct an expert survey, which entails a physical visit to the field to select the frequency channel on which the loss of transmitted packets will be minimized. This approach has major drawbacks, as research has shown that the optimal channel may vary over time [*Kerkez et al.,* 2009; *Watteyne et al.,* 2010]. It has been shown that in arboreal environments propagation of radio waves is adversely affected [*Oestges et al.,* 2009], and fluctuations in the surrounding environment can cause poor performance on channels that originally performed well. An option to address this time-varying behavior involves equipping WSNs with the ability to *channel hop*. In such deployments, motes within the network randomly select one of the available channels every time a transmission occurs, rather than persistently transmitting information on a single channel. While still an area of ongoing research [*Kerkez et al.,* 2009; *Watteyne et al.,* 2010], and requiring slightly larger implementation overhead, channel hopping has been shown to reduce the effects of multipath propagation and external interference, thus improving network reliability and battery lifetime

A number of commercial WSN solutions are available [*Dust,* 2011; *Ember,* 2011; *Jennic*, 2011; *Libelium,* 2011; *Nevis,* 2011; *Sentilla,* 2011], each offering distinct platforms with unique functionality, network stacks, and standards. The implementation of algorithms for network topology, routing, and channel hopping is a nontrivial task, but advances in WSN technology have alleviated, although not completely removed, many of these challenges. No simple off-the-shelf solution exists for environmental monitoring applications, and hardware performance can vary significantly depending on site conditions.

Deployments of WSNs have ranged from data center HVAC control [*Bell and Federspiel,* 2009] structural health monitoring [*Kim et al.,* 2007; *Rice and Spencer,* 2008], and military applications [*Culler et al.,* 2001]. Notable deployments for environmental monitoring purposes include habitat monitoring [*Hart and Martinez,* 2006; *Mainwaring et al.,* 2002; *Ramanathan et al.,* 2006], permafrost detection [*Hasler et al.,* 2008], the study of mountain ranges [*Ingelrest et al.,* 2010], and for purposes of snow depth monitoring [*Rice and Bales,* 2010]. A successful short-term deployment of a WSN for the monitoring of hydrologic phenomena was conducted by *Trubilowicz et al.* [2009], who noted that the technology they tested lacked ease of use and reliability. Further examples of wireless deployments for ecological and environmental monitoring were conducted by *Etzel and Braun* [2005], *Porter et al.* [2005], and *Szewczyk et al.* [2004]. Most recently, the efficacy of a large *ZigBee*-based network for monitoring soil moisture variability was also verified by *Bogena et al*. (2010). The hardware and software implementations in these studies varied significantly, making it difficult to gauge which

hardware, software, and network protocol implementation will perform the best for a given deployment. The current lack of the out-of-the-box usability, coupled with the vast choice of available hardware and network protocols, has the potential to become a severe time drain for scientists designing an environmental monitoring system.

2.3 Methods

2.3.1 Network design

Twenty-three locations were selected and instrumented to monitor water-balance variables in a remote, forested, headwater catchment along a 1.5-km transect in the Southern Sierra Critical Zone Observatory (CZO) **(**$37^o04'$ *N*, $119^o11'$ *W***),** which is co-located with the Kings River Experimental Watershed (KREW) [*Bales et al*., 2011; *Levia et al*., 2011]. KREW spans the rain-snow transition zone, with lower elevations receiving more precipitation as rain. The WSN is located in a northern, relatively higher elevation sub-catchment of KREW, where the majority of annual precipitation falls as snow. Site elevations in the WSN-instrumented catchment range from 1950 to 2010 m, with landscape varying over dense mixed-conifer forest (76-99%), open meadows, and across mixed chaparral, and barren land cover.

In the summer of 2009, sensors were placed at those 23 locations prior to WSN design (Figure 2.2). The locations were selected in the field to reflect variability in catchment-wide physiographic parameters, such as aspect, elevation, and canopy cover. Particular focus was given to previous studies [*Faria et al*., 2000; *Molotch and Bales*, 2005a; *Musselman et al., 2008*; *Rice and Bales*, 2010], which identified these major physiographic parameters as driving explanatory variables for snow depth. Each location was instrumented with snow depth, solar radiation, and relative-humidity sensors (Table 2.1). One-meter-deep soil pits were excavated at each location and the face of each pit was instrumented with soil moisture and matric potential sensors at 10, 30, 60, and 90 cm depths, unless physically prohibited by bedrock. The soil moisture sensors were installed in undisturbed soil, and the matric potential sensors were installed in nearby disturbed soil at the same depth. The soil profiles were then backfilled and hand compacted to maintain the original soil horizons and density as much as possible (see *Bales et al*., 2011 for more detail and Volumetric Water Content calibration procedure). A *Judd* snow depth sensor was also mounted 3 m above ground surface on a 75-cm cantilever beam.

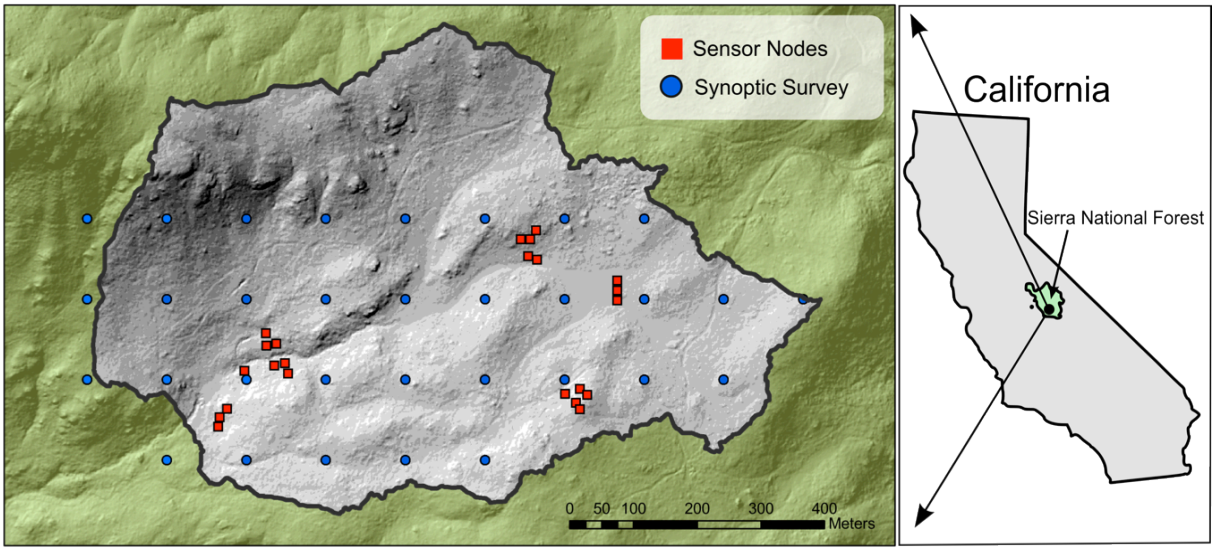

Figure 2.2 Site layout showing location of sensor nodes and synoptic survey sampling points.

Parameter	*Sensor*	*Manufacturer*	*Accuracy*
Snow depth	Ultrasonic	Judd Communications	±1 *cm*
Volumetric water content	EC-TM	Decagon	±3% VWC
Matric potential	MPS-1	Decagon	Calibration dependent (max ±40%)
Solar radiation	LI-200	LI-COR	±5%
Humidity and Temperature	SHT15	Sensirion	±2%, ±0.5˚*C*

Table 2.1 Sensor characteristics.

Along with two additional temperature measurements, this provided over 15 sensors readings at each measurement node (Figure 2.3), giving a total of over 300 sensors for the entire network. Our deployment makes a distinction between *sensor nodes*, which are motes interfaced with the sensors and data-logging infrastructure (Figure 2.3), and *repeater nodes*, which only contain motes and a power source, placed to ensure mesh redundancy, and to transfer data between locations which would otherwise be out of range.

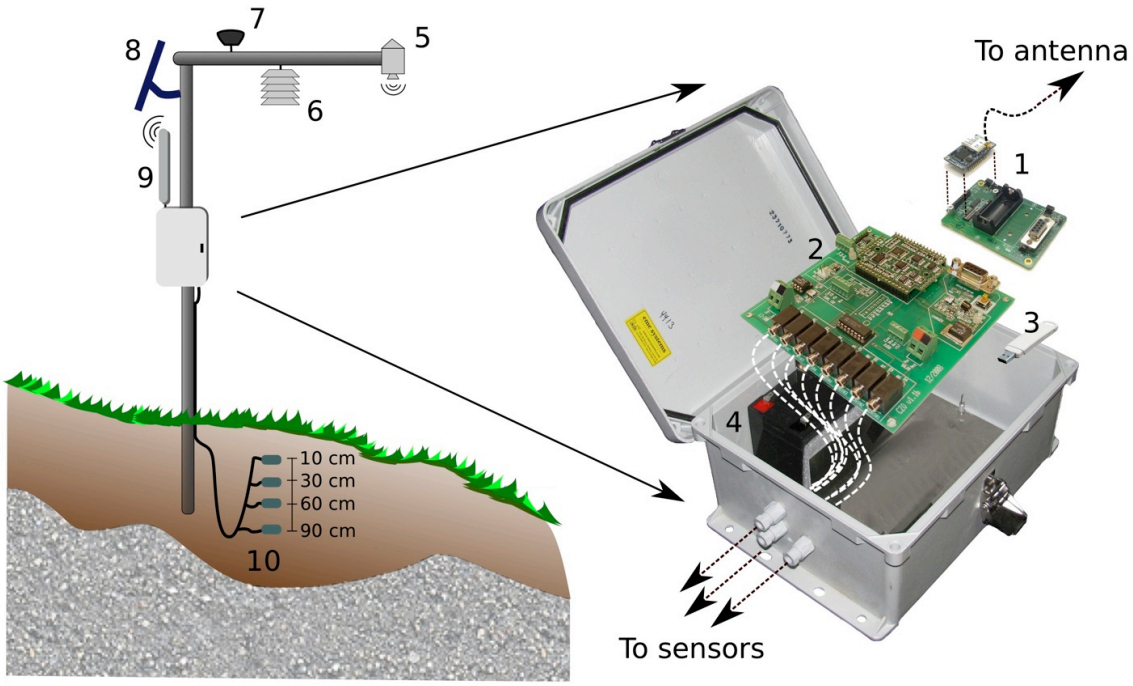

Figure 2.3 Sensor node architecture. (1) Mote, (2) custom data-logger to interface the sensor array, (3) on-site memory storage, (4) 12V battery, (5) snow depth sensor, (6) humidity and temperature sensor, (7) solar radiation sensor, (8) 10W solar panel, (9) external 8dBi antenna, (10) four soil moisture, temperature, and matric potential sensors at varying depths.

The WSN deployment was carried out in two steps. A smaller-scale prototype deployment was conducted in September 2009, relying on manufacturer-specified transmission distances as guide to mote placement (Figure 2.4). This network was mainly composed of sensor nodes (red points in Figure 2.2) and a few repeater nodes. Site-specific network statistics were collected to evaluate the performance of the WSN prototype. An analysis of these statistics, using a set of WSN metrics, was then used to inform a network-wide redesign. Further *repeater nodes* were then added to the network, and existing *repeater nodes* were re-located to ensure desirable network performance, establishing the final network configuration of 57 WSN nodes.

2.3.2 Synoptic surveys and variability analysis

To evaluate the ability of the WSN to capture catchment-scale hydrologic variability, data from gridded synoptic surveys were used to provide ground truth for the distribution of soil moisture and snow depth. The surveys were carried out in the larger CZO, but this analysis will only focus on the measurements specific to the sub-catchment instrumented by the WSN. The survey grid points were selected to evenly cover 125-m intervals (Figure 2.2). The points were located in the field using a *Magellan* handheld GPS unit with 3-5 m accuracy. A snow depth survey was carried out on April 7-9, 2010. At each point, three depth measurements were taken 5 m apart in a N-S transect using a snow depth probe. When possible, each of the three measurements was taken under varying canopy cover (under canopy, canopy drip edge, and open terrain). Two soil moisture surveys were also conducted on June 14-17, and September 6-10 2010 using a hand-held *Hydrosense* soil moisture system. Five synoptic measurements were

taken within a 30x30 cm area and average values were used. Synoptic measurements were integrated over the top 20 cm of soil. These readings were compared to the average of the 10 cm and 30 cm readings taken by the *Decagon* soil moisture probes in the network. Prior to the survey, the *Hydrosense* system was calibrated to soil obtained from the site, and compared to the Decagon soil moisture probes. It was determined that no significant difference in the measurement methods was apparent. In few cases, field conditions and accessibility did not permit some survey points to be sampled.

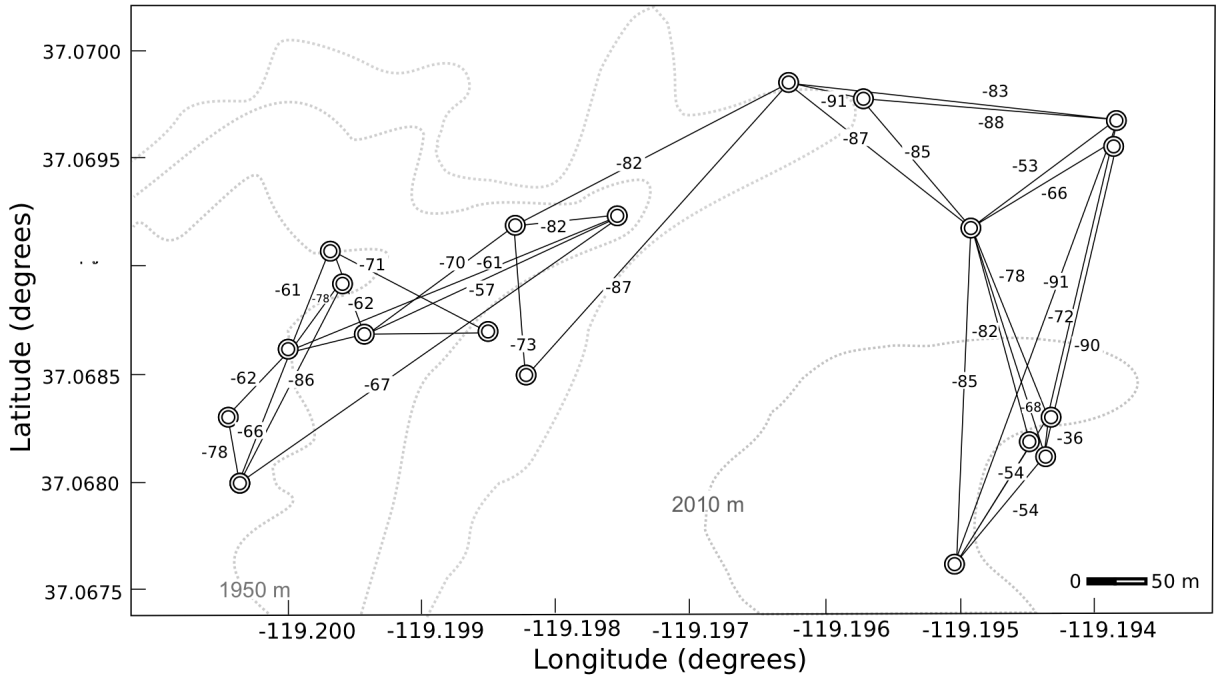

Figure 2.4 Snapshot of the prototype network. Circles indicate locations of motes. The figure lists path-specific RSSI values in dBm (decibels referenced to 1 mW).

One-m^2 resolution LIDAR data sets were also collected both for snow-on and snow-off conditions [*Guo et al., 2010*]. The snow-on LIDAR data set was collected on March 22-23, 2010, and the total snow cover was calculated by subtracting the snow-off raster from the snow-on raster in the ArcGIS software package. To assess the ability of the LIDAR data set to capture catchment-wide snowdepth distribution, and to validate the use of the LIDAR data for purposes of this study, a comprehensive ground-truth survey was conducted on the day of the snow-on LIDAR flight. Manual snowdepth readings, spaced 10m apart in a cross-shaped transect, were taken at forty locations in a central part of the catchment and tagged via hand-help GPS units during the snow-on LIDAR survey. Along with the snowdepth readings taken at the sensor nodes, these points revealed an average error of less than 10 cm when comparing LIDAR-derived snowdepth data to ground-truth readings. Furthermore, the snow-off LIDAR raster was used to delineate the catchment instrumented by the WSN using the *Spatial Analyst* toolbox in *ArcGIS*. The resulting delineation is shown in Figure 2.2 as a shaded relief map.

An analysis was also carried out to investigate the temporal trends in snow depth and soil moisture variability. The results were compared to previous studies, which evaluated forest canopy effects on snow depth variability [*Jost et al.* 2007], the effects of soil depth on volumetric

water content (VWC) variability (as measured by the Coefficient of Variation - CV) [*Famiglietti et al.,* 1998; Grant et al., 2004], as well as the dependence of VWC variability on mean VWC [*Bogena et al.,* 2010; *Famiglietti et al.,* 1998; *Grant et al.,* 2004; *Vereecken et al., 2007*; *Wackerly et al, 2008*]. A geostatistical analysis was also performed on the snow depth and soil moisture data sets [*Bogena et al.,* 2010; *Western et al.,* 2002; Deems et al., 2006] to evaluate effects of spatial scaling on sampling design.

The distribution of aspect and slope captured by the network and synoptic surveys reflected those of the more comprehensive LIDAR survey (Figure 2.5), but the network did not capture some of the higher elevations (1990-2040m) located in the northern portion of the catchment (Figure 2.2). These locations were excluded from the sensor placement in favor of sampling nearer the stream. The sensor nodes were also positioned under various canopy covers, with 30% of nodes under canopy (UC), 30% in open terrain, and 40% at tree drip edge (DE).

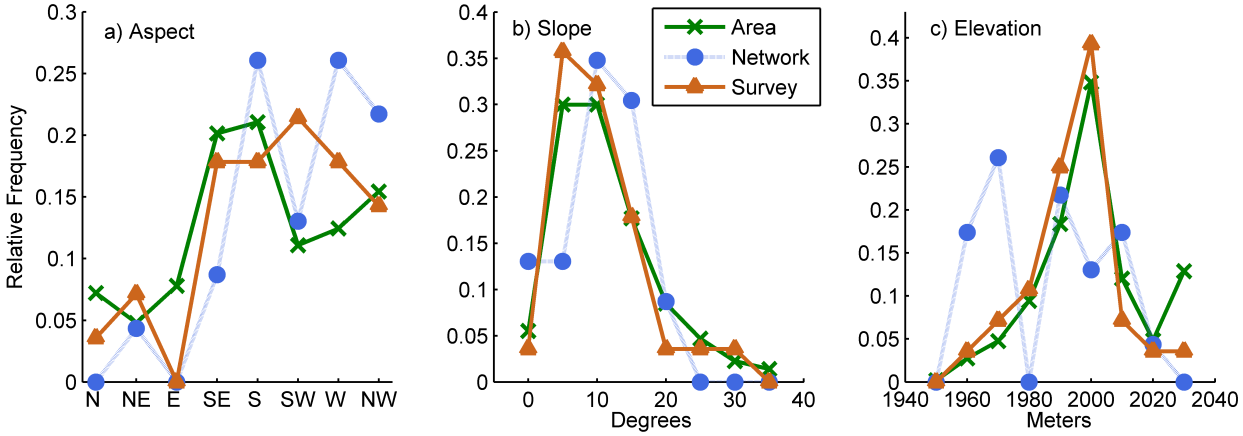

Figure 2.5 Comparison of physiographic parameters captured by LIDAR, against those covered by the sensor-nodes, and synoptic survey.

2.3.3 Hardware

The fluctuating humidity, heavy storms, and cold winter temperatures at the site called for special considerations during the design of the monitoring hardware. Reliability requirements demanded specific attention to the underlying network-control algorithms, as well as the selection of the appropriate transmission frequencies. The system uses wireless devices developed by Dust Networks [*Dust,* 2011], a company that primarily develops wireless data transfer technologies for industrial automation applications. The hardware, in our case the *M2135* mote, is rigorously tested at extreme environmental exposures to meet stringent industrial guidelines, motivating its use in the harsh conditions of the Sierra Nevada. The core software backbone of the Dust WSN relies on the *Time Synchronized Mesh Protocol* (*TSMP*) architecture [*Dust* 2006; Pister and Doherty 2008], which closely follows a number of well-specified industrial wireless protocols [*ISA* 2009; *Song et al.* 2008]. An open source version of TSMP is also available for implementation on most hardware platforms [*Watteyne et al.,* 2012].

These industrial protocols are built upon three main principles: *time synchronization*, *channel hopping*, and *automated multi-hop smart mesh formation*. Motes within a *TSMP* network are synchronized to within sub-millisecond accuracy. This enables accurate time stamping of sensor data, and permits motes to keep their radios off more than 99% of the time by

only communicating according to a dynamically determined scheduled. While such tight time synchronization facilitates an extremely low radio duty cycle, and thus significantly lowers battery usage, it also permits for channel hopping. Along with scheduling transmissions, motes in a *TSMP* network account for fluctuations in the radio space by scheduling a different frequency channel for each transmission. This approach effectively utilizes all 16 available channels defined by *IEEE802.15.4* [*IEEE* 2009]. This mitigates the previously mentioned effects of multipath fading and external radio interference by ensuring that no two mote pairs communicate on the same channel at the same time. This significantly differs from *ZigBee* based networks [*ZigBee Alliance 2008*], where one frequency channel is often used for the entire network, and limited synchronization can lead to intra-network radio interference due to lack of an explicit transmission schedule. The *TSMP* protocol also features dynamic smart-meshing algorithms, which allow each mote to automatically join the network, and adaptively communicate with multiple neighbors, without the need for specialized router nodes seen in *ZigBee* networks. The adaptive nature of this WSN also changes network *paths* (links, or node-to-node routes) dynamically to reflect possible network interference from outside sources, or changing field conditions. The availability of the *TSMP* features, all of which, to our best knowledge, are not implemented on other commercial platforms, motivated the decision to use Dust Networks hardware for this deployment. Each mote is powered by a 3V battery and has a manufacturer-specified battery life of over two years.

The WSN featured one base station, or *network manager*, that acted as a central network controller. The network manager communicated with a low-power, embedded Linux computer via Ethernet and the two devices exchanged commands over an extended mark-up language (XML) interface to actuate motes within the network, log sensor readings, and gather network statistics. The embedded computer was located at the base of a 50-m tower used for eddy correlation. A cellular modem positioned 25m up the tower provided Internet connection for the transfer of real time data to an off-site location. Power to the network manager, embedded Linux computer, and cellular modem was provided by a bank of 12V batteries and a 120 Watt solar panel. Careful selection of low-power components enabled battery-powered operation of the base station without the need for outlet power.

Anomaly detection within the network was conducted by custom software that was written to interact continuously with the network manager to monitor network statistics and sensor readings, informing network operators of impending battery outages, sensor failures, or other events of interest. To further improve transmission distance, physically hardened, high-gain 8dBi antennas were mounted 3m above ground surface at each network node.

While commonly used data logging components, such as the *Campbell Scientific CR1000* line [*Campbell Scientific,* 2012], could be interfaced with the wireless hardware, a more cost effective, and more customizable data-logging board was designed to control and power the sensor array, as well as to form the data into serial packets required for transmission by the Dust Networks mote. The board is the EME Systems *OWL2pe* data logger [*EME,* 2011], and interfaces nine analog and twenty digital inputs, allowing for additional sensors to be attached in the future. The board requires about 240 mW of power while actuating sensors, but spends 99% of its time in an efficient low-power mode. The *OWL2pe* can be self contained, or controlled remotely via the WSN, and built in flash memory allows the assembly to log data locally in case the WSN becomes unresponsive. Lack of accessibility to the physical network, especially during heavy snow periods, as well as the possibility of long-lasting diminution of solar radiation, played a significant role in the design of the power-management infrastructure. Monitoring

requirements called for data to be collected on ten to fifteen minute intervals, and conservative preliminary calculations indicated that a 12 Volt, 7 Amp-hour lead-acid battery, and a 10 W solar panel, could reliably power the data logger and sensor assembly. To mitigate the effects of extremely cold temperatures, the repeater nodes were equipped with 3.6V lithium thionyl chloride batteries. The power consumption of network components is shown in Table 2.2.

Component	*Voltage*	*Current – A*	*Duty Cycle (%)*	*Avg. Current*
Mote	3	0.02 (<$2x10^{-5}$)	<1	$6x10^{-6}$
Data logger	12	0.02 ($2x10^{-4}$)	1	$4x10^{-4}$
Sensors	5-12	0.1	1	$1x10^{-3}$
Network manager	12	0.06	100	0.06
Cellular modem	12	0.4 (<0.1)	50	0.2
Embedded	5	0.3	100	0.3

Table 2.2 Power consumption of various components in the network assuming a 15-minute sampling interval.

2.3.4 System Cost

At the time of deployment the cost of each mote was under $50 USD, with an additional $500 for a sensor-node, including custom data logger, housing, mounting, and external antenna. Base-station components, including the network manager, *Linux* computer, cellular modem, enclosure, batteries and solar panel totaled about $3000. Sensors used in the study comprised the biggest cost, varying in price from $25 to $600 each, the most expensive of which was the snow depth sensor.

2.3.5 WSN Metrics

The packet delivery ratio (PDR), and received signal strength indicator (RSSI) are two metrics that provide particular insight when evaluating WSN performance [*Al Basset Almamou* et al., 2009; *Karl and Willig,* 2005]. PDR is a metric that captures the overall communication efficiency of a link between two WSN nodes, and is defined as the number of successfully transmitted packets divided by the total number of transmitted packets. It can be thought of as the probability that a transmission between two nodes will succeed. Needless retransmission of a packet can take a significant toll on battery resources. Conservation of network resources ideally demands a PDR of 100% on all links used by the network. An understanding of the interplay between spatial network coverage and deployment-specific PDR characteristics is thus required. Since a primary motivation of using WSN hardware for environmental monitoring is the ability to span large areas, optimization of a field deployment demands a balance between the need to maximize battery life (power cost) while minimizing the number of required motes (capital and operational costs).

RSSI, conventionally measured in decibels and referenced to 1 milliwatt (dBm), represents the power in a signal when it arrives at the receiving antenna. As a general rule, the RSSI decreases as the distance between two nodes increases (inversely to the square of the distance between two nodes). Given the logarithmic scale, every 3 dBm decrease reduces the

received power by approximately one half, and more negative dBm values correspond with lower signal strength. The true behavior of RSSI over node-to-node distance depends upon a number of factors, such as multipath radio propagation and environmental obstructions but can be roughly captured by an idealized Friis antenna propagation model [*Friis,* 1946; *Friis,* 1971; *Kraus,* 1988]. Since this theoretical propagation model is valid only in *free space*, and does not take into account the effect of multipath radio propagation, the model cannot, in most cases, be used to accurately predict the behavior of most real world phenomena. Generally, WSN hardware has a manufacturer-predefined RSSI threshold (sometimes also referred to as *receiver sensitivity*) below which signals cannot be reliably decoded by the hardware. Changes in the physical environment affect the received signal strength, thus causing RSSI to fluctuate over time. The operational RSSI threshold (one that gives reliable performance in the field) can often exceed, or under-perform, manufacturer specifications, thus placing the onus of classifying link reliability onto the WSN operator. Since RSSI threshold is often used to specify reliable transmission distances, optimizing network performance becomes the job of the network operator (researcher) and depends on understanding the link-specific RSSI behavior for any given deployment. Improper determination of these thresholds can lead to network collapse and the need for reconfiguration during the actual deployment.

Motes have the ability to enter any number of possible sleep states [*Karl and Willig,* 2005; *IEEE,* 2009], where significant portions of the mote are shut off to conserve energy. For example, when a mote is not transmitting or receiving data, it turns off its radio while placing the micro-processor into a low-power state. This significantly reduces power consumption on the battery. A particular problem resulting from an extended lack of communication occurs when motes enter a deep-sleep, or hibernation, state. In such cases, motes that have not been able to establish communications with the network manager through their neighboring motes for a period of time, power off the majority of resources, and wake up only sporadically to rejoin the network by listening to advertising beacons from neighbors. This feature, although not standard on all platforms, was designed to conserve power in industrial mesh networks. When implemented, field experience has shown that it can take significant portions of time (hours to days) for motes to rejoin the network upon entering the deep sleep state. In relatively low node-density networks, such as those used for environmental monitoring applications, a mote in a deep sleep state is effectively removed from the network, which has the potential to severely compromise the mesh topology, and can lead to outages of entire subsets of the network. This behavior further underscores the need to understand path specific PDR and RSSI characteristics for ensuring network robustness, and avoiding network outages due to drops in connectivity.

2.4 Results

2.4.1 Prototype deployment

During on-site evaluation of the motes it became apparent that the initial prototype network layout (Figure 2.4) did not perform up to manufacturer-provided specifications. The specifications indicated an RSSI threshold of -89dBm, with a prescribed mote-to-mote spacing of up to 200-300 meters. In practice, few reliable paths at such distances were observed, and RSSI values between -80 and -90 dBm were experienced for significantly shorter network links. Network statistics were collected over a 25-day period, beginning with 10 motes on September 15, 2009 and expanding to 19 motes by September 18 (Figure 2.4). RSSI and PDR value were averaged over 15-minute intervals for each path in the network, providing a range of values to

determine proper hardware thresholds and evaluate network performance. Over this 25-day period, the number of paths (or links between nodes) varied as the WSN automatically adjusted the network topology to account for fluctuations in communication reliability (Figure 2.6). Generally, a larger number of paths are desirable since it permits for greater diversity of routes on which data packets can be transmitted. The spike around September 18 corresponds to the deployment of the full 19-sensor prototype network. As more motes joined the network, the number of paths increased correspondingly. The average network RSSI dropped as a result, implying that the newly formed network paths possessed low connectivity. This correlation was observed for the remainder of the 25-day evaluation of the network prototype. Ideal performance would have kept average RSSI steady, regardless of the number of newly created network paths, indicating that all paths reflect similar communication reliability. Further investigation into possible causes of network behavior showed that both the average network-wide RSSI and the average relative humidity experienced fluctuations over time, but did not appear to be positively correlated (Figure 2.6).

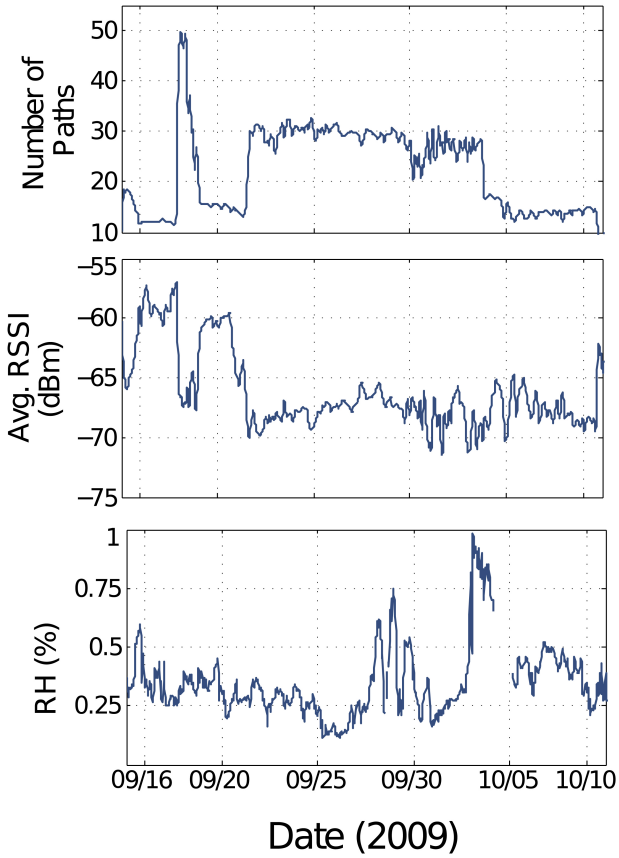

Figure 2.6 Network behavior over a 25-day period, starting on September 15, 2009. Network resets are evident as the sudden spikes in number of network paths. The average RSSI of nodes in the network dropped as more paths were created (more nodes joined the network). A rainstorm, evident around October 3, coincided with a significant drop of paths from the network.

The western portion of the prototype network experienced a complete outage on September 19, leaving only a few network paths operational (Figures 1.4, 1.6); no prior indication of an impending network outage was evident. During the outage, the total number of

paths in the network dropped significantly and abruptly. This can be attributed to a bottleneck effect within the network. In a meshed topology, *child nodes* far from the network manager need to transmit data through neighboring *parent nodes* (nodes closer to the network manager), which in turn pass the data further along the mesh until it reaches the network manager. On September 19 critical parent nodes in the mesh drifted into a hibernation state due to a drop in connectivity (most likely due to low RSSI values exhibited by network paths in the center of Figure 2.4), and child nodes that depended on these connections lost connectivity to the network manager. Child nodes in the western portion of the network were thus forced to enter a hibernation state as well. To avoid this bottleneck phenomenon, redundancy and robustness needed to be built into the network to ensure that no single outages can cause the remainder of the network to fail. This bottle-necking behavior further explains why average network RSSI was inversely proportional to the number of paths. Following this outage, only stable links with high RSSI values remained in the network.

The mote hibernation state was designed to conserve power during network communication outages, since transmission during such instances of low connectivity can significantly shorten battery life. Once in the hibernation state, it could take days for the mote to rejoin the network depending on field conditions, and a physical reset was required to bring motes out of hibernation, and to re-establish the entire WSN mesh. Once the network hardware was reset a week into the prototype deployment, the number of network paths increased and stayed constant for the following two-week period.

Relative humidity did not seem to affect the overall behavior of the network (Figure 2.6), but a rainstorm around October 3, 2009 coincided with another network collapse and a significant drop of nodes from the network. To mitigate future network collapse, the *TSMP* protocol settings could be modified to lower the threshold for entering the hibernation state, or additional repeater motes could be added to ensure that required RSSI thresholds would be met. It was deemed that the low-power sleep mode was necessary for efficient network operations, and that a redesigned network layout would offer a better alternative. This was again confirmed when noticing a significant drain on mote battery life after forced resets.

Part of the pre-deployment site analysis involved determining the field RSSI as a function of node-to-node distance. Figure 2.7 presents a logarithmic plot of RSSI vs. path distance for the prototype deployment, and for comparison, the idealized Friis propagation model [*Friis,* 1946], which governs how the hardware-specific RSSI between two nodes should change assuming perfectly isotropic antenna and *free space* conditions. The measured relationship was not monotonic, nor sharply delineated, and showed considerable uncertainty, which was most likely a function of multipath effects on radio propagation. The *in situ* network statistics showed a significant reduction in transmission distance compared to the manufacturer-specified 200-300 m outdoor range. As expected, the observed average RSSI decreased as the distance between nodes increased, but a significant spread existed around the mean (shown by an estimate of the 95% confidence bounds). The mean RSSI reached the manufacturer specified threshold of -89 dBm at a transmission distance of 150 m, but a portion of paths performed below this threshold at the same distance. At 115-120 m separation, a portion of network paths had the tendency to drop below this threshold, under-performing manufacturer specifications of a 200 m range. For distances below 25 m, the network links outperformed the predictions made by the Friis propagation model but mean RSSI significantly deviated from the idealized Friis model after the 25-m mark.

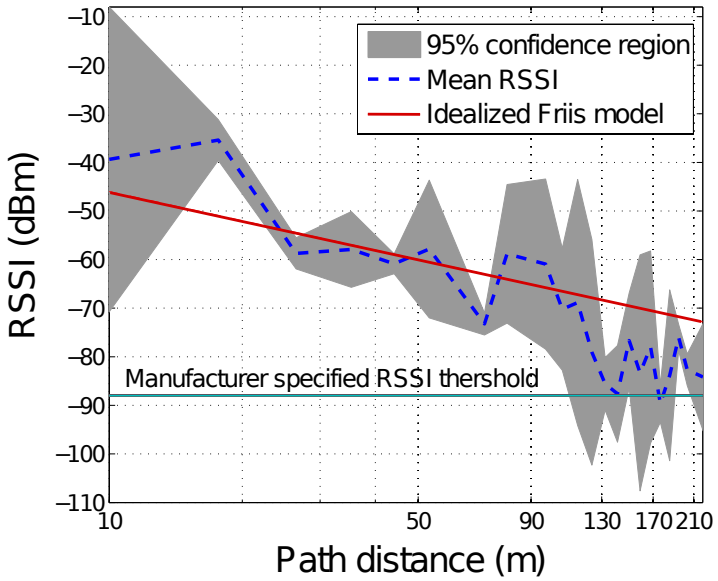

Figure 2.7 Plot of RSSI vs. path distance collected over a 25 day period. The RSSI decreased with an increase in path distance, and significantly fluctuated around the mean. A number of *in situ* radio links experienced RSSI values below the manufacturer-specified -89 dBm threshold at approximately 100-120 m.

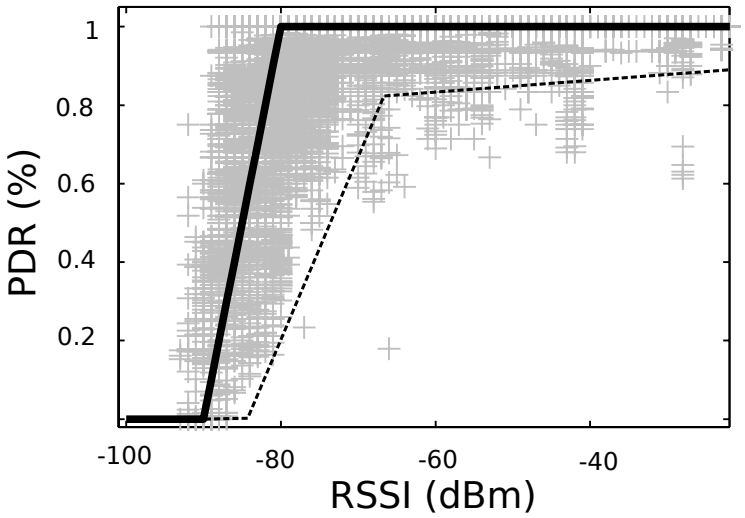

Figure 2.8 A plot of PDR as a function of RSSI. Gray points reflect observed behavior, while the solid black line represents an idealized *waterfall* fit. Most packets are transmitted successfully for an RSSI above -80 dBm, after which there is a sudden drop in PDR. The dashed line provides a conservative, and lower bound design curve.

Figure 2.8 presents another, and perhaps better, indicator of network performance - an analysis of PDR (%) as a function of RSSI (dBm). Although fluctuations of PDR around specific RSSI values existed, it is evident that for values of RSSI over -80 dBm, the PDR remained around 100%, implying high reliability and that few, or no retransmissions were required to successfully deliver network packets to their destination. For values below this RSSI

threshold, paths experienced sudden drops of PDR, falling entirely to 0% for RSSI values of -90dBm and lower. This sudden drop is often referred to as *waterfall* behavior when describing to the effect of RSSI on PDR.

2.4.2 Post-deployment analysis

Following the prototype deployment, an analysis of network metrics was used to inform a re-design. More repeater nodes were added to the network, and existing repeater nodes were re-located to improve performance and prevent network outages, leading to the final network configuration shown in Figure 2.9. After re-design, the network did not experience any more outages. A histogram of 30 days of PDR values for all paths in the network (Figure 2.10) a month after network reconfiguration, shows that about 80% of total PDR values were within a desirable 85-90% performance range, with more than half of all paths experiencing 100% PDR (particularly nodes spaced less than 50 m apart). Nodes communicated not only with their closest neighbors, but also with more distant nodes to maintain the mesh network topology. In the final deployment, network nodes had, at all times, a stable connection to at least two neighbors located 50 m or closer, and the average network-wide PDR remained within an 85-95% window for the remainder of the network lifetime.

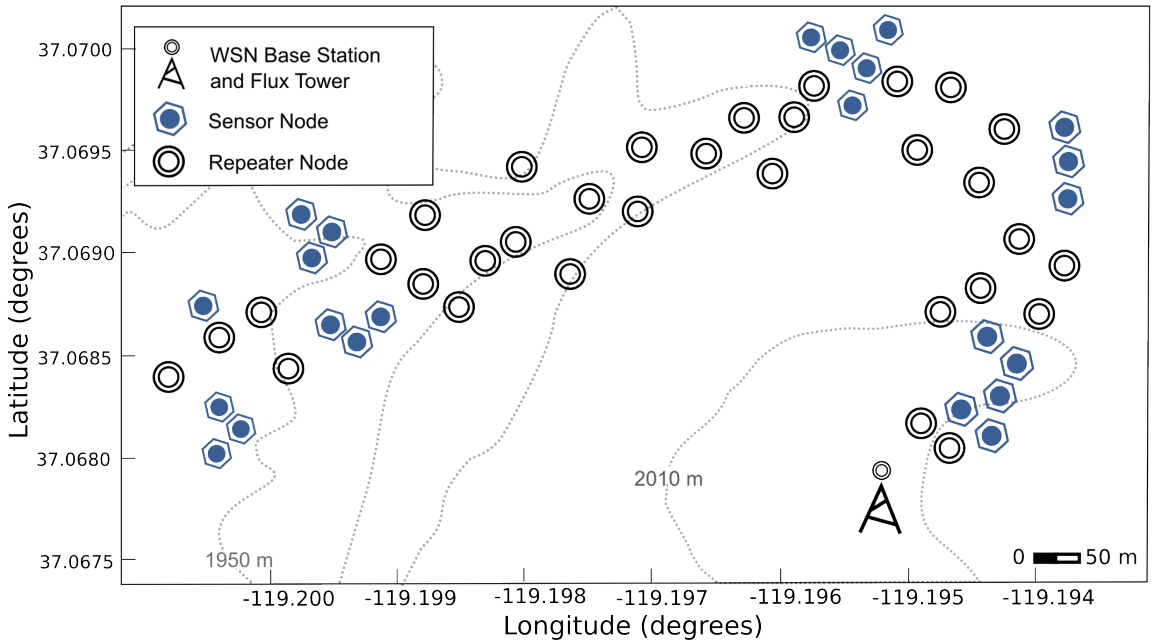

Figure 2.9 Final wireless sensor network layout. The base station, located at an eddy flux tower, houses a network manager and acts as a central data aggregation point.

To illustrate more some short-term dynamics of the network, Figure 2.11 shows a sample of PDR evolution for three randomly selected paths in the final network configuration. Though average network PDR remained high, path *A* exhibited a sudden drop at the four-hour mark, only to recover within one hour. Path *B* exhibited fluctuations of PDR for the entire sample period.

While requiring multiple retransmissions of the packets, this path did not drop from the network and its PDR never dropped below 80%. The figure also shows a new network path *C* being established during this 24-hour window, which is a positive indication that the self-healing network algorithms dynamically allocated paths to take advantage of more stable connections, thus alleviating otherwise relatively less stable connections and creating further redundancy. The formation of new paths between nodes was caused by fluctuation in the radio space, and is further motivation for using an adaptive, self-healing network.

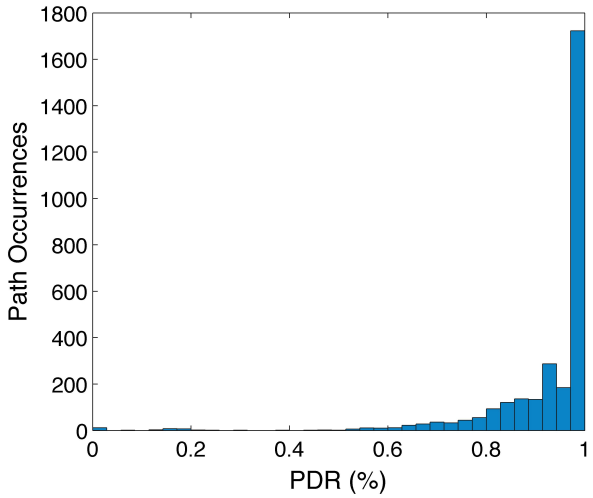

Figure 2.10 A histogram of PDR values for all paths in the network following a reconfiguration of the network. About 80% of all network paths are performing within the desired 85-90% design value, and over 50% of all paths are at 100% PDR.

During network operation, there were some instances of paths longer than 200 m experiencing RSSI values greater than -60dBm, while some paths less than 50 m operated with RSSI values of -80dBm. This behavior was not consistent across the network, and could not be directly attributed to physiographic features of terrain (such as elevation differences, or line of sight), thus further underscoring the unpredictability of the wireless environment.

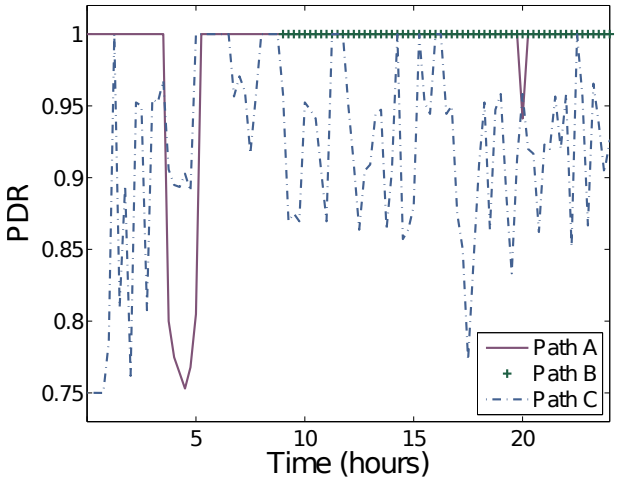

Figure 2.11 The fluctuations of PDR over time for three specific paths in the network over a 24-hour period.

2.4.3 Network measurement validation

The mean snow depth measured by the network on March 22-23, 2010, was 139cm, versus a 145-cm mean from the LIDAR data (Figure 2.12a). On April 7, 2010, the network measured a mean snow depth of 142cm, versus 140 cm for the synoptic survey. Although the network and LIDAR data showed a similar cumulative distribution for March 22, the network measurements of snow depths varied between 70-242cm, while the LIDAR snow depths varied between 0-300cm. The cumulative distributions of snow depth for the March LIDAR, March Network and April Networks shared a similar, near-normal shape, while the cumulative distribution of the April 7 survey was biased towards relatively high depth values.

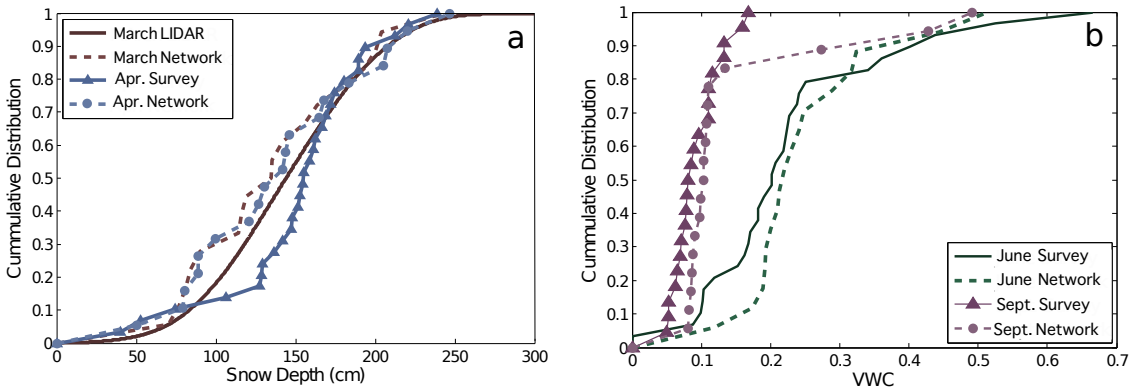

Figure 2.12 Cumulative distribution functions comparing network readings of a) snow depth, and b) soli moisture to those obtained by synoptic and LIDAR surveys.

For the June soil moisture study, the network VWC measured a mean of 0.23 (Stdev. 0.14), and synoptic survey measured 0.25 (Stdev. 0.09). The cumulative distributions in Figure 2.12b show the close resemblance between VWC values recorded during the June survey and the June Network measurements. For September 2010, the network reported a mean VWC of 0.15 (Stdev. 0.12), compared to the 0.09 (Stdev. 0.03) in the synoptic survey (Figure 2.12b). Three network readings, in particular, showed relatively high VWC during the September survey. The respective cumulative distribution of VWC showed a close correspondence between the September survey and network reading up until the inclusion of those relatively higher VWC values captured by the network. These three specific sensor nodes were located in flat-aspect terrain, in open meadows.

A histogram of snow depth distribution, based on the March 22 LIDAR point-cloud, shows that snow depth was near-normally distributed (Figure 2.13). This distribution was obtained after removing a small set of outliers (< 0.1%); i.e. negative values, and extremely large positive values (depths > 3.5m). Given this near-normal distribution of snow depth, as well previously observed tendencies of spatial VWC to follow a normal distribution [*Grant et al., 2004; Famiglietti et al., 1998; Western et al., 2002*], a two-sample *t-test,* and two-sample *f-test* [*Wackerly et al, 2008*], were carried out at a 5% significance level ($\alpha=0.05$) to assess the equality of the means and standard deviations between the surveys and network readings. With the exception of the September soil moisture study, these test confirmed null hypothesis H_0 (namely, the equality of the mean and standard deviation between survey and network data). The difference between the September soil moisture survey and network data was deemed statistically significant, with *attained significance values* (*p-values*) of $p=0.046$ and $p<0.0001$ for

the mean, and standard deviations tests, respectively. A non-parametric *Wilcoxon rank sum test* [*Wackerly et al, 2008*] was also conducted, yielding *p=0.0625* for the comparison of the September data sets.

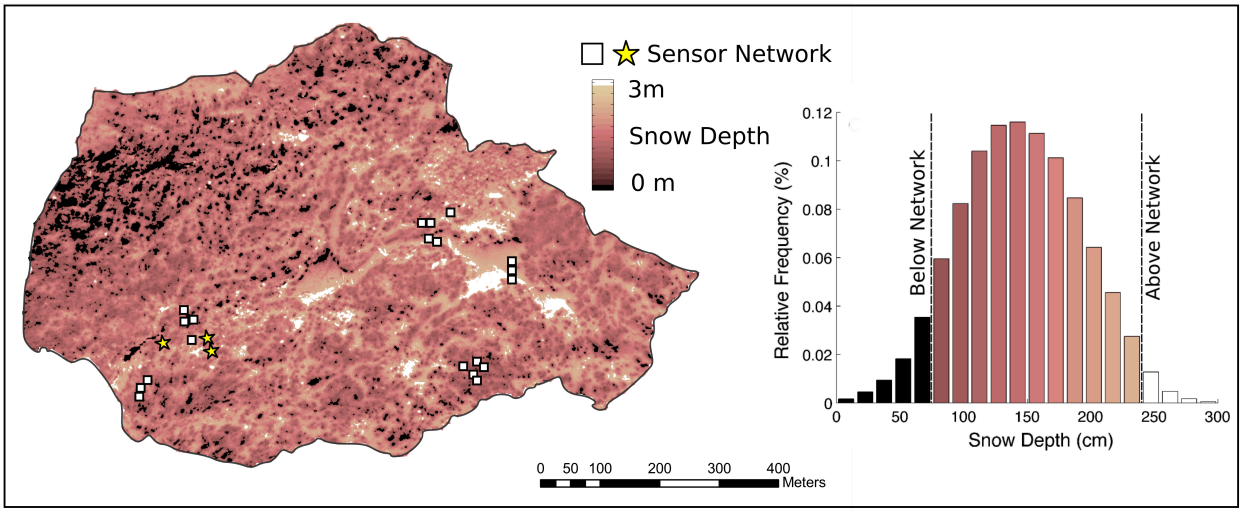

Figure 2.13 Total snow cover calculated from a snow-off and snow-on LIDAR data sets at $1m^2$ resolution (March 22-23, 2010). White squares indicate location of sensor nodes. The yellow stars indicate locations of the three sensor-nodes that persistently recorded snow depth values closest to the mean of both the LIDAR-, and synoptic-survey. Histogram of snow depth values calculated from LIDAR data is also shown. Partitioned regions indicate low, and high values not captured within the range the WSN measurements.

Figure 2.13 also draws a distinction between those portions of the LIDAR data that fell below and above the respective minimum and maximum measurements captured by the network. These extreme readings, although not captured fully by the sensors nodes, did not have a significant impact on the overall mean. The majority of locations with snow depth values below those captured by the network were located in the northwestern portion of the network (Figure 2.13). Larger patches with snow depth values above those captured by the network were located in the centrally situated, open meadow (Figure 2.13). The distribution of physiographic attributes of un-instrumented locations, for which the set of depth values fell below those captured by the network, did not show signification deviations from those observed on the overall site (data not shown). In the case of both the LIDAR and synoptic surveys, three particular sensor-node locations (13, yellow stars), situated in a north-facing aspect, and across varying canopy cover, reflected measurements very close to the catchment mean ($\pm$ 9 cm).

2.4.4 Snow depth variability

The majority of snow depth readings for water year 2010 (WY2010, starting in October 2009) reached an average peak accumulation of 160cm in April of 2010, although a subset of nodes experienced this same average peak in the middle of January (Figure 2.14a). Snowmelt timing showed notable variability, with a three-week span between first and final melt-out dates at sensor-node locations. Nodes located on flat facing aspects generally saw faster melt rates

than those on north- and south-facing aspects, while average accumulation showed little variability across different aspects.

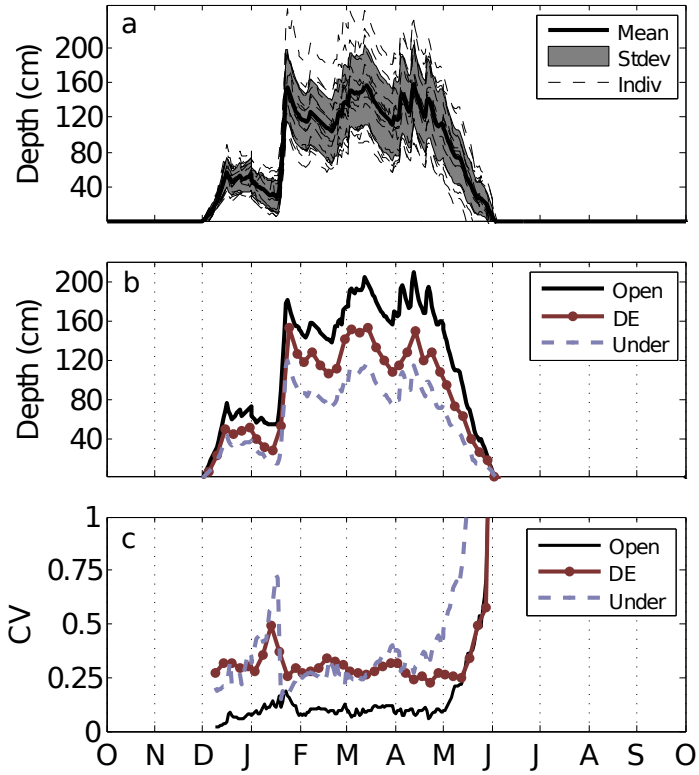

Figure 2.14 Temporal snowdepth trends October 2009 to October 2010. a) Average snow depth (one standard deviation shown in in gray, individual sensors indicated by dashed lines), b) average snow depth for various canopy covers (open, drip edge, and under canopy), c) Coefficient of variation over time for snow depth and various canopy covers.

Forest cover was the major explanatory variable of snow depth variability across the catchment. Sensor nodes located under canopy, measured 40cm less than those located under drip-edge, and 80cm less for open-cover (Figure 2.14b). Under-canopy snow depth showed greater variability, as measured by CV, compared to nodes located in open terrain (Figure 2.14c). This trend in variability exhibited temporal stability, with under-canopy nodes persistently reflecting relatively higher variability throughout the entire season. The group of nodes located along tree drip-edge reflected almost identical variability patterns compared to under-canopy nodes during the snow accumulations seasons. During snowmelt, however, snow depth variability at drip-edge nodes began to diverge from that of under-canopy nodes, until matching the variability exhibited by nodes located in open terrain.

2.4.5 Soil moisture variability

At most sensor nodes, peak VWC levels of 0.31 to 0.32 were reached at the beginning of June 2010, coinciding with the melt-out of the snowpack (Figure 2.15a). A subset of nodes experienced peak VWC in the middle of an October 2009 rainfall event. The remainder of VWC

fluctuations corresponded with the snowmelt cycle. On average, sensor nodes located on flat aspects experienced VWC 5%-10% greater than those situated on north and south facing aspects. The major contribution to this higher VWC was from sensor nodes placed in the flat meadow regions, where catchment melt water accumulated and some locations experienced VWC of up to 0.60. On average, nodes placed under canopy and in open areas reflected similar VWC behavior, while nodes placed along the drip-edge exhibited relative larger VWC (by 0.05).

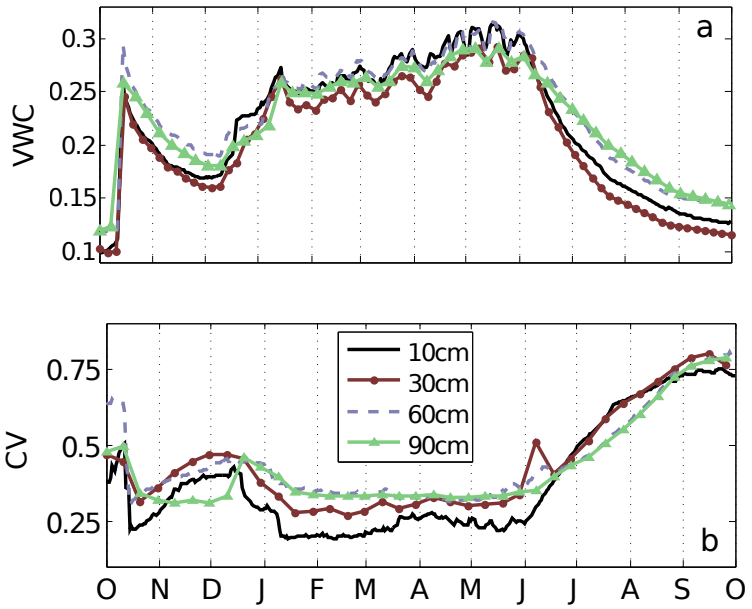

Figure 2.15 Temporal VWC behavior October 2009 to October 2010. a) Average soil moisture across various solid depths, b) VWC coefficient of variation at various soil depths.

In general, average variability (CV) of soil moisture increased with decreasing mean VWC following snowmelt (Figure 2.15b). Shallower soils showed greater VWC fluctuations during the snowmelt season, with VWC fluctuations attenuating with increased soil depth. At a 10cm depth the relationship between mean soil moisture and the coefficient of variation showed the convex-down pattern (peak at 20%VWC) seen in other studies *[Bogena et al.,* 2010; *Famiglietti et al.,* 1998; *Western et al.*, 2002; *Vereecken et al., 2007*] (data not shown). The same analysis showed that no clear relationship between mean VWC and CV was apparent for deeper soils.

During snow accumulation and melt, soil moisture variability increased with soil depth, with soils at 10cm exhibiting lower variability relative to deeper soils (Figure 2.15b). Following snowmelt, soil moisture variability of the shallow soils exhibited a more pronounced increase in variability. As average soil moisture decreased towards the end of the water year, deeper soils once again exhibited greater variability.

2.4.6 Geo-statistical Analysis

The snow depth values extracted from the March LIDAR data showed a very well defined experimental variogram, with a range of 30m and a sill of γ^2=0.2m^2 (data not shown). This snow depth variogram was very similar to that seen in *Deems et al. (2006)*. The variogram

for the April 7 snow depth survey showed no pronounced support, with a constant sill of $\gamma^2=0.2m^2$. Variograms were also generated over time for all snow depth readings recorded by the network. While the sill of the sensor node variogram changed over time to reflect depth fluctuations during the season, a support of 30-40m was evident throughout the analysis. The sill of the variograms generated by the network data also matched that of the LIDAR and synoptic surveys during the respective analysis periods. The same variogram analysis was carried out for soil moisture surveys and networks readings. Even after de-trending for various physiographic parameters, the analysis did not find a clear spatial relationship between lag and VWC variability, such as those seen in previous studies [e.g. *Bogena et al.* 2010; *Famiglietti et al.* 2008].

2.5 Discussion

2.5.1 Wireless sensor network

We define an optimal network deployment as the one that maximizes PDR, while minimizing RSSI (implying greater transmission distances). On average, the need to retransmit data will then be reduced, while spatial coverage of the network will be maximized. Based on this design criterion, Figure 2.8 shows that this value corresponds approximately to -80dBm (although there begins to be a significant spread of PDR for values below -70dBm). This *waterfall* behavior is a further indicator that a manufacturer-specified RSSI *receiver sensitivity* threshold dBm should not be employed as a proxy for reliability. In the case of this study, it is apparent than no packets can be received beyond the manufacturers specified threshold of -89 dBm, but PDR begins to exhibit non-desirable qualities well before this threshold. As such, the manufacturer specified RSSI threshold thus presents a worst-case performance indicator, rather than serving as a guide for reliable communications. A deployment-specific waterfall plot will however go a long way towards helping to design a reliable WSN.

For our deployment -70 dBm was chosen as a conservative, but reliable, RSSI threshold to account for adverse effects due to unexpected wintertime fluctuations. Given the measured -70 dBm needed for acceptable communications reliability, a design path distance was then extracted from Figure 2.7. For further reliability, the lower 95% confidence bound, rather than the mean RSSI, was used as the design criteria. For the installation, this resulted in the derivation of a conservative 50-m mote-to-mote spacing. Realistically, however, our redesign could have employed a larger node-to-node spacing (100-120m), and still performed well, since we did observe reliability for such path distances in the final network. Given the expectations of a harsh winter season, and the possibility of nodes getting damaged during heavy snowstorms, we chose a 50m node-to-node spacing to alleviate any concerns of potential network outages. While our approach is conservative in this regard, we expect that a more temperate, but equally large catchment could be instrumented efficiently with significantly fewer repeater nodes.

It is evident from Table 2.2 that the mote is the smallest power consumer of the entire sensor node assembly. The choice of snow depth sensor, which consumed a major portion of power at each node, was based on the sensor performance and reliability during a number of previous studies [*Bales et al.,* 2011]. A lower-powered ultrasonic sensor alternative (such as the one in *Varhola et al.,* 2010), could have reduced battery size and solar panel requirements, in which case radio power consumption would begin to play a larger role in overall energy efficiency.

The current availability of power at each sensor node would have made it possible to employ higher powered radios, and thus reduce the number repeater nodes. Given the radio

propagation characteristics described previously however, increasing transmission power significantly affects energy consumption, but does not proportionally increase transmission distance [*Friis,* 1971; *Kraus,* 1988]. It can be shown that in many cases, the use of lower-powered radios and efficient repeater nodes consumes overall less power than a system with higher powered radios but fewer repeaters [*Karl and Willig,* 2005]. This becomes a major motivating factor when considering the need to power the entire wireless network on batteries in remote locations with limited access. In our case, the low cost of repeater nodes, along with their long battery life (potentially five years on two AA batteries) and low maintenance continues to motivate our use of the present WSN system. Such low power use and reliability is a particular benefit of the *TSMP* protocol, with its tight time synchronization and extremely low duty cycle.

We propose a three-step design procedure to optimize a WSN deployment (Figure 2.16). A WSN deployment can be separated into *pre-deployment*, *deployment*, and *post-deployment* phases, which, when carried out properly, will ensure robust and reliable network communications while maximizing battery lifetime and transmission distances (and thus reducing the number of required motes).

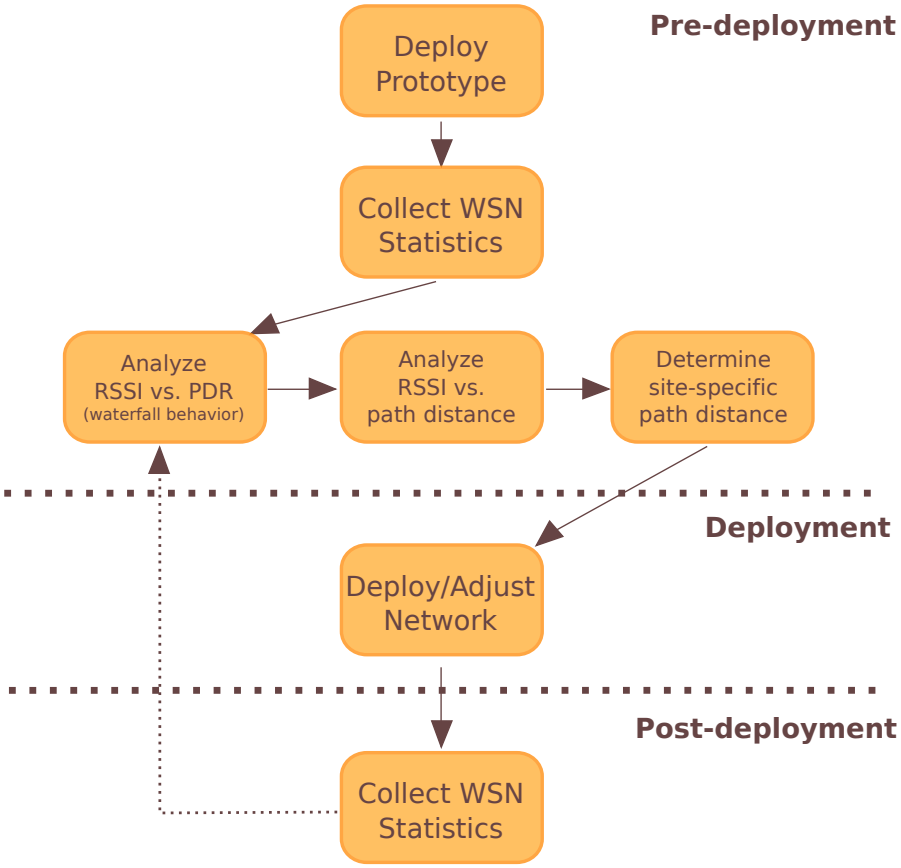

Figure 2.16 The three phases of a WSN deployment.

First, in the *pre-deployment* phase, a prototype network (or a subset of the actual network) should be deployed a priori to gauge the actual performance of the wireless hardware under operating conditions. The pre-deployment analysis should be carried out in environments similar to those of, or preferably at, the actual deployment site, with an attempt to cover a range of terrain parameters. Networks statistics, specifically RSSI and PDR, should be collected for

several days, or longer, to capture possible fluctuations in performance. This will place manufacturer specified performance thresholds into context and will permit site-specific network behavior to be evaluated prior to a full-scale deployment. A *waterfall* plot of the effects of RSSI on PDR (such as the one in Figure 2.8) should be constructed to extract the site-specific RSSI threshold. The primary design objective entails maintaining stable network links (high PDR values) and ensuring robust network performance. This will prevent network outages and collapse due to the inability of motes to communicate. While maximizing network-wide PDR will reduce re-transmissions, and thus also prolong battery life, there is a monetary and labor trade-off associated with placing many repeater nodes. A PDR of 100% is thus not desirable, and may, in fact, be physically unattainable. Calculations should be conducted to evaluate the effect of average PDR on site-specific battery resources, but in most cases a PDR value of 85-90% provides a realistic performance goal. Such an analysis will maximize node-to-node spacing, while reducing drain on battery resources, and will facilitate robust and reliable network links. It is critical to develop a statistically meaningful RSSI threshold from the PDR behavior to account for possible fluctuations in performance. This threshold can then be used to determine transmission distances (in this case from Figure 2.7), and will provide a reliable estimate of the number of nodes that will be required to cover the sensing region.

Second, the *deployment* phase of the final network (placement of sensing nodes, repeater nodes, or star-layout) should be carried out based on the previous analysis of network statistics, keeping in mind that these statistics provide only a bound on performance behavior. Signal attenuation and multipath propagation characteristics may adversely affect specific network links. Brief field checks during the actual placement of nodes should be made to ensure that each network path meets the desired RSSI threshold. Often, simply moving nodes by a few meters will establish a more robust link.

Third, in the *post-deployment* phase, collection of network statistics should be continued as long as necessary to evaluate the *in situ* performance of the network and to capture possible fluctuations in connectivity. Depending on the observed behavior, these statistics should be further analyzed, and adjustments to the network can be made to achieve a desired performance, by adding, removing, or re-locating WSN nodes.

In our case, the data obtained through the analysis of the network prototype was used to gauge the future behavior and provided a set of valuable thresholds upon which to re-design the network. Analysis of such results would have saved a significant amount of time if conducted prior to the deployment of the network. A network design procedure using the above steps will, in most cases, regardless of hardware choice or manufacturer, ensure that WSN performance will meet the demands of an environmental monitoring application, by maximizing path reliability, battery life, and spatial coverage.

While the outlined methods apply to a broad range of available WSN platforms, initial hardware selection should still be strongly guided by built-in network features. The choice of WSN hardware for this deployment was based on the ability of the *TSMP* protocol to facilitate long battery life, as well as the ability to channel hop and to set up and maintain an adaptive self-healing, multi-hop mesh network. Regardless of its major benefits, the complexity of this protocol has prevented many manufacturers of WSN hardware from its implementation. There is a significant cost if the direct user is required to design and implement low-level control of the WSN hardware, as well as the complex algorithms necessary for efficient channel hopping and mesh topologies. While lower frequencies (e.g. 900MHz) can also be used for WSN operations, and may increase transmission distance, availability of these products is trending downward due

to regulations in Europe and Asia. Transmission distances can be improved by using high-gain antennas to replace the more traditional mote whip antennas. Close attention to details, such as maintaining sufficient battery voltage at each mote, will additionally help to avoid network blackouts in winter. As most battery chemistries are adversely affected by colder weather, battery specifications should be checked to ensure the minimum amount of current required by the mote can at all times be drawn from the battery, regardless of external temperature.

2.5.2 Snow depth and soil moisture variability

This study further confirms the importance of canopy cover as a major explanatory variable of catchment-scale snow depth distribution. More so, it validates the claim that snow depth variability is highest under canopy [*Jost et al.,* 2007]. While those authors investigated the influence of forest cover on snow distribution during a number of surveys, analysis of the spatiotemporally dense data obtained during our study shows that there is a clear temporal stability in this canopy-dependent variability – under-canopy snow depth variability was persistently higher than that of open sites throughout WY2010.

Due to effects of partial canopy cover, sensor nodes located at drip edges of trees showed relatively high snow depth variability during the accumulation season, similar to that of under-canopy nodes. These nodes likely experienced lower long-wave radiation than under-canopy nodes, and but were exposed to higher short wave radiation in the spring. This further validates the claim by *Jost et al., (2007)* that it is important to consider detailed forest inventory, rather than the binary variable of forest cover, to derive better estimate of snow cover variability.

Our findings on the effect of soil depth on VWC variability are contrary to those of *Bogena et al.* (2010), but agree with *Grant et al.* (2004), who noted an increase in variability at deeper soil depths in a comparably sized mountain catchment. Similar to *Grant et al.* (2004), we attribute this effect to increased heterogeneity of soil properties at greater depths. We also share agreement with *Grant et al. (2004)* on the lack of a correlation between mean VWC and variability (CV), except for near surface measurements. This agreement suggests that there are unique characteristics of VWC variability in mountain catchments that do not follow common findings in a number of other studies [Famiglietti et al., 1998; *Western et al*., 2002; Bogena et al., 2010].

2.5.3 Network measurement validation and sampling design

The ability of the WSN to capture catchment-wide snow depth mean and variability has been confirmed through statistical and qualitative comparisons with the data obtained by the LIDAR and synoptic surveys. Given this agreement, it is reasonable to suggest that a stratified placement strategy based on evenly instrumenting major physiographic parameters performs well with regard to characterizing the distribution of catchment-wide snow depth at the km^2 scale. Some discrepancies between the snow depth data sets can be attributed to general sensor calibration error. Reflection due to canopy cover may also have introduced noise into the LIDAR data set.

The approximate 30-40 m support of the snow depth variogram resulting from the geostatistical analysis of both LIDAR and network data explains why snow depth surveys did not yield distinct variograms. The 125m spacing between survey locations did not capture smaller scale variability. If similar length variogram support is validated through future plot-scale

surveys, the resulting finding will underscore sampling design: to optimize placement of snow depth sensors, a stratified sampling strategy should be used to cover the range of physiographic parameters on the site, but nodes should not be placed within a variogram range (in this case 30m) of each other since values within this range can be confidently estimated through geostatistical methods (e.g. Kriging).

Given the near-normal distribution of snow depth exhibited in the LIDAR data set (Figure 2.13), an under-sampling of the catchment would not likely capture the extreme values of snow depth. Although the mean and variance estimates were not significantly affected, an example of such under-sampling is given by WSN measurements on March 22, which did not capture the relatively low values in the northwestern region of the catchment, and relatively high values in some of the open meadow regions (dark and light locations in Figure 2.13). Future placement of a set of sensor nodes in these areas is expected to minimize this bias in coming studies. The nature of WSN technology will permit for the seamless integration of these sensing locations into the existing network. Application of non-linear classification techniques [*Balk and Elder,* 2000; *Molotch et al.,* 2005b] could also reveal if the un-instrumented regions can be more accurately predicted from the instrumented subset of the catchment.

The existence of three sensor-node locations (Figure 2.13, yellow stars) that consistently reflected snow depth mean suggests that it may be possible to identify key measurement locations which capture mean snow depth across the catchment. This notion has been addressed in previous studies [*Molotch and Bales, 2005*], but analysis of multi-year data is required to validate this claim. These points are not associated with any particular physiographic features that could explain this behavior, suggesting that specific spatial coordinates, rather than physiographic attributes, could also serve as indicators of the overall catchment mean.

The discrepancies between the distributions of soil moisture for the synoptic survey and network data were likely influenced by effects of spatial scaling. If the network readings represent the catchment wide VWC distribution, then VWC variability increases over time as the soil begin to dry out after snowmelt. This increase in variability may not have been captured by the 125m spacing of the synoptic survey. This is particularly true in the case of the September VWC study, in which the network captured high VWC values in meadow regions, thus shifting the estimate of the mean, and increasing overall variability. The layout of the synoptic survey may have under-sampled the wet meadow regions, while the spacing between survey locations may have missed the short-scale variations in VWC. Sampling error could also have been introduced by the hand-held GPS units used in the study.

Although discrepancies exist between the VWC values captured by the network and synoptic surveys, results suggest that a strategically placed sensor network can effectively characterize the distribution of catchment-wide soil moisture at the km^2 scale. This follows in agreement with *Grant et al.* (*2004*), who noted that catchment-wide soil moisture could effectively be characterized by few samples. As *Western et al.* (2002) noted however, it may be of more interest to understand the detailed spatial arrangement of VWC, and effects of physiographic parameters on soil moisture. Such a detailed spatio-temporal analysis [*Famiglietti et al.,* 1998*; Williams et al.,* 2009; *Bogena et al*. 2010*]* requires a more spatially dense synoptic data collection.

2.6 Conclusions

When compared to synoptic surveys and dense LIDAR data, a sensor-placement strategy based on coverage of physiographic parameters can effectively capture the mean and variability

of snow depth across a catchment. Snow depth was also shown to vary significantly based on canopy cover. This variability exhibited temporal stability when comparing under-canopy and open terrain measurements. VWC variability increased significantly following snowmelt, making it difficult to gage sampling design given the density of the synoptic survey VWC data. While the network appeared to capture this increase in variability in the springtime, more detailed plot-scale surveys will be carried out in the future to validate this behavior. A weighted sampling approach, based on the relevance of physiographic parameters and physically based soil moisture, and snowmelt dynamics, may be more suitable for characterizing the VWC behavior at the km^2 scale.

Although WSN technology is continuously improving, no off-the-shelf solution exists for the harsh conditions experienced during most environmental monitoring campaigns. Additionally, few detailed findings are available in previous WSN studies to make conclusive statements regarding their observed performance. Results gathered from the 57-node WSN mesh described in this chapter partially fill that gap, by demonstrating a cost-effective means by which to instrument a large, and remote area in complex terrain. Low-power and reliable performance can be achieved by proper evaluation of hardware behavior in the specific field conditions. This deployment successfully covered a 1.5-km transect with a conservative 50 m node-to-node distances, and continues to reliably transmit data from more than 300 sensors every 15 minutes.

The sheer scale of future monitoring campaigns will demand many network nodes, making it impractical to optimize every single link. *In situ* network behavior must be quantified to derive average indicators of performance. Gathered over a span of several days, PDR and RSSI should be sufficient to estimate operational path behavior. It is expected that these network characteristics will vary based on site-specific terrain. Unless obvious physical obstacles exist (e.g. large rock outcroppings or metal surfaces blocking a path), the average maximum transmission distance extracted from such an analysis should be sufficient to inform network design.

An open mindset to current limitations of the hardware, and monitoring of network statistics lead to an iterative deployment approach, which can be split into three distinct phases. Since maintaining reliable network performance is the most important design imperative, trading transmission distance for a robust, hardened network is the proper decision. More repeater nodes can be added to improve connectivity and increase connection redundancy, but only collection and analysis of network statistics can shed light on actual causes of network failure, and lead to rational strategies for improving performance. Using this approach, the network presented in this study experienced significant boosts in performance, and no further collapses.

The implementation of the above deployment strategies should lead to successful WSN deployments of up to 100 nodes, and spanning 1-2 km^2. Environmental monitoring deployments of networks beyond this size should use multiple network managers, each responsible for networks of up to 100 motes. This will reduce the likelihood of large-scale network outages caused by bottleneck effects, while simplifying the identification and management of network behavior. Such architecture permits for manageable approach to scaling, when considering the instrumentation of large areas such as entire basins, or mountain ranges. The network described in this chapter is a prototype for monitoring such large areas, and future work will investigate the feasibility and methods of scaling this deployment strategy to cover significantly larger areas. Using these heuristics, deployment time can be significantly reduced, thus allowing more resources to be devoted to the actual scientific program.

Chapter 3: Sensor Placement Strategies for SWE Estimation in the American River Basin

The previous chapter validated the ability of distributed WSNs to resolve snowdepth variability across catchment-scale regions. When scaling this approach to monitor much larger basin-scale areas, individual networks should be deployed independently, and the resulting data from each network must be integrated to resolve hydrologic processes across the basin. As described previously, the deployment of clusters of networks ensures a manageable operational approach. In this sense, each network cluster can be abstracted as a single sensor at the basin scale, providing estimates of hydrologic variables for each catchment. These catchment-wide readings can then be weighted accordingly to derive a basin-wide estimate. The purpose of this chapter is to develop a sampling methodology (or sensor placement strategy) by which individual sensor network locations can be selected to most optimally derive the state of a hydrologic parameter across the wider basin. Of primary interest will the measurement and estimation of snow water equivalent across the American River Basin in the Sierra Nevada of California.

3.1 Introduction

Water managers often rely on a few point estimates of snow water equivalent (SWE) in a given mountain basin to inform streamflow forecasting. Accurate estimation of SWE across a basin, and thus forecasting of snowmelt runoff, particularly in a changing climate, is impeded by both limited on-ground measurements and high inter-annual variability in snow accumulation and ablation. Snowpack variability is driven by the intersection of climate variability with complex physiographic spatial patterns [*Bales et al*., 2006].

Though SWE is measured at more than 1700 points across the western U.S. through a combination of automated snow sensors and manual snow courses, these measurements are insufficient to resolve SWE variability at the basin scale [*Bales et al*., 2006; *Rice and Bales*, 2010]. Deployed with the intent of informing statistical streamflow predictions, the majority of current snow courses and snow sensors are placed in easily accessible, flat, open, and relatively mid-elevation locations where snow cover is persistent. These physiographically homogeneous locations result in redundant observations that may not reflect basin behavior as a whole. Moreover, it has been shown these sensors are not representative of immediately surrounding areas [*Molotch and Bales*, 2006], and that the bias between the surrounding terrain and snow sensors are of greater magnitude during the ablation period [*Meromy et al.*, in press]. This strategy under samples high elevations where snow can linger into the summer months, and low elevations where melting can occur throughout the winter [*Bales et al.*, 2006]. Current operational locations do not reflect the actual spatial distribution patterns of SWE, nor the spatial behavior of snowpack accumulation or ablation patterns.

Ideally, SWE can be estimated at a resolution that reduces the subgrid heterogeneity to a level where the majority of variability can be modeled explicitly [*Blöschl*, 1999]. Smaller-scale (<100 km^2) SWE-distribution approaches, using various geostatistical techniques such as *binary regression trees* and *Kriging*, have been implemented successfully, but require dense measurements which are typically confined in space and time [*Balk and Elder*, 2000; *Erikson et al*., 2005; *Erxleben et al*., 2002].

Satellite-derived snow-covered area (SCA) measurements aid in distributing ground measurements at the basin scale, and are appealing from a real-time operational hydrology perspective [*Painter et al.*, 2009; *Rice et al*., 2011]. Remotely sensed SCA products have been used in conjunction with ground-based SWE measurements via various interpolation strategies to build spatial SWE estimates [*Fassnacht et al*., 2003; *Molotch et al*., 2004; *Harshburger et al*., 2009]. These estimates have proven reliable for predicting existing snow stations, but do not create representative estimates of basin-wide SWE distribution due to the inherent bias in ground measurements. It is also not known to what extent the improvements in water balance estimates will quantitatively affect basin-scale water forecasting.

Energy-balance and precipitation-based models have been implemented on the basin scale and shown to compare well to snow-sensor measurements [*Shamir and Georgakakos*; 2006]. Reconstruction models provide an energy-balance based retrospective tool for estimating SWE distribution without using ground observations [*Cline et al*., 1998; *Molotch and Margulis*, 2009; *Rice et al*., 2011]. These models use satellite derived melt-out dates and SCA time series to back calculate SWE using an energy balance for each grid element. While this method does provide a retrospective, unbiased, and independent estimate of SWE historical distribution, it can not be used to make real-time SWE estimate.

Sensor placement falls under the larger umbrella of sampling design, a robust and growing field within statistics and computer science that has seen limited reported application to snow hydrology. Expanding the current snowpack sensing infrastructure requires the determination of how many sensors (or sampling locations) are required and where they should be placed within the basin to maximize the amount of captured information. *Molotch and Bales* (2005) and *Rice and Bales* (2010) searched for optimal sensor placements that capture the mean and variability of SWE across a larger grid element, but did not evaluate potential resulting placement error or the amount of information captured by would-be placements. While it is generally acknowledged that improved strategic sampling is needed to build better large-scale sensor deployments and SWE estimates [*Bales et al*., 2011], no quantitative methods or performance metrics have been proposed.

This chapter will use a multi-year historical reconstruction of SWE for the American River basin in California's Sierra Nevada to identify sensing locations for accurate, real-time estimation of basin-wide SWE. This research addresses three questions. First, what properties of basin-wide SWE distribution are temporally stationary? Second, how can we use these properties to inform sensor placement? Finally, can we develop a quantitative sampling method to generate sensor placements based on historical SWE observations? A more general aim is to develop a ground-based sampling strategy that will become a core element of a larger water-information system for the Sierra Nevada and other mountain basins.

3.2 Background

3.2.1 Sampling strategies

Given a large set of observable locations, selection of the “most informative” subset is a problem faced within numerous disciplines. Selecting ideal sensing locations for observation of spatial phenomena, in our case SWE, is a direct instance of this problem. One sensor-placement approach is to assume that sensors exhibit disk-shaped sensing regions, and to search for a set of sampling locations that provides ideal spatial coverage; this approach is formalized as the “art gallery problem” in the computer science literature [*Gonzalez-Banos and Latombe*, 2001].

Within snow hydrology, the concept of circular impact areas has been applied to estimate unobserved locations through the use of inverse distance weighting (IDW), optimal distance averaging (ODA), and Kriging [*Fassnacht et al.*, 2003; *Carroll*, 1995]. These techniques are often combined with methods such as linear regression against elevation to derive more robust estimates of SWE [*Carroll and Cressie*, 1996]. While these studies have validated the efficacy of these approaches for estimation purposes, these principles have not been applied to sampling design.

It is well documented that SWE and snow depth vary with physiographic variables, such as elevation, aspect, slope, and canopy cover [e.g., *Fassnacht et al*., 2003]. A reasonable sampling strategy would be to "evenly sample" across each predictor variable. Such strategies have been described within Computer Science. (see *Latin hypercube sampling* [*Mackay et al.,* 1979]). Unfortunately, it is difficult to quantify the effectiveness of such a placement scheme since evenly sampling physiographic features may not adequately capture the relative importance of any given feature.

An alternative approach is to couple sensor placement with modeling. This involves developing a statistical model of a phenomenon based on previous observations, and then designing a placement scheme that improves model performance. The statistics community has addressed the question of maximizing the quality of parameter estimates in linear models through *sampling design* [e.g., *Atkinson, 1*988]. Algorithms have also been developed to select sampling locations to yield the best possible "error reduction" in linear models [*Das and Kempe,* 2008].

A specific modeled-based approach, developed by *Cressie* [1991], involves developing a *Gaussian Process* model for an underlying spatial phenomena, either from expert knowledge or a pilot deployment. Use of this inference framework assumes underlying Gaussian distributed data. The power of this Bayesian technique stems from its ability to provide variance estimates for predictions made at un-instrumented locations, allowing us to frame a sensor placement approach by selecting those locations which most reduce model uncertainty across the estimated field. The following section will layout a detailed derivation of this framework.

3.2.2 Spatial Estimation of SWE

Given measurements provided by a sparse set of sensors, a common problem in distributed sensing deals with the estimation of a process at the set of remaining uninstrumented locations. In our example of SWE sensing, this involves using SWE measurements from a number of deployed WSNs to estimate SWE values in the remainder of the basin. One common approach is to use simple linear regression [*Wackerly,* 2008; *Fassnacht et al.*, 2003]. These formulations, however, do not stress the uncertainty of the prediction, or the underlying covariance structure of the data. As will soon become evident, for sampling design it desirable to obtain these uncertainty estimates along with the prediction. A powerful way to obtain a prediction of a spatial phenomenon, as well as its corresponding uncertainty, is through the use of a *Gaussian Process* (GP) model [*Rasmussen*, 2006], which can be derived by embedding the standard linear model into a Bayesian framework.

3.2.3 The Bayesian Linear Model

Begin with a training set D containing n observations,

$$D = \{(x_i, y_i) \mid i = 1,...,n\}, \qquad (1)$$

where **x** is an input vector of dimension D, and **y** is the observed one-dimensional dependent (output) variable. If all input vectors are part of a $D \times n$ *design matrix*, and the corresponding outputs are collected in the vector **y**, the training set can be written as $D = (X, y)$. The Bayesian linear regression model is then given by

$$f(x) = x^T w, \quad y = f(x) + \varepsilon \qquad (2)$$

where $f : x \to y$ maps an input vector to an output value, and w is the parameter (weight) vector. In this formulation it is assumed that the observed values of y contain an additive normally distributed noise ε of mean zero and variance σ_n^2 such that

$$\varepsilon \sim N(0, \sigma_n^2), \qquad (3)$$

where N(·,·) denotes a normal distribution. *Rasmussen* (2006) shows that through this noise assumption, when given the parameters and *design matrix*, the probability density function of the observation (also known as the likelihood) takes the form

$$p(y \mid X, w) = N(X^T w, \sigma_n^2 I). \qquad (4)$$

In this case, a zero mean prior, with covariance Σ_p is placed over the weight vectors, such that

$$w \sim N(0, \Sigma_p). \qquad (5)$$

The application of Baye's rule allows for the calculation of the posterior distribution of the weights, once the set of input and output vectors has been observed. It can be shown that the resulting relation takes the form of a Gaussian distribution:

$$p(w \mid X, y) \sim N(\frac{1}{\sigma_n^2} A^{-1} X y, A^{-1}) \qquad (6)$$

where $A = \sigma_n^{-2} X X^T + \Sigma_p^{-1}$. Now it is possible to make a prediction for a test input x_* by calculating a weighted average over all possible parameter values. The predictive distribution is given by

$$p(f_* \mid X, y, x_*) = \int p(f_* \mid x_*, w)\, p(w \mid X, y)\, dw = N\left(\frac{1}{\sigma_n^2} x_*^T A^{-1} X y, x_*^T A^{-1} x_*\right). \quad (7)$$

Equation 7 not only provides a prediction for a given input, but it is now also possible to derive a confidence estimate for this prediction.

3.2.4 Function Space Projections and Gaussian Processes

A Gaussian process (GP) is a collection of random variables, any number of which have a joint Gaussian distribution. The previously described regression technique can be posed as GP regression to reconstruct an underlying signal f given measurements with noise ε. Given a process $f(x)$, the GP is fully described by it's mean function $m(x)$ and covariance function $k(x, x')$:

$$m(x) = E[f(x)]$$

$$k(x, x') = E\left[(f(x) - m(x))(f(x') - m(x'))\right]. \quad (9)$$

To denote the distribution over functions, the GP is written as

$$f(x) \sim GP\left(m(x), k(x, x')\right). \quad (10)$$

In our example application, $f(x)$ will be the estimated SWE, and the input vector x is the observation of SWE made by a WSN. Without loss of generality, the mean function can be taken to be zero. Then, the previously presented Bayesian linear regression model is a GP described by

$$E\left[f(x)\right] = m(x) = x^T E\left[w\right] = 0$$

$$E\left[f(x) f(x')\right] = \mathrm{cov}(f(x), f(x')) = k(x, x') = x^T E\left[w w^T\right] x' = x^T \Sigma_p x'. \quad (11)$$

For more information on GP's see *Rasmussen* (2006).

3.2.5 Near-optimal Sensor Placement

When choosing from a set of possible sensing locations, a quantitative sampling design seeks to identify those sampling locations that maximize the amount of acquired information (minimize the overall uncertainty). The mathematical framework of *Information Theory* is one means by which to quantify the amount of uncertainty that can be informed within a process [*Shannon*, 1948]. A direct measure of uncertainty is given by *Entropy* [*Pierce*, 1980].

Adapting the notation of *Krause et al*. (2007), the problem can be formulated as finding the n best sensing locations, given a finite subset V of possible locations, and a finite set of already instrumented indices A. The Entropy of this placement depends on the probability

structure, and more specifically the covariance, of the underlying sensing space. The conditional entropy of the unobserved locations $V \setminus A$ after placing the sensors at locations A is given by

$$H(X_{V\setminus A} \mid X_A) = -\int\int p(x_{V\setminus A}, x_A) \log p(x_{V\setminus A} \mid x_A) dx_{V\setminus A} dx_A \qquad (12)$$

where X_A and $X_{V\setminus A}$ are sets of random variables at locations A and $V \setminus A$. Minimizing this property has the effect of finding the set of sampling locations A^* which are most uncertain about each other [*Kemppainen et al.,* 2008; *Lee and Queyranne*, 1995]. More formally,

$$A^* = \operatorname{argmin}_{A \subset V : |A| = n} H(X_{V\setminus A} \mid X_A) \qquad (13)$$

This optimization problem is also known an *D-optimal Design*, and has been shown to be NP-hard [*Karuse et al.,* 2007]. The main drawback to this approach is that the selection process favors locations which are most uncertain about each other, to the extent that it places sensors so far away from each other as to push them near the boundary of the sensing region. Sensing near boundary regions of a sampling space has the tendency to waste information, since it may not fully utilize the entire coverage area of a sensor.

This approach can be augmented by reformulating the problem to search for the subset of sensor locations that most significantly reduce the uncertainty about estimates in the remainder of the space. Let the space of discrete locations be $V = S \cup U$, where S is the set of positions where placement is allowed, and U is the set of locations where no placement is possible. The goal then is to select sampling locations to reduce uncertainty at uninstrumented locations $V \setminus A$. Formally, this equates to computing the optimal set A^* where

$$A^* = \arg\max_{A \subseteq S : |A| = n} H(X_{V\setminus A}) - H(X_{V\setminus A} \mid X_A). \qquad (14)$$

This approach has the intended effect of maximally reducing the entropy of the uninstrumented space. The above expression was introduced by *Guestirn et al.* 2005, and is equivalent to maximizing the Information Theoretic quantity of *Mutual Information* (MI). *Krause et al.* (2007) prove that the problem of maximizing Mutual Information is NP-hard, but provide a alternative polynomial time near-optimal algorithm with a constant factor guarantee. Algorithm 3.1 is greedy, adding sensors in a sequence by choosing the next sensor to maximally increase the mutual information of the placement. More formally, the goal is to select a sensor y to maximize

$$MI(A \cup y) - MI(A) \qquad (15)$$

which the authors show to be equivalent to

$$H(y \mid A) - H(y \mid V \setminus (A \cup y)). \qquad (16)$$

such that

$$A^* = \arg\max_{A \subseteq S : |A| = n} H(y \mid A) - H(y \mid V \setminus (A \cup y)). \quad (17)$$

In general, model-based sampling design presents a difficult computational problem, in many cases leaving little alternative than the use of greedy sensor placement algorithms. Rather than finding the optimal set of sensor locations, greedy algorithms place sensors one at a time, with each step incrementally maximizing the Mutual Information of a placement. Although greedy algorithms cannot guarantee a truly optimal solution, they are considered near-optimal, and provide performance guarantees, ensuring that algorithm performance in no worse than some fraction of the optimal solution [*Guestrin et al*., 2005; *Krause et al.,* 2007].

input: Σ_{VV} (covariance matrix), n (number of sensors), $V = S \cup U$ (locations)

$A \leftarrow \varnothing$

for $j = 1$ **to** n **do** $\tilde{A}$

  for $y \in S \setminus A$ **do** $\delta_y \leftarrow \dfrac{\sigma_y^2 - \Sigma_{yA}\Sigma_{AA}^{-1}\Sigma_{Ay}}{\sigma_y^2 - \Sigma_{y\tilde{A}}\Sigma_{\tilde{A}\tilde{A}}^{-1}\Sigma_{\tilde{A}y}}$

  $y^* \leftarrow \arg\max_{y \in S \setminus A} \delta_y$

  $A \leftarrow A \cup y$

Algorithm 3.1: Greedy algorithm for maximizing mutual information [*Karuse et al.* 2007]. For notational convenience, $\tilde{A} = V \setminus (A \cup y)$. Mutual information is calculated using elements of the covariance matrix of the locations.

3.2.6 Regionalization (clustering) of data

In certain conditions the parameterization of a single statistical model may not be sufficient to explain the variability of an entire data set. In the case of SWE, it may be necessary to use multiple models, each of which describes a subset of the observed data. Effective grouping of data into homogenous subsets has been shown to improve model performance [*Chipman et al*., 2002; *Hannah and Dunson*, 2011], and is a common pre-analysis step within hydrology investigations [*Serreze et al*., 1999; *Clark et al*., 2001]. Observations are often grouped by geographic proximity [*Moore and McKendry*, 1996], but more complex statistical methods have also been applied to snow data such as principal component analysis [*McGinnis*, 1997; *Cayan*, 1996], and self-organizing maps [*Fassnacht el al*., 2010].

Wu and Lui [2012] searched for statistically similar regions within soil moisture data using a coarse-grained ordering approach. This method assumes that while overall moisture levels will vary significantly throughout space and time, certain locations will typically be drier or wetter than others. This phenomenon is measured by sorting from least to greatest, by wetness, all locations at each time step and analyzing the stability of this ordering over time. Such ordering is ultimately used to divide the study area in homogeneous regions, or clusters, roughly corresponding to wetter and drier areas which can be explained by separate statistical models.

3.3 Methods

An eleven-year spatially distributed SWE data set for the American River basin, with a grid size of 500 x 500 m, was used to determine fixed measurement locations for estimating SWE during snowmelt under a range of inter-annual climate conditions. The eleven years of data were then divided into training and testing sets, the first eight years (2001-2008) were used to train the sensor placement algorithm and to develop a historical SWE covariance matrix, and the next three years (2009 -2011) were used for evaluation of the placement algorithms.

3.3.1 Study area

The American River flows westward from the crest of the Sierra Nevada, with its three forks emptying into Folsom Reservoir. Basin elevations range from 200 m at Folsom reservoir to over 3000 m in the Desolation Wilderness of the Eldorado National Forest. Snowfall typically is greatest and persists the longest at higher elevations. Of the 4741 km^2 above Folsom, 2155 km^2 is above 1500 m, and 242 km^2 is above 2400 m. Pixels with mean elevations below 1500 m were not used in the analysis, as these elevations do not significantly contribute to overall SWE in the American River basin (Figure 3.1).

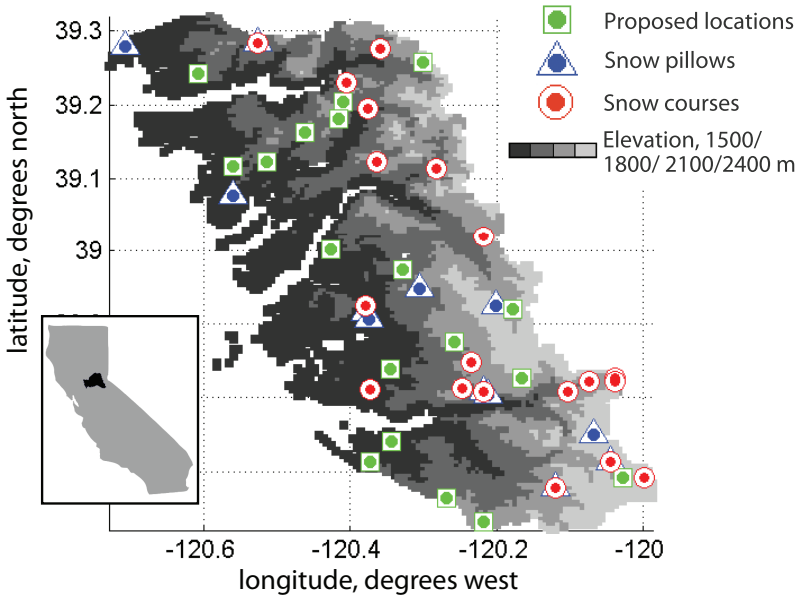

Figure 3.1 The American River Basin above 1500 m elevation, divided into approximately 7000 500-m square pixels. Snow-pillow and snow-course locations are obtained from California Cooperative Snow Survey. Optimal sensing locations selected through the method proposed in this paper are shown as green squares.

3.3.2 Data

The spatial SWE dataset used in this analysis was developed using a pixel-by-pixel energy-balance reconstruction of snowmelt, summed back in time from when snowcover disappeared to the beginning of seasonal snowmelt. Methods are summarized below, with details given in *Rittger et al.* [2011] and *Rittger* [2012].

The reconstruction method combines the satellite-derived snow-cover-depletion record of the melt rate to retrospectively estimate how much snow had existed at every pixel. The technique has been validated in the Sierra Nevada [*Cline et al*., 1998; *Rittger et al*., 2011] and applied to large basins at multiple scales [*Molotch and Margulis*, 2008; *Rice et al*., 2011]. SWE was estimated using a snowmelt model that combines both energy-balance and temperature-index methods [*Brubaker et al.*, 1996]. Air temperature and incoming solar and longwave radiation were downscaled from 1/8° NLDAS2 reanalysis data [*Cosgrove et al.*, 2003] to adjust for elevation, topography and vegetation [*Dubayah*, 1992; *Dozier and Frew,*1990; *Garen and Marks* 2005]. Reflected solar radiation was estimated using MODSCAG albedo [*Painter et al*., 2009] and outgoing longwave radiation was estimated using the Stefan-Boltzmann equation assuming the snow temperature was the minimum of air temperature and 0° C.

Fractional SCA used to scale the potential melt estimate for each pixel was derived from the MODSCAG model [*Rittger et al.*, 2012; *Dozier et al*., 2008]. Validation of maximum SWE was performed in a 3x3 grid cell region around each snow pillow and snow course by calculating the mean error in the nine cells and the best cell within that region. For snow pillows the error RMSE was 242 mm and 186 mm with a mean difference of 19 mm and 3mm. For snow courses the RMSE was 296 mm and 230 mm with a mean difference of 77 mm and 18 mm. Snow pillows are considered a better validation because they measure snow daily or hourly, while snow courses only measure SWE on the first of each month. The MODIS satellite became operational in the year 2000, limiting the present work to 11 years of reconstructed data. Figure 3.2 shows the reconstruction product in the American River Basin through the 2006 melt season. Prior to modeling, SWE training data were standardized each day by subtracting the pixel mean and dividing by the standard deviation. Since mean SWE varied significantly each season, this standardization step was necessary to extract stationary features of SWE variability. Historical monthly snow-course and daily snow-sensor data were obtained from the California Data Exchange Center (http://cdec.water.ca.gov). There are 12 snow pillows and 26 snow courses in the basin (Figure 3.1).

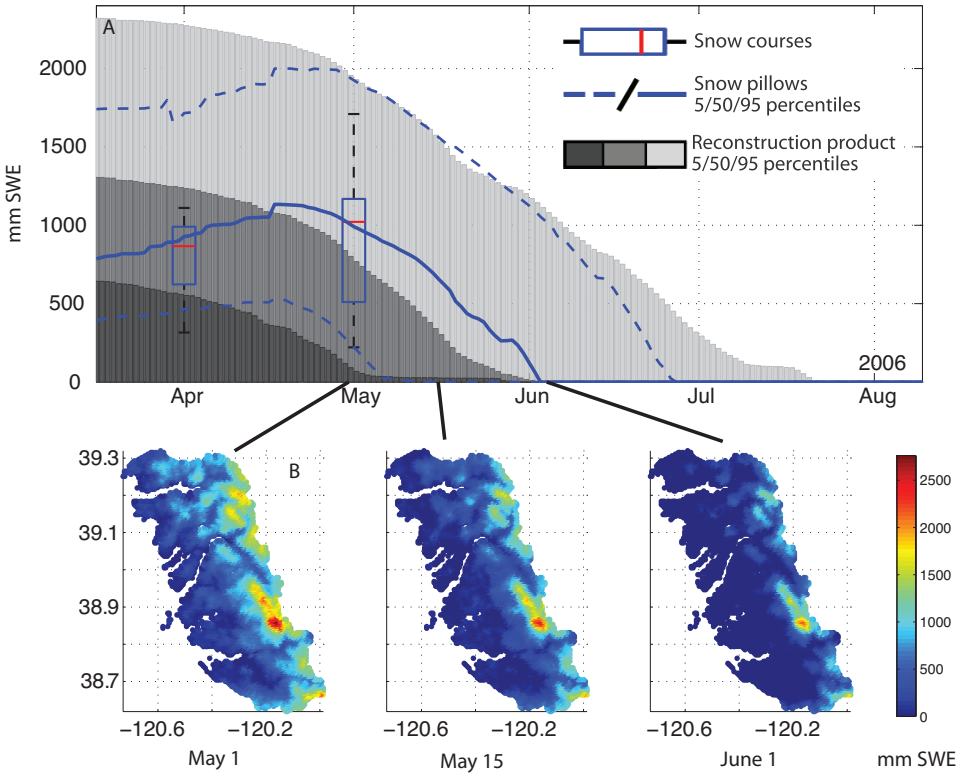

Figure 3.2 (a) Snow course, snow pillow, and interpolation product over 2006 melt season (b) Time slices of interpolation product shown for May 1, May 15, and June 1.

3.3.3 Clustering analysis

Multiple clustering methods were performed, investigating the existence of stationary subsets to inform sensor placement within the American River SWE data. The first and simplest method, called *global clustering,* treated the entire American River basin as one cluster on which to conduct sensor placement. The second method, *geographic clustering*, intended to capture smaller-scale SWE variability by delineating the larger basin into six third-order sub basins. The third method, *rank clustering*, clustered similar regions in space through the analysis of historical behavior of each SWE pixel. Daily SWE data were ordered from least to greatest, resulting in a "rank" for each pixel for every day studied. The sum of ranks was taken for all days studied, and pixels were then clustered based on sum rank values using a K-means algorithm as the number of clusters used was varied between two and eight. K-means is a simple clustering algorithm that minimizes the Euclidean distance between each clusters' centroid in K-dimensional space and the data within that cluster [*Mackay*, 2003]. K-means used here was one dimensional, as a simple tool to ensure that the most similarly ranked pixels were assigned to the same clusters. Rank stability was evaluated by plotting histograms of pixel rank [*Wu and Lui*, 2012]. It was observed that although ranking were not necessarily stable throughout the year, rankings during melt seasons were relatively stable (Figure 3.3). It is difficult to establish a consistent "melt-season" definition year to year; during our analysis simply using April 15- June 1 produced the stability needed to carry out our analyses.

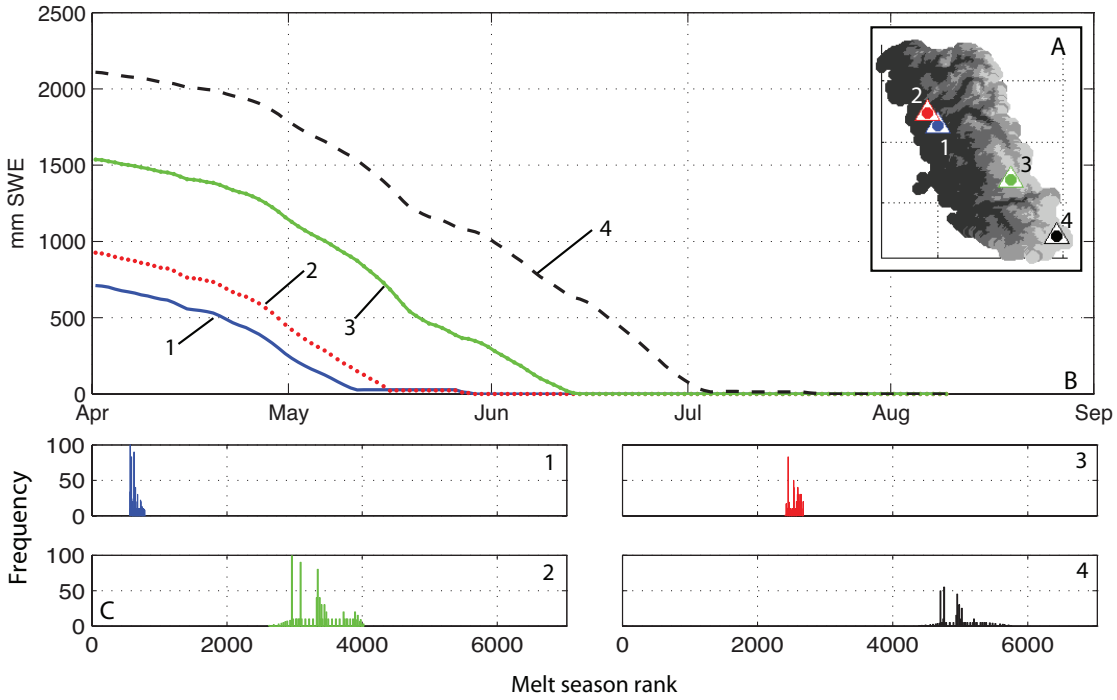

Figure 3.3 (a) One pixel chosen randomly from each elevation band (b) time series SWE for each pixel (c) Ranks across melt season.

3.3.4 Dimensionality reduction

Both GP modeling and a number of other placement schemes rely heavily on the covariance matrix of the input data. Performing the needed operations with covariance matrices requires the matrix to be well-conditioned and close to full rank. Poorly conditioned or rank deficient covariance matrices can result from two or more pixels having SWE time series that are too similar. If all the information known about one pixel is very close to that of another, it becomes computationally impossible to store the difference within a covariance matrix that must describe the entire space. This results in covariance matrices that require computationally impossible numerical precision to be operated on.

A solution to this problem is to reduce the dimension of the input data – group highly statistically similar pixels together, and use representative pixels from each independent group in modeling and placement algorithms. Dimensionally-reduced data are at the core of our placement and modeling algorithms, and then projected back onto the full space for final predictions. A K-means algorithm was applied to all the clusters resulting from the previous section, isolating sufficiently independent time series, from which a historical covariance matrix could then be constructed. It should be noted that unlike the K-means approach that was used above to build rank based clusters, K-means was used here in a higher dimensional space to find and group pixels with the most similar time series. While we are analyzing approximately 7000 pixels, the nature of the input data allows us to effectively operate on 50-300 representative pixels per cluster. An important implication here is that the nature of the underlying dataset ultimately determines the precision with which sensors can be placed.

3.3.5 Sensor Placement

Sensor-placement algorithms were applied within clusters, and constructed under the premise of maximizing information from historical data. The approach was evaluated by placing sensors using historical training data for 2001-2008, and evaluating the placements on data sets in years 2009-2011. It is important to note that years 2009-2011 were used strictly for validation of the placement approach, and were not in any way used to inform the placements or estimations. Two sensor-placement schemes were compared: *Mutual Information* (MI) based placement and random placements chosen from the dimensionally reduced subset described in the previous section. A MATLAB toolbox for *Submodular Function Optimization* was used to carry out the optimization problem for the MI-based placement [*Krause* et al., 2010].

In the current analysis, "sensor" refers to an instrumented 500 × 500 m pixel assuming that the average SWE in that area is accurately measured by the sensor. In practice, a "sensor" of this scale consists of 10-20 individual, fixed snow-measurement nodes (see chapter two) that strategically sample the pixel and are combined to give a pixel-scale SWE value [*Rice et al.,* 2011; *Bales et al*., 2011]. When comparing global placement to cluster-based placements, one must decide in what order to add sensors to clusters. Here, sensors were simply added numerically by cluster number (e.g. for 5 clusters and 12 total sensors placed, clusters 1 and 2 each have three sensors, and clusters 3 through 5 each have two sensors).

3.3.6 Modeling

To evaluate potential placements, Gaussian Process models were used to make predictions of SWE based on historical observations as well as observed locations selected for each placement scheme. The 2001-2008 data were used to build covariance matrices, effectively "training" the Gaussian Processes model. The covariance matrix was then used in conjunction with the locations selected by each sensor placement to derive predictions for years 2009-2011 at un-instrumented locations using equation (1). Predictions were then back-standardized using the sample mean and variance of observed locations from the training set. This constraint meant that it was necessary to place at least two sensors within each cluster, in order to calculate a sample standard deviation.

3.3.7 Evaluation

The root mean square error (RMSE) between the original SWE product (years 2009-2011) and model output was used to evaluate the efficacy of a given placement. RMSE was used to compare placement approaches, the effect of adding more sensors, as well as the effect of varying the number of clusters. To evaluate placements across multiple time steps, average RMSE was calculated as:

$$\mathrm{AvgRMSE} = \frac{1}{\mathrm{J}}\sum_{\mathrm{j}=1}^{\mathrm{J}}\frac{1}{\mathrm{T}}\sum_{\mathrm{t}=1}^{\mathrm{T}}\sqrt{\frac{\sum_{\mathrm{i}=1}^{\mathrm{n}}(\hat{\mathrm{x}}_{\mathrm{i}}-\mathrm{x}_{\mathrm{i}})^2}{\mathrm{n}}} \qquad (18)$$

Where J is the number of trails conducted, T is the number of time slices used, n is the number of pixels analyzed, $\hat{x}$ is estimated SWE, and x is actual SWE. This formulation of RMSE permitted us to cumulatively evaluate RMSE across three years. RMSE was only calculated for

pixels with actual SWE above ten mm, with the assumption that nearly snow free pixels can be realized real time via satellite SCA products. For random placements, ten trials were conducted and mean RMSE was computed ($J = 10$), while MI based placements required RMSE only to be computed once ($J = 1$).

3.4 Results

3.4.1 Rank-cluster analysis

Rank-based clustering behavior was consistent throughout the eight-year analysis period. Figure 3.4 shows results from the 2006 season, which was comparable to those seen in other years. Rank-based clustering resulted in clusters that roughly followed basin elevation bands (Figure 3.4a). Higher variability was typically observed in higher-elevation clusters (Figure 3.4b). On April 1st, 2006 (shown here to be representative of the beginning of melt) standard deviations varied between approximately 400 mm for the highest SWE cluster (cluster 5, 2100 mm peak SWE) to 150 mm for the lowest SWE cluster (cluster 1, 700 mm peak SWE). In the 2006 season, each cluster exhibited a noticeably unique melt-out date ranging over a two-month span from late April to late June. Rank-based clusters divided the bimodal SWE distribution of the whole basin (April 15, 2006) into near-Gaussian distributions of various spreads (Figure 3.4c). Rank-based clusters divided physiographic predictor variables primarily by elevation (Figure 3.5), generating near-Gaussian distributions of elevation, with the exception of the lowest elevation-cluster, which was cut off at 1500 m as described in the methods section. No notable correlation was seen between elevation and aspect within clusters.

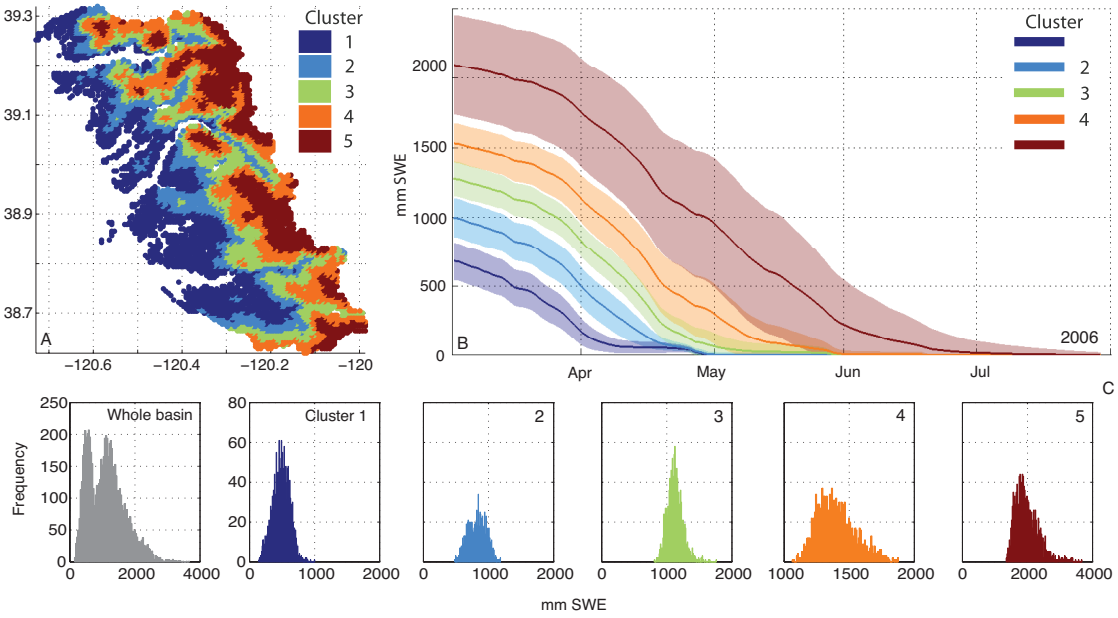

Figure 3.4 (a) Five rank based clusters (b) Cluster mean and standard deviation over 2006 season (c) Histograms of SWE by cluster for April 15.

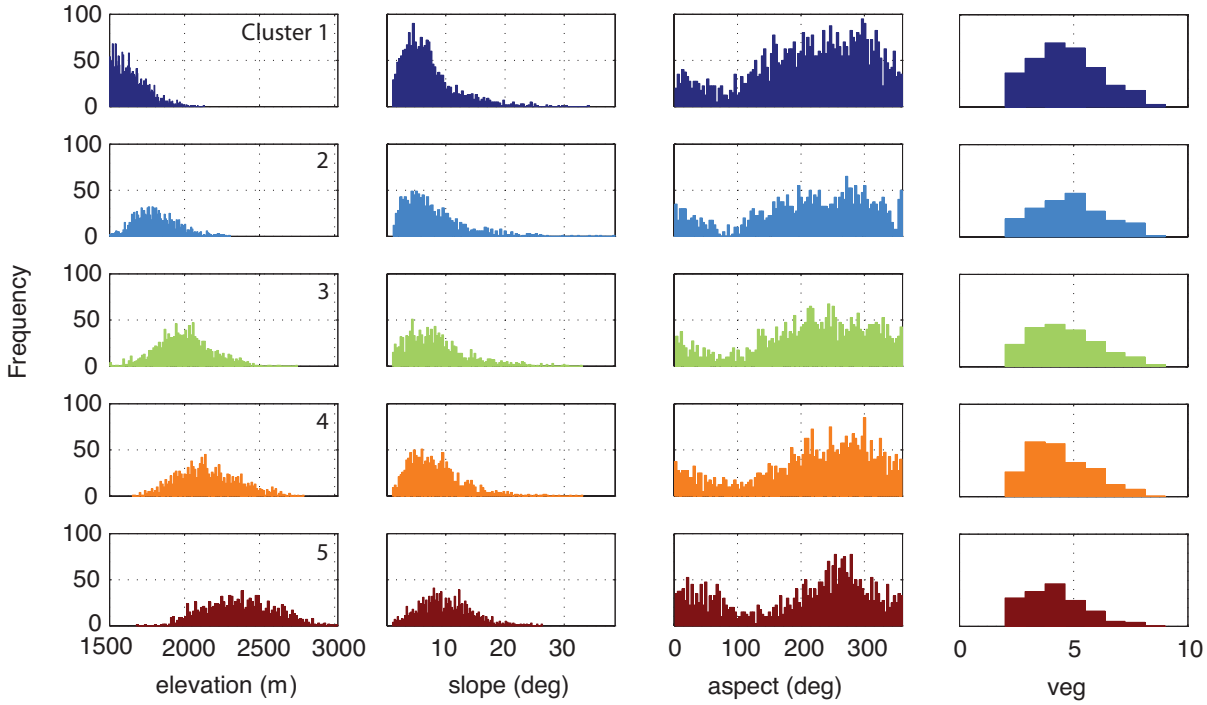

Figure 3.5 Histograms of physiographic variables divided by rank-based cluster.

Inter-annual differences between rank-based clusters ranged from 22 to 41%, with a median value of 31% (Table 2.1). This is a measure of stationarity that evaluates if a pixel retained its cluster assignment between these two periods. Differences were larger when comparing results for fifteen day periods within one year, with values of 0.13-0.67 and a median near 0.5 (Table 2.2). Cluster assignments varied significantly when comparing different periods during a single season, but agreed when comparing multiple seasons to each other.

2002	2003	2004	2005	2006	2007	2008	2009	2010	
0.302	0.315	0.397	0.340	0.342	0.336	0.312	0.316	0.363	2001
	0.253	0.369	0.300	0.242	0.313	0.296	0.328	0.278	2002
		0.290	0.223	0.250	0.287	0.316	0.304	0.291	2003
			0.308	0.344	0.339	0.366	0.329	0.408	2004
				0.269	0.323	0.340	0.304	0.316	2005
					0.307	0.302	0.288	0.317	2006
						0.317	0.270	0.390	2007
							0.275	0.365	2008
								0.407	2009

Dry
Normal
Wet

Table 3.1 Inter-annual cluster difference, fraction of pixels in a different cluster between years

Apr 15-30	May 1-15	May 15-31	June 1-15	
0.322	0.548	0.634	0.671	Apr 1-15
	0.443	0.551	0.592	Apr 16-30
		0.196	0.281	May 1-15
			0.132	May 16-31

Table 3.2 Intra-melt-season cluster difference, fraction of pixels in a different cluster between periods

3.4.2 Modeling Results

Analysis of sensor placement methods showed that average RMSE over the three years studied (2009-2011) decreased with the number of sensors placed (Figure 3.6). Diminishing returns occurred between 20 and 30 sensors placed, after which placing further sensors, or taking additional samples, yielded little improvement in RMSE.

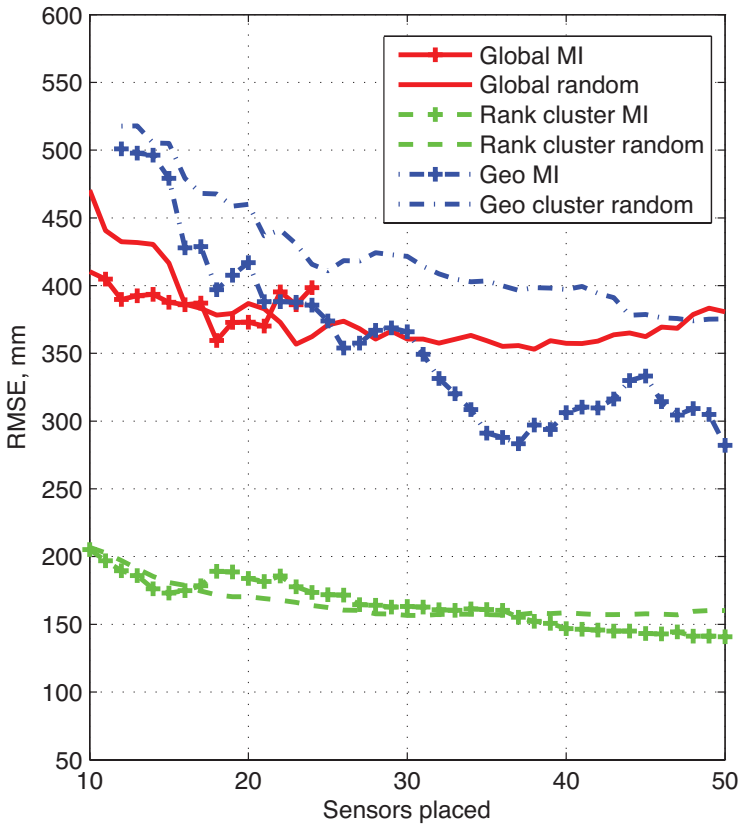

Figure 3.6 Average RMSE vs. sensors placed , averaged over 2009-2011 melt seasons. Rank-based results shown for 5 clusters.

Sensor placement within a rank-based clustering approach outperformed geographic clustering and global placement. Within each clustering method, MI based sensor placements generally offered improvements in RMSE when compared to random placements within the same clusters. MI-based placements within geographic clusters offered an RMSE improvement

between 0 and 150 mm compared to a random approach, while MI-based placements within rank-based clusters yielded between 30 mm greater and 30 mm lower RMSE than random placements.

To investigate temporal RMSE behavior, and taking the 2009 season as a typical year, daily RMSE for rank-based and geographic clustered placements showed a general decrease approaching melt-out of the snowpack. Global placement showed unstable and high RMSE values (Figure 3.7). Rank-based cluster RMSE was consistently below 200 mm, while geo clustering had RMSE values below 500 mm. Rank-based clustering, along with MI-based sensor placement showed the best RMSE performance throughout the season.

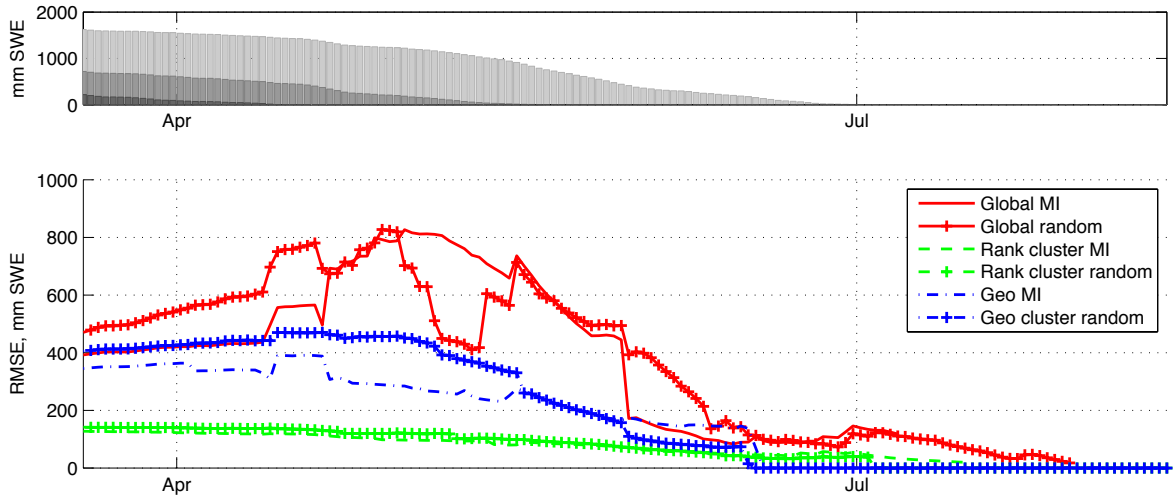

Figure 3.7 (a) 5, 50, and 95th percentiles of SWE across whole basin (b) RMSE over 2009 melt season. Rank-based cluster results shown for five clusters.

Given the performance of a rank-based clustering approach, an analysis was carried out to determine the effect on RMSE when varying the number of clusters while keeping the number of sensors per cluster steady. RMSE generally decreased as more rank-based clusters were used in the placement (Figure 3.8), where the use of one cluster is equivalent to a *global* clustering approach. RMSE decreased from a range of 400-600 mm in the case of *global* placement, to less than 200 mm when using eight clusters. Diminishing returns in RMSE became apparent in the three to six cluster range, after which further rank-based clustering offered no significant improvement in RMSE. MI-based placement offered slight improvements over random placement.

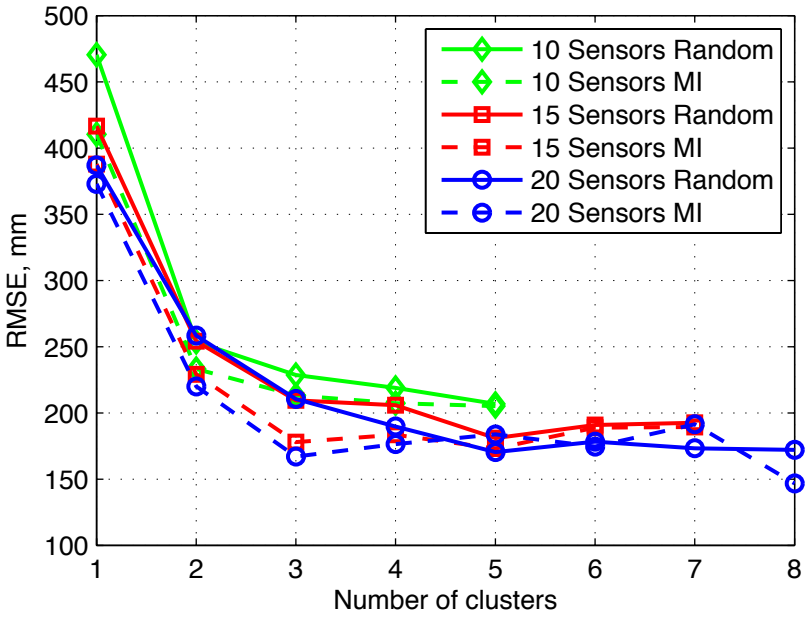

Figure 3.8 Average RMSE vs. number of rank based clusters. At lease two sensors are required per cluster.

Figure 3.9 shows the final proposed sensing regions resulting from coupling rank-based clustering with MI-based sensor placement. As discussed, the nature of the input data limits the precision of the covariance structure. Once optimal locations were selected using the placement algorithm, they were back projected from the dimensionally reduced data set into the original, higher-dimensional data set. Thus, rather than calculating single optimal locations, the algorithm produces optimal sensing areas where each pixel is considered equivalent. In Figure 3.9 each round corresponds to placing one sensor within each cluster; all locations indicated in a single round are equally valuable.

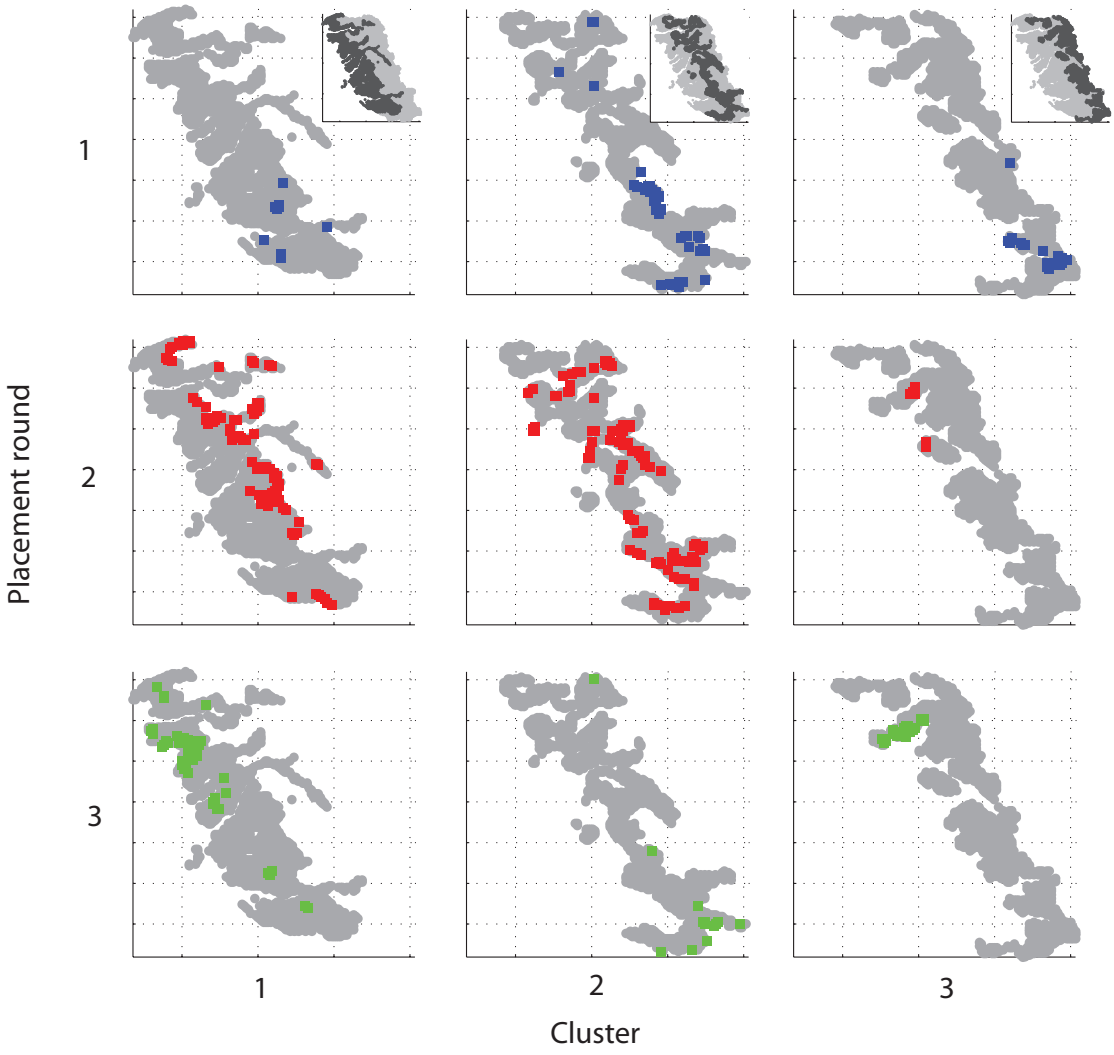

Figure 3.9 Recommended regions for placements in three rank-based clusters using *Mutual Information* based placement algorithm. Round one is the most informative placement, round two the second most and so on. One sensor should be placed within a given region each round. All proposed locations within a given region are equally valuable.

3.5 Discussion

3.5.1 Temporal stability through rank based clustering

A major goal of this analysis was to seek out stationary metrics to inform sensor placement. As seen through the reconstructed data, SWE is non-stationary with regard to annual mean and variance, making it difficult to infer one statistical model to cover the full range of annual spatiotemporal fluctuations. Some level of stationarity does however need to be extracted

to make use of historical observations for sensor placement. It is well known that SWE is correlated to a number of stationary physiographic variables such as elevation, slope, aspect, and vegetation [*Fassnacht et al.*, 2003]. However, this does not mean that the correlation between SWE and its physiographic predictors is stationary. As an example, Figure 3.10 shows perhaps the most common correlation metric used when studying SWE, the slope of the regression line taken between SWE and elevation. While this relationship may be robust over a single time step, it is temporally non-stationary, both inter- and intra-melt season. This non-stationarity impedes the use of these model parameters for purposes of sensor placement.

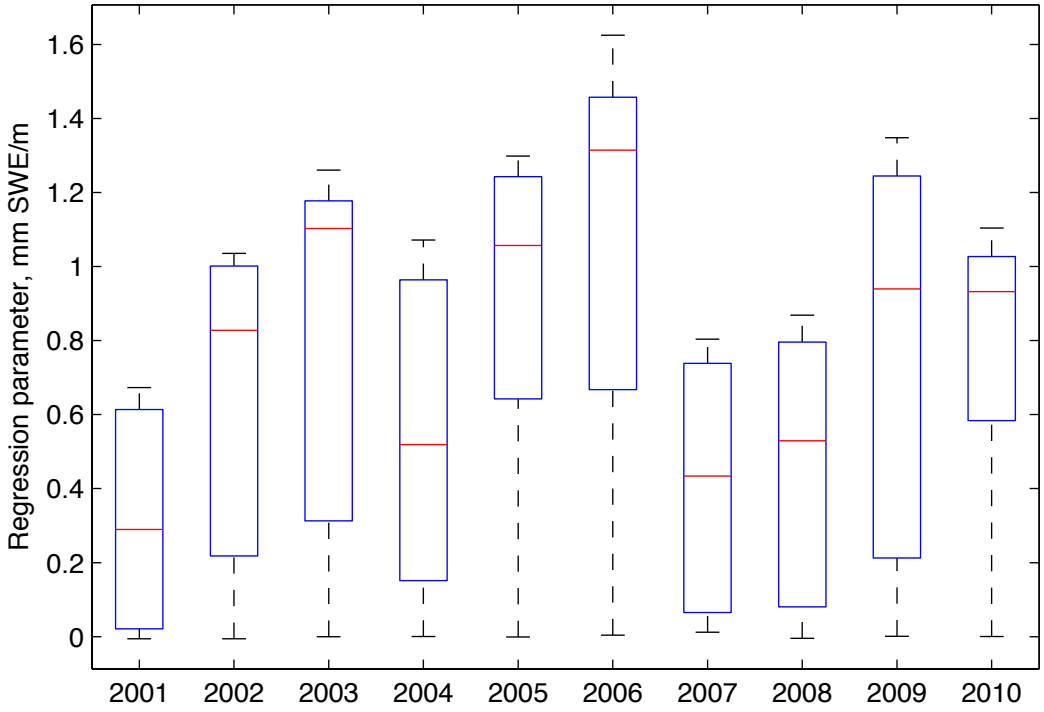

Figure 3.10 Elevation vs. SWE regression parameter computed daily across each melt season for American River Basin.

Rather than finding absolute trends in the data, rank-based clustering isolates subsets of that data that exhibit relatively similar behavior through time. Tables 3.1 and 3.2 show that rank-based clustering consistently isolated common clusters between years, with worst-case inter-annual cluster differences not exceeding 41%, while other measures of basin-wide SWE, such as regression against elevation (Figure 3.10), vary more than one order of magnitude inter-annually. Although SWE itself is quite variable, certain locations have consistently higher or lower SWE than others – this phenomenon is well exploited by rank-based clustering. This provided a strong starting point for sensor placement efforts. While Table 3.2 does show higher intra-annual differences, these can largely be attributed to the melt out of pixels toward the end of a season: 48% of pixels with snow during the April 1-15 period had no snow by June 1. Implementing cluster analysis on a multi-year data set maximizes the use of all historical data while providing a balance between long-term trends and inter-seasonal variability.

Figure 3.5 indicates that rank-based clusters divide primarily on elevation, a reasonable result considering the strong correlation broadly seen within the literature [*Fassnacht et al.,* 2003,

Bales et al., 2006]. Of interest here is that elevation is not simply binned into bands, but divided into overlapping near-Gaussian distributions. A similar result is observed when examining SWE histograms for each cluster, which also exhibit near-Gaussian distributions (Figure 3.4). This trend may have strong implications for sensor placement, and could partially explain why a random sensor-placement approach performed well.

3.5.2 Placement strategy within clusters

Dividing the sampling space using rank-based clustering is the first step of the placement approach described here. Within the derived clusters, further steps were taken to identify sensors placements that best inform future predictions. Dimensionality reduction, a computationally necessary step, is inherently part of the placement process. Typical clusters (1000-7000 pixels) were reduced to approximately 50 representative pixels to form computationally workable covariance matrices. It is within this reduced space that MI and random placement methods were carried out. Allowing selection exclusively from this subset means that we are already guaranteed reasonably independent observations, regardless of how placements are chosen in the reduced space. In practice, working in this reduced (and standardized) space limits the potential benefits of a MI placement approach, explaining the inconsistent RMSE improvements generated by MI placement over random placement seen in Figures 3.6, 3.7, and 3.8. Similar results were observed by *Wu and Lui* [2012] on soil-moisture data.

While MI based placement did not always necessarily outperform random placement, the approach does offer other benefits. Perhaps the most significant is a performance guarantee. The RMSE results of random sampling approaches were the mean of many trials, spanning a range of good and poor placements. It may be entirely feasible that a real-world random sampling approach could select one such poor placement. On the other hand, mutual-information-based placements are unique for a given covariance matrix, and will always select the same sensor locations in the same order, consistently guaranteeing the indicated performance. A second benefit of the MI-based sensor-placement approach is the ability to start with existing snow-pillow and snow-course sites, and then select additional sites through the placement algorithm. This should further improve estimation performance by maximizing the utility of data obtained from existing sites. The validation of the proposed algorithms used a historical covariance structure to predict SWE at future years based on observations made though relatively sparse sensing locations. Given the statistical nature of the modeling approach, it is feasible to update the covariance structure through real-time observations, and to add more sensors as needed. Given reported RMSE performance and consistent near-optimal placement guarantees, we recommend an MI-based approach.

In some cases the RMSE slightly increased though the addition of more sensors (Figure 3.6). These perturbations can be explained through the real-world nature of the data set. RMSE is not a monotonic function, and the validation data set may contain noise or outliers. Existence of outliers or data-set noise can adversely affect the RMSE behavior.

3.5.3 Number of Clusters

Figure 3.8 shows that three to eight clusters should be sufficient to model SWE and conduct sensor placement within the American River basin. Modeling with more clusters does not improve RMSE performance significantly. Furthermore, from the perspective of good

statistical practice, it is desirable to reduced model complexity to avoid the possibility of over fitting. With the current historical data sets we thus propose that three to five rank-based clusters are an optimal choice for the American River. The number of available sensor should then be divided equally throughout each cluster, or added proportionally to the variance of clusters, as done by *Wu and Lui* (2012).

3.5.4 Overall Sampling Approach

Based on the RMSE performance of the placement approaches implemented here, we are confident that one pixel in each of the regions shown in Figure 3.9 should be instrumented to derive better spatial estimates of SWE. In practice, specific points can be chosen within each proposed area based on accessibility or site preference. Figure 3.1 shows one such ideal placement resulting form our proposed method.

The method proposed in this chapter is designed to select sensing locations across a large area (over 1000 km^2). Each sensor location was treated as a 500 ×500 m pixel. It was assumed that instrumenting a pixel provides a noise-free representative measurement for that area. The locations selected in this paper should be thought of as strategically placed sets of individual nodes, with the set of nodes together forming the spatial sensor. Realistically, SWE is known to vary significantly at small scales. As such, it is unlikely that a single measurement within a 500-m cell will accurately represent mean SWE of the area [*Molotch and Bales*, 2005; *Rice and Bales*, 2010]. Multiple sensors (low cost depth sensors coupled with pillow measurements) will be required to reliably instrument one pixel. Sensor-placement strategies for smaller-scale areas thus need to be considered as part of a basin-scale design. Thus the overall strategy involves a hierarchical placement: i) optimal locations at the large scale (more than one square kilometer) should be selected based on the methods presented in this paper, and ii) estimates of mean SWE at each one of these locations should be derived though a secondary placement scheme at the smaller scale to capture local variability in aspect, vegetation, and energy balance. Each of the locations selected in this paper can then be thought of as a spatial sensor made up of a set of individual sensor nodes. We are confident that this presents the most realistically manageable sensor placement approach when considering the future instrumentation of large basins such as that of the American River. The blending of these strategically placed on-the-ground measurements with broad-coverage satellite and aircraft measurements will offer unprecedented estimates of snowpack, soil moisture, vegetation state and energy balance, and snowmelt [*Rice et al.*, 2009].

3.6 Conclusions

The approach in this paper provides a sampling strategy that uses historical data to select sensing locations that will be most informative for future real-time estimation of SWE in the American River basin. While the spatio-temporal SWE patterns are a non-stationary process, a rank-based clustering approach derives a set of regions that remain relatively stable over time. Following standardization of the annual SWE data, the variability of SWE within rank-based clusters was explained through a stationary covariance structure. This covariance structure was used to inform a quantitative placement scheme, by greedily placing sensors to maximize *Mutual Information* within each cluster. By splitting the existing data set into a training and test subsets, our approach reduces the RMSE of the estimation as a function of the number of sensors placed and outperforms non rank-based clustering approaches. Furthermore, there is a point of marginal

returns, after which placing more sensors does not improve estimation performance significantly. We recommend the use of rank-based clustering, coupled with placement based on MI, as the near-optimal choice to determine the number of required sensors, and their respective locations. Our future work will focus on deploying sensor network clusters in the locations proposed in this paper, and will evaluate the efficacy of these networks to estimate SWE in the American River basin.

Chapter 4: A Hybrid System Model of Seasonal Snowpack Water Balance

The measurement and estimation of basin-wide SWE is the first step toward informing water resource operations in mountain basins. These areal estimates of SWE must also be projected into the future to develop insights into the resulting snowpack runoff, snowpack states, and basin-wide streamflow. It is expected that improved data gathered by hydrologic WSNs will open the door to a new suite of physically-based models. One such model is introduced in this chapter.

4.1 Introduction

We presents a direct application of hybrid systems for the modeling of a natural phenomenon, and shows that through the hybrid systems framework it is possible to model extremely complex system dynamics, such as hydrologic process in mountains basins, as set of much simpler switched dynamics. We derive a mathematical model for the dynamics of snow, and show how hybrid systems can be used to blend heuristics and physical laws within a coherent mathematical framework. Our model is motivated by water supply shortages in the western United States, and provides a powerful new tool to aid in critical water management decisions.

The physical snowpack is known to be a highly nonlinear system which, to be modeled properly, requires solutions to nonlinear partial differential equations (PDEs). This often leads to comprehensive, but highly parameterized models with many state variables. We show, however, that it is possible to cast the problem in the framework of hybrid systems, allowing us to describe the evolution of the snowmelt process as a decoupled system of three regimes. We will highlight how a hybrid system model can simplify the complex dynamics of offered by similar PDE approaches without sacrificing model accuracy. Casting our model in this framework brings a new sense of intuition to the modeling process, something which is lost in more complex PDE or heuristic approaches. In the past, successful applications of hybrid systems theory have ranged from studies of bio-molecular networks, control of electronic vehicle management systems, and cognitive radio networks [*Ripaccioli et al.,* 2009; *Alur et al.,* 2001; *Ghosh and Tomlin,* 2001; *Borri et al.,* 2010]. Our model not only brings further mathematical rigor, and new analytical tools to hydrologic sciences, but validates the appropriateness of hybrid systems for the modeling of real-world phenomena.

4.2 The Physics of Snow

4.2.1 Physical Behavior

The physical snowpack is an inherently unstable system, continuously subject to fluxes in mass, energy, and momentum. Mass and energy are introduced into the system in the form of water, in either a solid (snow), liquid (rain), or vapor form. Additionally, energy is exchanged with the snowpack through evaporation, soil heat flux, as well as solar and thermal radiation. As energy is introduced onto the snowpack, the surface melts, and the subsequently produced liquid water, compelled by gravitational forces, begins to flow through the snow matrix. Depending on

the temperature and physical characteristics of the snow below the melt surface, this water may either refreeze onto the snow grains below, or flow into or onto the underlying soil surface. As water melts on the surface and moves into the underlying snow matrix, the latent heat generated by the refreezing of the water is used to heat the snow below the melt surface. It is thus possible to classify the bulk equations governing snowmelt using standard conservation laws, along with constitutive relations for water flux through porous media [*Bear*, 1988; *Turns*, 2006].

At a basic level, snowmelt is thus governed by a coupled interaction between energy and mass balance components:

$$\frac{dE}{dt} = f(M(t), E(t)) \tag{1}$$

$$\frac{dM}{dt} = g(M(t), E(t))) \tag{2}$$

where M and E reflect specific mass[1] and energy components. It is also customary to introduce a momentum conservation relation, which tracks the movement of water and deformation of the snow matrix. Since our main goal is to model snow water equivalent, rather than track front propagation within the snow matrix, we take the approach of Marks et al. (1998), and model compaction in a prescribed, rather than explicit, manner. An energy balance governs the flux of energy into the system, usually in the form of solar and thermal radiation, latent heat transfer, and soil and evaporative fluxes. General conduction theory along with the heat equation can be used to model the temperature evolution through the snowpack. Realistically, temperature distribution in the snowpack is non-uniform, but due to the inability to sense these distributions, most authors assume an isothermal snowpack (e.g., [*Colbeck*, 1972; *Colbeck*, 1980]).

Water flow through snow is a non-linear process, complicated not only by microscopic characteristics of the snow pack, but also by the tendency of the water to refreeze as it flows through the snow matrix. Water flow through porous media, described in detail by [*Bear*, 1988], can be broadly classified into two regimes: saturated flow, and unsaturated flow. Water flow near the surface of a melting snowpack, where a layer of wet snow exists, is governed by Darcy's law [*Bear*, 1988]. The flow is proportional to the gravitational pressure gradient, as well as microscopic properties of the snow grains. Once water leaves this region, flow in the dry region can be modeled by the nonlinear Richard's PDE [*Bear*, 1988], which is highly dependent on microscopic properties, and for which a closed form solution does not exist. It can thus be seen how the coupling of all these interactions has the potential to make the proper modeling of snowpack a non-trivial process. Prior to engaging the inner workings our proposed model, it will be advantageous to discuss previous efforts from the snowmelt modeling literature.

4.2.2 Previous Work

A number of authors have explored the empirical and mathematical modeling of snowmelt processes. A series of notable works was published by Colbeck, in which the author explored both the thermodynamic phenomena guiding snowmelt, as well as water flow through

[1] Mass per unit area. It should be noted that from here on, any references or calculations involving mass are to be taken as references to

the snow matrix [*Colbeck*, 1971; *Colbeck*, 1980; *Colbeck*, 1978; *Colbeck*, 1972; *Colbeck*, 1975]. In these papers, snow was treated as a porous medium consisting of ice, water and air. A mass balance approach along with a model for water flow through porous media, was used to arrive at a PDE that guides the evolution of the pack, and whose solution depended upon a finite difference approach. Various other authors extended this work by realizing that snow is a medium in which the ice matrix and the water phases are, depending in the energy input, interchangeable. Sellers showed, by building upon [*Kelley et al.,* 1986; *Morland et al.,* 1990], that snow can be modeled as a phase-changing porous medium, where the water resulting from melt, is interconvertible with the rigid ice matrix [*Sellers*, 2000; *Sellers*, 2001]. Additionally, the author developed a method by which to track the propagation of the melting snow surface and the percolating interior boundary, by noting that the problem can be formulated as a generalized Stefan condition (or, a free moving boundary) [*Crank*, 1984]. Tseng et al., guided by a model for water flow through subfreezing snow by Illangasekare et al. (1990), arrived at a similar solution, but formulated the problem using volumetric supply terms [*Tseng and Illangasekare,* 1994]. More recent work on the subject was conducted in a comprehensive three part paper by Bartelt and Lehing, in which the authors modeled all the phenomena described in the previous section by making simplifying assumptions about snow micro-structure. The complexity of their extensive model, however, warranted a solution only by a finite element approach [*Bartelt and Lehning,* 2002]. Currently, two popular point based modeling packages exist: SNOBAL, and SNTHERM [*Marks et al.,* 1998; *Jordan*, 1990]. The former is a two layer heuristic model that tracks mass and energy fluxes well, while the latter is a highly parameterized 20 state PDE model that tracks the compaction and formation of layers within the snow-matrix in detail. Our approach, however, offers an elegant hybrid solution, preserving the motivating heuristics, while retaining mathematical rigor.

4.3 Modeling the Snowpack

The complexity of the previously described PDE approaches has the tendency to remove general intuition regarding the physics of the snowmelt process. The procedures usually entail lengthy mathematical derivations along with a high degree of parameters, arriving at a solution that is difficult to intuitively comprehend. Additionally, many of the inputs into these models are not readily available (e.g. snow grain diameter), and must be estimated. Our proposed model addresses these concerns by splitting the snowmelt process into stages, each of which is governed by relatively simple dynamics. The evolution of the entire pack in time can then be modeled as a switched system between these stages. A number of reasonable assumptions will justify simplified dynamics in each of the three proposed regimes. The snowpack is at all times assumed to have a uniform bulk temperature T (see section 4.4.1 for details). Furthermore, we assume that the snowpack can be modeled as a single-layer homogeneous mixture of ice, water and air with bulk density ρ_S. Most authors have taken a multilayer approach to model melt over long periods of time [*Marks et al.,* 1998; *Jordan*, 1990], but since this model aims to forecast weekly SWE values rather than track long-term structural metamorphosis, a single layer was deemed to be adequate. Model validation will show this assumption to be appropriate. The following sections will decouple the snowmelt process into three regimes, and will derive the necessary equations to describe the dynamics of the snowpack in each regime.

4.3.1 The Sub-freezing Snowpack

During this stage no liquid water is present within the snow matrix. The snowpack is completely frozen and the volume occupied by the system is strictly filled with ice and air. We denote E as the energy introduced into the system. A positive energy input will heat the surface of the pack until it reaches a temperature of $0^o C$, upon which the ice at the surface will begin to melt and percolate into the snow matrix. The water will then refreeze, and the latent heat released by refreezing will in turn raise the temperature of the surrounding snow. This process steadily progresses until the water front eventually reaches the bottom of the snow pack. At this point, the entire snowpack is assumed to be at a bulk temperature $T = 0^o C$. Since the primary purpose of our analysis is to model the water content of the snowpack, it is not necessary to track the evolution of this interior saturation boundary, but rather to identify the point at which the snowpack reaches a bulk temperature of $T = 0^o C$. Following the approach of [*Illangasekare et al.*, 1990] and using general conduction theory [*Turns*, 2006], it is possible to identify the point in time at which the snowpack will reach an isothermal state at $T = 0^o C$. The relation giving the change in temperature of the snowpack during this regime can be given by

$$\frac{dE}{dt} = M_s C_s \frac{dT}{dt} \tag{3}$$

where E is the energy input into the system, M_s is the mass of snow, and C_s is the specific heat of snow. We take a proscribed approach to snow pack metamorphosis, merging approaches of [*Marks et al.,* 1998; *Martinec and Rango*, 1991] and note that during this stage of snowmelt, change in snowpack height is governed by an exponential decay model. Analysis of field data shows this to be a reasonable assumption.

The change in density of the snowpack is given by:

$$\frac{d\rho_s}{dt} = \frac{A}{B\left(1 + \dfrac{\rho_s(t)}{A - \rho_s(t)}\right)^2} \tag{4}$$

where ρ_s is the bulk density of the snowpack, and A and B are easily obtained model parameters. We derive (4) in Appendix 4.8.1. When the bulk temperature of the snowpack reaches $T = 0^o C$, the system jumps into the second regime of snowmelt described next.

4.3.2 The Isothermal Sub-saturated Snowpack

At this stage the snowpack has reached an isothermal state, where the bulk temperature is $0^o C$ throughout. It is physically impossible for the bulk temperature to climb beyond this value. Therefore, any energy input into the system at this point is used to convert (melt) snow at $0^o C$ to water at $0^o C$. Two more equations can now be added to reflect that ice is being converted to water at a rate proportional to the energy input:

$$L_f \frac{dM_i}{dt} = -\frac{dE}{dt} \tag{5}$$

$$L_f \frac{dM_w}{dt} = \frac{dE}{dt} \tag{6}$$

where M_w and M_i are the mass constituents of ice and water of the pack, and the mass of snow is given by $M_s = M_i + M_w$. L_f is the latent heat of fusion of ice. Water is still not exiting the snowpack because it is held in place by capillary action within the ice matrix [*Bartelt and Lehning,* 2002]. The volumetric water content of snow θ_w, the volume of water over the total volume of snow, is given by

$$\theta_w = \frac{V_w}{V_{total}} = \frac{M_w/\rho_w}{M_s/\rho_s} \tag{7}$$

where ρ_w and ρ_s are the density of water, and snow, respectively. θ_w has to climb above the threshold value θ_r (typically 1%) for water to begin flowing out of the snowpack [*Marks and Dozier,* 1992]. Thus, in our model the snow pack begins to drain once the volumetric water content condition has been satisfied, or

$$\theta_w \geq \theta_r \tag{8}$$

or, equivalently, when

$$\frac{M_w/\rho_w}{(M_i + M_w)/\rho_s} \geq \theta_r. \tag{9}$$

4.3.3 The Isothermal Saturated Snowpack

At this stage, the volumetric water content has reached the critical threshold value, and water can begin to flow out of the snowpack. Similarly, if the volumetric water content drops below θ_r the system relapses back into the second regime. Snowpack density and change in the ice mass are still evolving according to the relations in the previous regime. The amount of water that can now leave the system is equal to the amount of ice being melted along with the amount being discharged due to any excess volumetric water content. The relation giving the change in water content of the snowpack is given by

$$(\rho_s - \rho_w)^2 \frac{dM_w}{dt}$$

$$= -\theta_r \rho_w \left(\frac{1}{L_f} \frac{dE}{dt} (\rho_s - \theta_r \rho_w) + M_i \frac{d\rho_s}{dt} \right). \quad (10)$$

We show in Appendix 4.8.2 that eq. (10) can easily be derived using the volumetric water requirement of regime three. Intuitively, the above relation accounts for the fact that as density of the snowpack is changing, the volumetric water content changes accordingly, and water must leave the system to ensure $\theta_w = \theta_r$.

4.4 Hybrid Systems Modeling

The previous section showed that is it possible to separate the natural process of snowmelt into three regimes, each reflecting physical states of the snowpack. As such, a hybrid systems model of snowmelt is an appropriate means by which to formalize the interaction between the above dynamics. Formulating the problem in this framework will allow for a model whose solution is dependent on a series of switched ordinary differential equations (ODEs). Our hybrid model captures the behaviors of more complex models, while ensuring that a sense of intuition regarding the snowmelt process is preserved. More information on hybrid systems can be found in [*Tomlin et al.,* 2000].

4.4.1 Preliminaries

We base our definition of hybrid systems on a class of hybrid systems given in [*Tomlin et al.,* 2000]. Definition:

An autonomous hybrid automaton H is a collection

$$H = (Q, X, Init, U, f, Dom, R, Y)$$

where

$Q = \{q_1, q_2, ...\}$ is a set of discrete states;

$X = \mathbb{R}^n$ is the continuous state space;

$Init \subseteq Q \times X$ is a set of initial states;

U is a set of continuous input variables;

$f : Q \times X \times U \to \mathbb{R}^n$ is a vector field;

$Dom : Q \to 2^{X \times U}$ is a domain;

$R : Q \times X \times U \to 2^{Q \times X}$ is a reset relation;

Y is a collection of continuous output variables.

A hybrid time set is a finite or infinite sequence of intervals $\tau = \{I_i\}_{i=0}^{N}$ such that $I_i = [\tau_i, \tau_i']$ for all $i < N$; and $\tau_i \le \tau_i' = \tau_{i+1}$ $\forall i$.

if $N < \infty$ then either $I_N = [\tau_N, \tau_N']$ or $I_N = [\tau_N, \tau_N')$;

A hybrid trajectory (τ, q, x) consists of a hybrid time set $\tau = \{I_i\}_0^N$ and two sequences of functions $q = \{q_i(\cdot)\}_0^N$ and $x = \{x_i(\cdot)\}_0^N$ with $q_i(\cdot) : I_i \to Q$ and $x(\cdot) I_i \to \mathbb{R}^n$.

An execution of a hybrid automaton H is a hybrid trajectory, (τ, q, x) which satisfies:

Initial Condition: $(q_0, x_0) \in Init$;

Discrete Evolution: $(q_{i+1}(\tau_{i+1}), x_{i+1}(\tau_{i+1}) \in R(q_i(\tau_i'), x(\tau_i'))$;

Continuous evolution: for all i,

1. $q_i(\cdot) : I_i \rightarrow Q$ is constant over $t \in I_i$, that is, $q_i(t) = q_i(\tau_i)$ $\forall t \in I_i$;

2. $x_i(\cdot) : I_i \rightarrow X$ is the solution to the differential equation $\dot{x}_i = f(q_i(t), x_i(t), u(t))$ over I_i, starting at $x_i(\tau_i)$; and,

3. $\forall t \in [\tau_i, \tau_i')$, $x_i(t) \in Dom(q_i(t))$.

4.4.2 A Hybrid System Model of Snow

Using the above definition, it is now possible to model the physical snowmelt process as a hybrid system:

- $Q = \{q_1, q_2, q_3\}$ is the set of discrete states, corresponding respectively to each regime of snowmelt.

- $X = \mathsf{R}^4$ is the set of continuous states, where for $x \in X$ we have

$$x = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} M_i \\ M_w \\ \rho_s \\ T \end{bmatrix} \tag{11}$$

where x_1 and x_2 are the mass of ice and water in the snowpack, x_3 is the density of the snowpack, and x_4 is the bulk temperature of the snowpack.

- The set of initial states

$$Init = Q \times \{x \in X : x_1 \geq 0 \wedge 0 \leq x_2 \leq \theta_r \rho_w \frac{x_1 + x_2}{x_3} \wedge 0 \leq x_3 \leq A \wedge x_4 \leq 0\}$$

where A is the same as in (4).

- The domains are given by

$$Dom(q_1) = \{x \in \mathsf{R}^4 \mid x_4 < 0\}$$

$$Dom(q_2) = \{x \in \mathsf{R}^4 \mid x_4 \geq 0 \wedge 0 < x_2 < \theta_r \rho_w \frac{x_1 + x_2}{x_3}\}$$

$$Dom(q_3) = \{x \in \mathsf{R}^4 \mid x_2 \geq \theta_r \rho_w \frac{x_1 + x_2}{x_3}\}$$

- The system dynamics for each state are given by

$$f_1(t,x(t),u(t)) = \begin{bmatrix} 0 \\ 0 \\ h(x_3(t)) \\ \frac{1}{C_s}\frac{u(t)}{x_1(t)} \end{bmatrix} \tag{12}$$

$$f_2(t,x(t),u(t)) = \begin{bmatrix} -\frac{1}{L_f}u(t) \\ \frac{1}{L_f}u(t) \\ h(x_3(t)) \\ 0 \end{bmatrix} \tag{13}$$

$$f_3(t,x(t),u(t)) = \begin{bmatrix} -\frac{1}{L_f}u(t) \\ g(x_1(t),x_3(t),u(t)) \\ h(x_3(t)) \\ 0 \end{bmatrix} \tag{14}$$

where

$$g(v,w,z) = \frac{-\theta_r\rho_w\left(\frac{1}{L_f}z\left(w-\theta_r\rho_w\right)+vh(w)\right)}{\left(w-\rho_w\right)^2} \tag{15}$$

and

$$h(z) = \frac{A}{B\left(1+\frac{z}{A-z}\right)^2} \tag{16}$$

and all other terms are constants as defined previously. The continuous input $u(t) \in U \in \mathsf{R}$ represents the amount of energy input into the system (given in W/m^2). The continuous output of the hybrid automaton is, at all times, given by $y(t) = (x_1(t) + x_2(t))/\rho_w$, reflecting the SWE of the snowpack.

$R(q_1,x) = (q_2,x)$ if $x_4 \geq 0$

$$R(q_2, x) = (q_1, x) \text{ if } x_2 \le 0$$

$$R(q_2, x) = (q_3, x) \text{ if } \frac{x_2/\rho_w}{(x_1 + x_2)/x_3} \ge \theta_r$$

$$R(q_3, x) = (q_2, x) \text{ if } \frac{x_2/\rho_w}{(x_1 + x_2)/x_3} < \theta_r$$

$$R(q, x) = \varnothing \text{ otherwise.}$$

A graphical representation of the hybrid automaton is given in Figure 4.1.

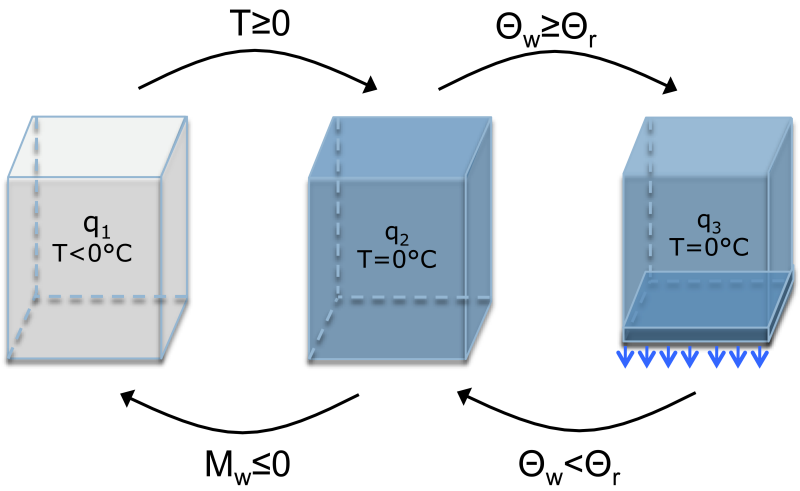

Figure 4.1 A visual representation of the snowmelt hybrid automaton. State q_1 is the first stage of snowmelt, where the entire snowpack exists at below freezing temperatures, and no liquid water is present. State q_2 represents the isothermal snowpack at $T = 0^o C$. Ice begins to melt, but water can not leave the snow matrix because it is held in place by capillary forces. Once the water content θ_w of the snowpack reaches a value greater than or equal to a threshold θ_r, water begins to exit the system and flow into the soil below. A negative energy input in q_2 causes existing water to refreeze. If all the water in state q_2 freezes, the system relapses into the sub-freezing state q_1.

4.5 Model Validation

Due to the relatively straightforward nature of the discrete mode dynamics, a detailed analysis of the system would not offer more information than is provided by a qualitative overview:

State q_1: The steady state behavior of the hybrid automaton during this discrete mode reflects only densification of the snow pack, as an energy input is required for the temperature to be affected. During this state, the density of the snowpack, x_3, is evolving towards a terminal density value A, given in (4). Introduction of an energy flux $|u(t)| > 0$ evolves the temperature of the snowpack proportional to the sign and magnitude of the energy input.

State q_2: The evolution of x_1 and x_2 is equal in magnitude, and opposite in sign, and is proportional to the energy input into the system. The continuous state x_4, the temperature of the snowpack, does not evolve and remains at a steady $T = 0^o C$, since ice and snow can not exist in solid form beyond this region.

State q_3: With a steady positive energy flux, the snowpack will eventually melt and the dynamics of state q_3 will ensure that the mass of ice and water go to zero. In q_3, the steady state behavior does affect mass balance, since densification of the snowpack x_3, changes the system's volumetric water content. As the snowpack density approaches the terminal value A, the steady state contribution of the x_3 dynamics to overall mass balance becomes negligible. Given a positive energy influx, the dynamics of x_2 evolve to ensure that the volumetric water content of the snowpack θ_w is at all times equal to the threshold value θ_r.

4.5.1 Initial Conditions and Parameterization

Values, such as the amount of liquid water, and the temperature of the snowpack, are difficult to measure in the field. Since these values are embedded in our state space, it becomes important to develop a proper approach to confidently estimate their initial conditions. The hybrid automaton was initialized with night-time values following a fresh snow-fall event. This maximized the likelihood of capturing the system during state q_1, as we assume the snowpack temperatures to be below freezing during such events. We believe that this method allows us to estimate initial states confidently, by assuming that at this stage, the temperature of the snowpack is equal to the air temperature, and the mass of liquid water in the system is zero. The density was initialized using physical measurements, and the mass of ice was calculated using density and snow height measurements. The free parameters, or physical constants, of the model are given in Appendix 4.8.3, and are based on commonly available physical constants, as well as known properties of snow in the Sierra Nevada. The model shows no significant sensitivity to initial conditions and parameters when they are constrained to physically realistic values.

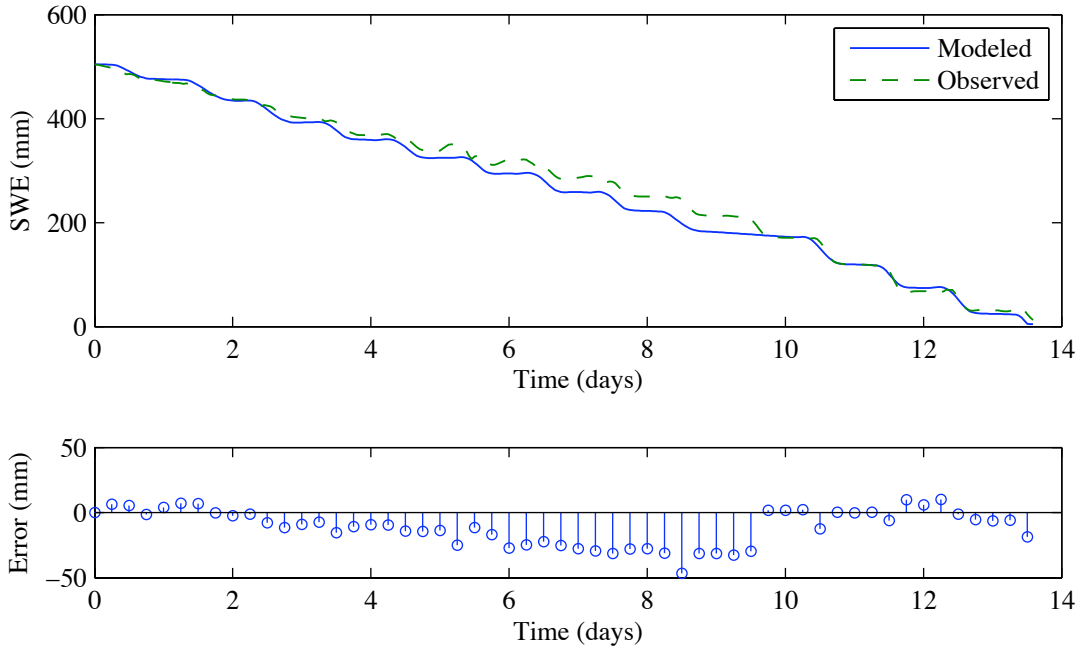

a) Observed SWE vs. Modeled SWE for a site in the Sierra Nevada during the 2004 snow season.

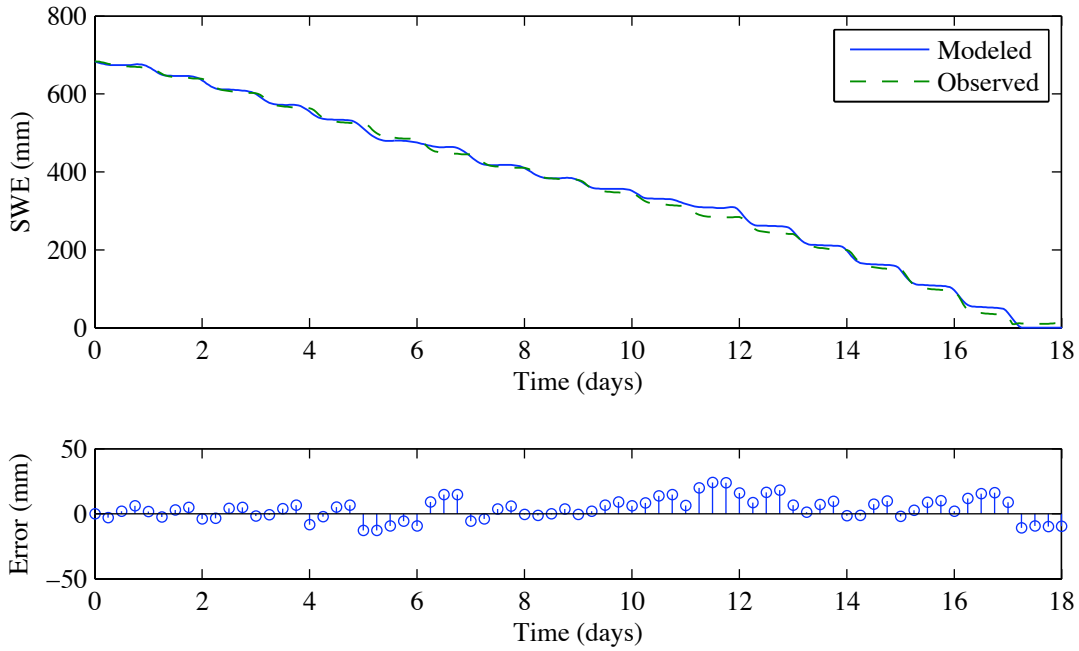

b) Observed SWE vs. Modeled SWE for a different site in the Sierra Nevada during the 2006 snow season.

Figure 4.2 Plots of Snow Water Equivalent (SWE) at two different sites during 2004 and 2006. Snow water equivalent is the height of water that would result if the whole snowpack were to melt instantaneously. Step-like features in the time series are indications of discrete mode switches of the hybrid automaton.

4.5.1 Validation Results and Analysis

By inspection of Figure 4.2, very close correspondence between modeled and observed SWE values becomes evident. Maximum error for a two week prediction of SWE, for both years, did not exceed 50mm of SWE. A measurement error in the range of 20mm is known to exist for the snow-depth sensor that was deployed in the study. Given the week long time horizon requirement imposed on this model, we regard these results to be indicative of very good model performance. The accuracy of the model is further validated by the fact that it is able to correctly identify values and trends in SWE for two different sites, and snow seasons. The point at which complete melt occurs is identified correctly as well. Overall, the model performs extremely well within the desired seven day prediction window, and is able to capture high-frequency daily variations in SWE. Figure 4.2a shows a slight model undershoot at the five day mark. At this time the observed SWE value displays an actual increase, indicating a snow fall event has occurred. This is an acceptable behavior, since our main aim is to forecast SWE given only current snowpack conditions. It is assumed that during the operational phase, the model will be used as a moving window to incorporate changes in incoming meteorological forecasts. The step like response of the SWE plot in Figure 4.2, indicates switching of the discrete system dynamics, and further intuitively validates the switched dynamics approach of our model.

4.5.2 Comparison to Existing Simulators

To fully identify the advantages of the proposed approach, a rigorous quantitative analysis and comparison with existing simulators is warranted. Given the currently available validation data, as well as the high dimensionality and parameterization requirements of both the SNOBAL [*Marks et al.,* 1998] and SNTHERM [*Jordan,* 1990] models, such an analysis was deemed beyond the scope of this particular paper. A series of test runs using SNOBAL found the model predictions of SWE to be inconsistent with the ground truth data. SNOBAL performed along the same error margin as the hybrid automaton simulation for the first four days of the melt cycle, only to significantly diverge through the remainder of the simulation period, predicting total melt with an error of seven days. This behavior is largely attributed to the model's apparent sensitivity to parameterization and snow layer initialization. It should also be noted that the source code of the SNOBAL model had to be modified to allow for inputs of solar radiation that were physically measured, but deemed unreasonable in the original model design. The SNTHERM model is, to our knowledge, one of the most comprehensive snowmelt simulator available. As such, it requires more inputs and parameters that were available to this study. Assumptions regarding these parameters could not be justified given the available data sets. A data set currently being collected in the Sierra Nevada, however, will permit for a more comprehensive error analysis of the proposed hybrid system model, along with providing confident estimates of model parameters for SNTHERM and SNOBAL.

Compared to existing simulators, computational tractability along with model analyzability stand out as the two main benefits gained by casting the snowmelt process in a hybrid systems framework. As mentioned previously, the primary aim of the proposed model is to reliably forecast weekly SWE given real-time data acquired by densely deployed wireless sensor networks. Spatial variation in snow cover will be accounted for by the distribution of

network nodes, and separate snowmelt simulations will need to be carried out for each measurement location. The relatively simple structure of the proposed model allows it to run extremely fast on a standard personal computer, and offers promise in terms of scalability. Compared to existing simulators, the proposed model also lends itself to a number of potential hybrid systems analysis techniques. The model originated with the aspiration to control dams in the California mountains based on the input of real-time meteorological and snow cover data. Casting the snowmelt process in hybrid systems framework will allow for investigation of controller synthesis methods for such purposes. Additionally, hybrid system estimation, and reachability analyses can be carried out to provide water management officials with valuable information regarding water storage and snowmelt runoff.

4.6 Conclusions and Future Work

In this paper, we derived a hybrid system model of the seasonal snowmelt water balance. This model is part of a larger project scope to control water management infrastructure and forecast available water supplies in the state of California. The main contribution of our paper hinges not only on its performance, but rather also on showing that it is possible to accurately model an extremely complex natural phenomenon using a system of simple switched dynamics. Compared to other approaches, our model significantly reduces the number of required state variables while delivering a high level of intuition to the modeling process. The approach also makes it possible to incorporate well known heuristics into a coherent mathematical framework. Dimensionality reduction is one of the main appealing features of the approach, when compared to traditional PDE models. It is believed that the dimensionality of the system could be further reduced in the future by taking into account that in each discrete mode, certain continuous states experience zero change. As (12), (13) and (14) show, the system could be modeled as a hybrid automaton with switching between state spaces of different dimensions. Future work will explore the choice of a new coordinate system which takes further advantage of known physical behaviors of snowmelt.

From a hydrologic perspective, future work will focus on investigating to which extent the model can accurately predict spatial distribution of SWE. Our current analysis was limited to point-wise forecasting of a monotonically decreasing SWE based on energy input measured at the exact location of the SWE estimate. To broaden the impact to flood and agricultural planning, our approach will be expanded to treat the spatial variation of the snowpack throughout a drainage basin. This approach will entail the interaction of multiple automata, each governing snowmelt in a specific region of the basin. Minor modifications to the current snowmelt automaton will account for thermodynamic interaction between the automata, and will permit water to flow into soil as well as into neighboring automata. New data sets are currently being collected, which will be used to validate the model over long-term freeze-thaw snowmelt cycles. The variability of energy fluxes on the model area, as well as appropriate terrain gridding techniques will be investigated. This model was designed to utilize real-time data provided by wireless sensor networks, which are currently being deployed. Estimation techniques will also be developed to ensure that the incoming input data is a valid representation of actual environmental conditions.

4.7 Derivations

4.7.1 Derivation of Bulk Density Evolution Equation

We are given the empirical observation [23] that the density of aging snow can be determined through the half saturation equation:

$$\rho_s(t) = \frac{A}{1 + B/t} \tag{17}$$

where A is the maximum saturation level (in kg/m^3) and B is the time at which half-saturation is achieved (days). We would like to derive a relation for the change in snow density, given the current density, from this relation. Given a density, we solve (17) for the corresponding time

$$t = \frac{\rho_s B}{A - \rho_s} \tag{18}$$

and then take the derivative with respect to t of (17) to get

$$\frac{d\rho_s}{dt}(t) = \frac{AB}{(B+t)^2} \tag{19}$$

after substituting (18) into (19) we get the desired relation in (4).

4.7.2 Derivation of Change in Water Mass Relation

When the system is in the saturated isothermal regime the mass of water in the snowpack at any time can be given by solving (7) for M_w to yield

$$M_w = \theta_r \rho_w \frac{M_i}{\rho_s - \theta_r \rho_w}. \tag{20}$$

Taking the derivative of (20) with respect to t gives

$$\frac{dM_w}{dt} = \theta_r \rho_w \left[\frac{\frac{dM_i}{dt}(\rho_s - \theta_r \rho_w) - M_i \frac{d\rho_s}{dt}}{(\rho_s - \theta_r \rho_w)^2} \right]. \tag{21}$$

Substituting with the relation in (5) and rearranging gives (10).

4.7.3 Model Parameterization

$C_s = C_{ice} = 2.05 \times 10^3 \, J/kg^o K$

$\theta_r = 0.01$

$L_f = 334 \times 10^3 \, J/kg$

$\rho_w = 1000 kg/m^3$

$A = 450 kg/m^3$

$B = 20$ days

Chapter 5: Impacts of Improved Snow Water Equivalent Estimates on Water Resource Operations in the American River Basin

The previous chapters introduced improvements for the measurement and estimation of snowmelt driven processes in remote mountain basins. While it is expected that these methods will enable improved scientific understanding of the spatiotemporal coupling of catchment-, and basin-scale hydrologic processes, a motivating example has still to showcase the downstream, societal impacts of the proposed methods. A convincing case of the benefits of improved SWE data has to be made to private and public stakeholders if investments into a more advanced snow monitoring infrastructure are to continue

A step towards motivating the future adoption of these techniques beyond the scientific community is to analyze the potential impact of improved snowpack estimates on large-scale water forecasting. This chapter reviews the current state of basin-scale streamflow forecasting in the American West, focusing especially on the American River basin of California (see Chapter Three). A forecast skill analysis will be carried out to evaluate the efficacy of currently used methods over the past decade, and it will be shown that improved snow water equivalent estimates will play a significant role in improving these predictions in the future.

5.1 Introduction

The allocation and use of water resources in the American West is heavily depended upon seasonal water forecasts made by a set of large agencies and smaller groups. Such predictions aim to provide long-term streamflow forecasts based on a combination of various models, streamflow observations, and hydrologic and meteorological datasets [*Hartmann et al.* 2002]. Three closely watched forecasts are currently made by the National Weather Service (NWS), the Natural Resource Conservation Service (NRCS), and the California department of Water resources (DWR) [*California Cooperative Snow Surveys*, 2012; *Rice and Bales*, 2010; *SNOTEL* 2012]. Each of these forecasts combines current and historical streamflow measurements, with data obtained from the set of available snow pillow and snow course measurement (Figure 1.1). The NRCS uses their own network of snow pillow observations, known as SNOTEL, while the other agencies rely on snow course depth measurements and snow pillow data provided by the California Cooperative Snow Surveys (CCSS) [*Bales et al.* 2006]. Depending on the agency, each forecast is derived through methods ranging from simple multivariate regression to complex physical process models. These forecasts are summarized in table 5.1.

Agency	Method	Data	Forecast offered
Natural Resource Conservation Service	Multivariate regression	Precipitation, streamflow, SNOTEL sites	Reservoir storage; Seasonal inflow; 70% and 30% exeedance probabilities
National Weather Service	Physical, process-based model for snow, soil, and flow routing	Precipitation, temperature, streamflow, and CCSS snow pillow sites	Ensemble streamflow predictions; minimum, maximum, and probable flows; exeedance probabilities
California Department of Water Resources	Multivariate regression	CCSS Snowcourses, precipitation, streamflow	Reservoir storage; Seasonal inflow;

Table 5.1 Water forecasting in the Western-US. Observed streamflow inputs for each method is provided by the United States Geological Survey (USGS).

5.1.1 Regression Techniques

Presently, the inputs used in the regression forecasts vary between agencies, but each agency typically agrees on the need to use SWE or snowdepth data, monthly precipitation, historical streamflow, and streamflow leading up to the month of the forecast date [*Garen* 1992; *McFadden* 2011]. There is an inherent assumption of stationarity in such predictions, as future streamflow forecasts are heavily dependent upon historical observations. This, in part, explains why these forecasts tend to have high errors during *wet* and *dry* periods, when observed snowmelt and precipitation are above the historical mean [*Bales et al.* 2006]. Additionally, many of the measurements of SWE and snowdepth are made in easy to access locations, which typically exhibit similar topographic features. As such, the input datasets become correlated, reducing the amount of information available to simple multiple regression models. While not used in official forecasts, *Garen* (1992) suggested an improved method based on *Principal Component Regression* [*Hastie et al*., 2001], with the aim of creating a new set of independent inputs variables that are a linear combination of the original inputs. Using this method, the author showed a notable improvement to forecast predictions, while maintaining month-to-moth forecast consistency. This method is currently being used effectively at a smaller scale by *McFadden* (2011) in forecasting natural runoff and stream accretion at watersheds supplying eight powerhouses in the Upper American River basin.

5.1.2 Physical Model

Due to historically poor performance of regression models in extreme *wet* or *dry* years, the NWS shifted to the use of conceptual hydrologic and hydraulic models to predict the melt of the snowpack, and to route the corresponding runoff downstream [*Day,* 1985]. While a number of distributed, physically-based snowmelt and hydrologic models exist (see Chapter Four) the de-facto standard for large-scale river and streamflow forecasting is the Extended Streamflow Forecast Model used by the NWS [*Curtis and Schaake* 1979]. Parameterized on historical data, and taking the inputs precipitation and temperature, the model facilitates the interaction between

multiple smaller sub-models including soil moisture accounting, snow accumulation and ablation and channel routing. Aside from the final output of streamflow, the model tracks all internals states for analytical purposes.

As a more general framework, the NWS model is composed of two main systems: the calibration system, and the operational forecast system. The former uses historical data to calibrate the model, with the assumption that the physically-based model parameters maintain stationarity across the calibration and modeling intervals. The latter provides the streamflow predictions which are then made available to public stakeholders. A probabilistic framework has also been added to the model, enabling the use of ensemble predictions to provide confidence intervals and exeedance probabilities. One major driving force behind the adaptation of this model was its ability to predict abnormal water years, such as the 1976-1977 California drought [*Curtis and Schaake* 1979].

5.1.3 Performance Analyses

Each forecasting agency conducts annual performance analyses (skill assessments), comparing observed streamflow at a set of major locations to the streamflow predicted by their respective models. Since these reports are only published for major basins, a number of other studies have also conducted more thorough skill assessments. Of major interest are the seasonal forecast, which attempt to predict the total volume of runoff between April 1st and July 31st. *Pagano et al*. (2004) analyzed the data from 4841 historical forecasts made primarily by the NRCS's regression models. The forecasts were compared to USGS streamflow data, using a normalized score, focusing specifically on forecasts made between January 1st – April 1st. The authors noted considerable improvement in forecast accuracy between these two periods (*Nash-Sutcliffe* score from 0.32 to 0.7), especially in regions of California, where January-March precipitation presents 46% of the annual total. The authors also call on cooperation between forecasters, by presenting a historical overview of the conflicts between NRCS (regression) and NWS (physical model), presenting studies by the two groups to show the superiority of each method.

Morrill et al. (2011) conducted a more advanced statistical skill assessment from 136 locations over a five decade span across the Western US. The authors found that below-average water-year flows are often over predicted, while above average flows are under predicted. They show a general improvement in forecast skill over this time period, but less that the 10% reported by previous studies. Additionally, the accuracy of each forecast was heavily dependent on region, with Arizona having an largest error of 55% on the April 1st forecast, and Wyoming the lowest with 20%. *Schaake and Peck* (2005) noted that by subdividing various sources of error (climate, data, and model), it is possible to improve overall forecast error by up to 30%. A long-standing attitude on improving regression forecast error, however, notes that large improvement in their accuracy could not be expected through any further refinement of the technique [*Shafer and Huddleston*, 1984].

Forecast verification in operational hydrology has been limited to date. This is especially true for the NWS physically-based model, making it difficult to gauge its performance over the past few decades. Given the complexity of this model compared to regression techniques, *Demargne et al.* (2009) lay out a framework for conducting skill analyses within the NWS model, noting specifically that the system needs to evaluate the quality of inputs and outputs at each step of the forecast process.

Franz et al. (2003) conducted an analysis of the NWS's physically-based model, generating forecasts for 40 years on the Colorado River Basin, which was sub-divided into an upper and lower basin. Results showed that forecast accuracy improved in later portions of the season. The upper basin forecast performed better. The authors attributed this performance to smaller hydrologic variability at higher elevations. The authors also underlined the need for further investigation to better understand the relative importance of the snowpack and its initial states on the forecast accuracy. It has also been shown that the snowmelt components of the NWS model is extremely sensitive to model inputs of precipitation and temperature, at times reaching errors of 70-80% [*Shamir and Georgakakos,* 2005]. *He at al.* (2012) showed that incorporating real-time SWE measurements can improve the performance of the NWS model, but theses measurements were still estimated from the sparse set of available snow pillow sites.

5.1.4 Impact of SWE Uncertainty

While the above studies have carried out skill assessments of the two forecasting frameworks, there has been considerably little emphasis on the value of spatially representative and more accurate snowpack data on the overall streamflow forecasts. Considering that the majority of streamflow in the state of California results from the snowpack, such an analysis is still needed. This chapter tries to fill that gap by evaluating the explicit benefit of improved snowpack data on reducing overall forecast error and uncertainty. An anticipated improvement to SWE estimates is expected to arise as a result of the implementation of the methods described in previous chapters.

The reconstructed SWE dataset from Chapter Three will be used as a *ground truth* estimate of SWE in the study area. In particular, we aim to answer four questions. Can forecast error of regression techniques and the NWS's physically-based model be improved through more accurate, spatially distributed SWE data? How does snowmelt modeled by the physically-based model compare to *ground truth* snowmelt? What is the sensitivity of the physically-based forecast model to errors in the SWE estimate? How does the physically-based forecast change when modeled snowmelt is replaced with observed, *ground truth* snowmelt? The study area and methods will be discussed and a skill assessment of previous forecasts for will be carried out.

5.2 Study Area

The North Fork of the American River comprises nearly one fifth of the total area of the larger American River basin. The American River basin was also the focus of the near-optimal sensor placement study in Chapter Three. As such, this chapter is a direct continuation of that study, focusing instead, however, on the impacts the improved SWE estimates on streamflow forecasting. The North Fork sub-basin was selected for analysis for two primary reasons: 1) there are no upstream reservoirs, thus allowing the analysis to focus on natural, un-buffered flow only, and 2) the operational parameters of the NWS forecast model for this basin were provided to us directly from its operators [*Hartman R.,* personal correspondence].

The North Fork of the American River basin is located in the northwestern Sierra Nevada of California (Figure 5.1). It comprises approximately 800 square kilometers, and spans an elevation gradient of 200 m to 2800 m. Beyond an approximate rain-snow elevation of 1500m, most precipitations falls in the wintertime as snow. The North Fork basin drains into the larger American River basin, whose topography is described in Chapter Three.

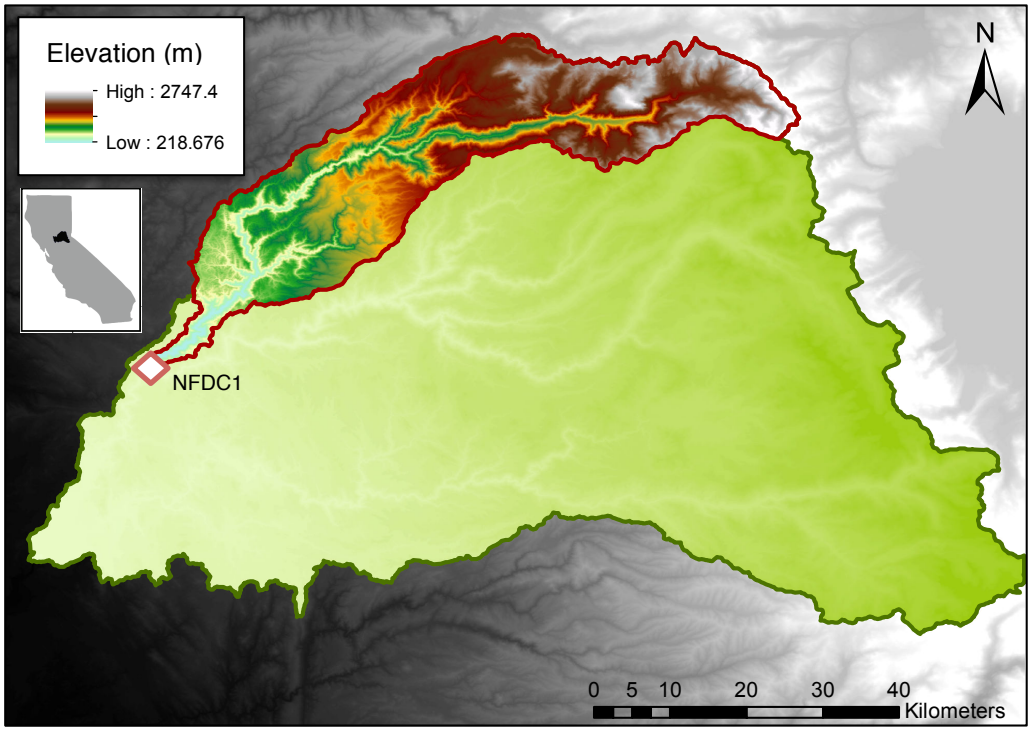

Figure 5.1 Study location. The north fork of the American River in the Sierra Nevada of California. The North Fork is delineated as a sub-basin of the larger American River basin (green).

5.3 Data

The spatial SWE dataset used in this analysis was derived from the pixel-by-pixel historical energy-balance reconstruction of snowmelt described in Chapter Three [*Rittger et al.,* 2011; and *Rittger,* 2012]. After careful consideration, it was decided that this spatially distributed data set offered the best independent ground truth of SWE across years 2000 to 2011. Using the *ArcGIS* software package, the data set was delineated for the North Fork basin to obtain daily values of snowmelt for 500m x 500m pixels. This historically reconstructed data set was treated for rest of the analysis as the *ground truth* snowmelt, assuming that such data will be available in real-time in the coming years, as more WSNs are deployed across the American River.

The drainage of the basin is instrumented by a USGS stream gauge (NFDC1, see Figure 5.1), which provided daily streamflow data over the past decade. Five decades of mean areal precipitation (MAP), and mean area temperature (MAT) were also obtained through personal corresponded with the NWS.

For purposes of evaluating the effects of improved snowpack data on regression methods, data from over thirty snow courses, made available by the CCSS, was assembled (Figure 5.2). Given the sparsity of these measurements in the North Fork basin, it was necessary to collect snow course data across the entire American River basin. For comparison with observed

streamflow, the past decade of official predictions of the DWR regression forecast for the North Fork were obtained from the department's website[2].

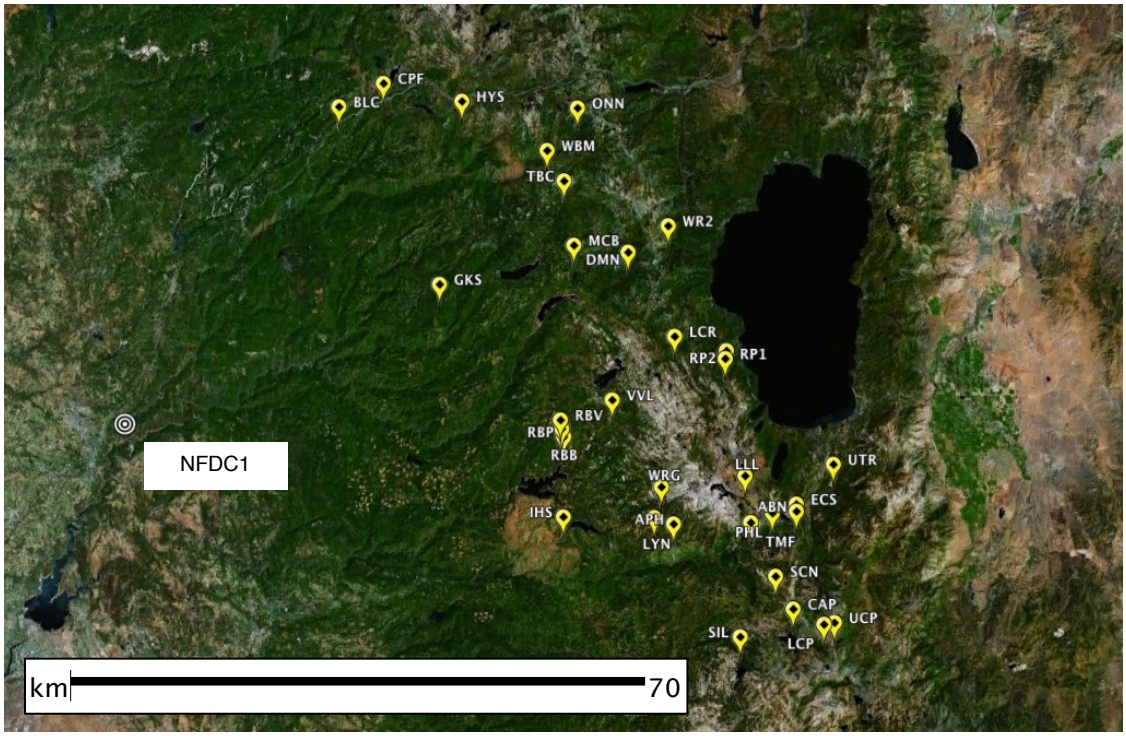

Figure 5.2 Snow course locations used to obtain snowdepth data. The drainage of the North Fork basin is circled (NFDC1).

5.4 Methods

An initial assessment of the forecast skill in the basin was carried out by comparing predicted values made by the DWR's regression forecast and by the NWS's physically-based model, to those observed by the USGS streamflow gauge at the basin drainage point (see NFDC Figure 5.1). While historical DWR forecasts were available, there was no historical records of the NWS forecasts. The predictions of the NWS forecast had to be explicitly modeled for the period of 2001-2011.

5.4.1 Principal Component Regression

The DWR regression forecasts rely on historical observations of stream flow (output), and corresponding historical observations of snow courses and precipitation (inputs). The problem can be formulated as:

$$y_i = \beta_0 + \sum \beta_j x_{ij} \qquad (1)$$

[2] Available at http://www.water.ca.gov/

where y is a vector of streamflow observations, X is the corresponding matrix of input variables, and β is the vector of model parameters. To make predictions, the model parameters β must be learned. In the case of most current water forecasts, this is done by solving the *Least Squares* problem [*Hastie et al*., 2001]. Problems may arise in the above formulation if portions of the input matrix X are highly correlated, leading to poor predictions, or in some cases making it impossible to carry out the matrix inversion that is necessary to solve for the model parameters. Highly collinear input variables are very evident in the snowcourse locations of the American River basin. Figure 5.3 plots the snowdepth observed at a number of selected snowcourses from Figure 5.2, showing highly correlated measurements.

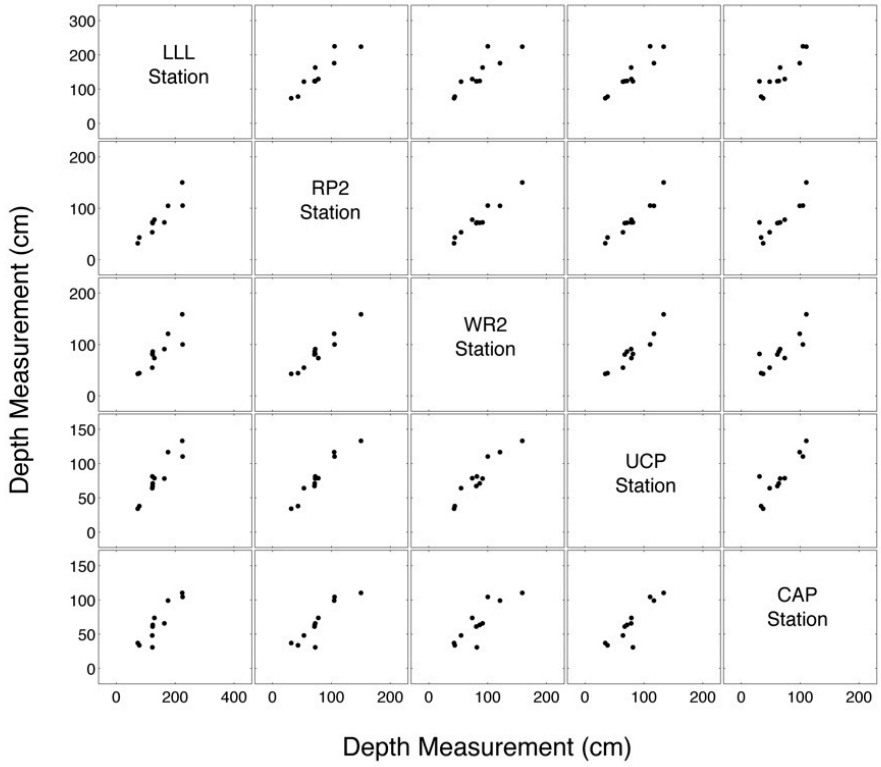

Figure 5.3 Correlation between the SWE measurements for five given stations.

To overcome this high correlation in the snow course data set, Garen (1992) proposed the use of principal component regression (PCR). Based on Principal Component Analysis (PCA), the method transforms the input data into an independent set of principal components, which are then used in the regression. The first step of this procedure entails carrying out PCA on the input matrix X, effectively projecting the inputs into a lower dimensional state, such that:

$$x_i = \sum_{j=1}^{p} \alpha_{ij} \xi_j \qquad (2)$$

where ξ_j is the principle component of x. The regression problem is then reformulated as:

$$y_i = \beta_0' + \sum_{j=1}^{p'} \beta_j' \alpha_{ij} \qquad (3)$$

where $p' < p$ and is chosen based on an analysis of the loading factors. In the case of the snow course inputs, this extra step in the regression will have the effect of maximizing the variance of the input data to improve the overall predictions. For more information on PCA and PCR please see *Hastie et al*., (2001).

5.4.2 NWS Forecast Model

The two primary components of the physically-based NWS model [*Day,* 1985] are the snow accumulation and ablation model (SNOW-17) [*Anderson*, 1973], and the Sacramento Soil Moisture Accounting and runoff model (SAC-SMA) [*Burnash et al.,* 1973]. By coupling the two models it is possible to develop estimates of SWE dynamics in hydrologic basins, and to project the corresponding melt patterns into the runoff model for the derivation of streamflow at the basin discharge point. Note that the Hybrid System snowmelt model that was developed in the previous chapter could also be coupled with SAC-SMA. The purpose of this chapter, however, is to develop insights into operational hydrologic procedures, thus warranting the use of SNOW-17.

Reflecting the NWS's operational use of the model, our model subdivides the larger North Fork basin into an upper and lower basin (Figure 5.4). The delineation occurs at around a 1500m elevation, which is the approximated annual rain-snow-transition elevation. The upper basin comprises 38-39% of the total basin area, and historically receives most of its precipitation as snow.

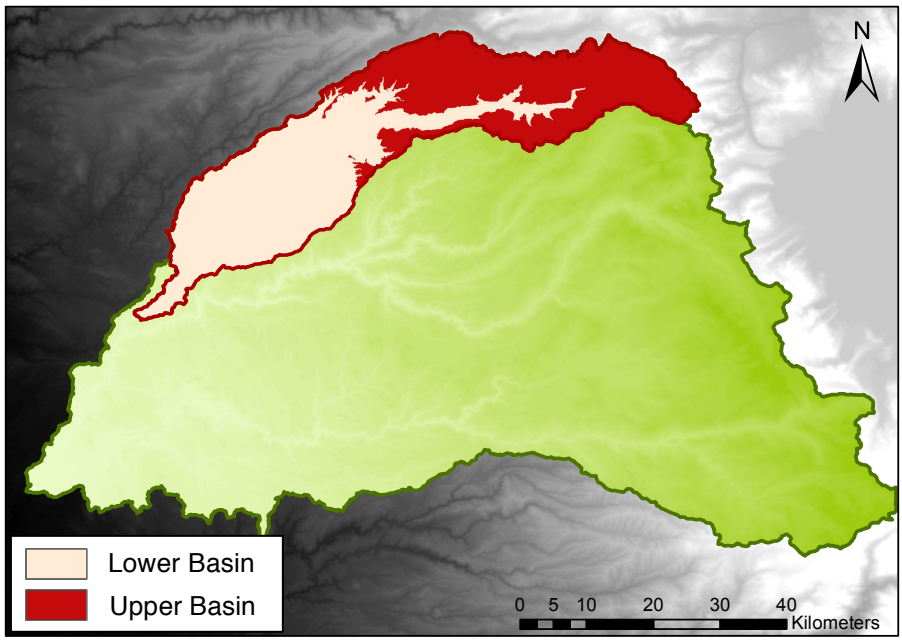

Figure 5.4 Delineations of the upper and lower basins used in the implementation of the NWS model.

5.4.3 Snow-17 Component

SNOW-17 is an areal model of snow deposition and snow melt. The model takes as its inputs the mean areal precipitation and temperature. Based on the areal temperature, precipitation is either

accumulated as snow, or partitioned directly into runoff. Energy exchange at the snow-air interface, liquid water storage, and ground melt are governed by first order relations, similar to those seen in Chapter Four of this thesis.

One major difference between SNOW-17 and the Hybrid System snowmelt model is that SNOW-17 models accumulation explicitly using a parameterized rain-snow transition elevation, which coincides with the average temperature that separates rain from snow. The density of the new snow p_n $(gm \cdot cm^3)$ is a function of air temperature and is given by the following relation:

$$p_n = \begin{cases} 0.05 & T < -15^{\circ}C \\ 0.05 + 0.00017 \cdot T^{1.5} & o.w. \end{cases} \qquad (4)$$

where T is the ambient temperature during snowfall. More information on the SNOW-17 dynamics is given in *Anderson et al.* (1973).

5.4.4 SAC-SMA Component

The Sacramento Soil Moisture Accounting Model is a lumped soil moisture accounting model that uses mean areal precipitation and temperature as inputs. It can also be provided with a number of optional inputs, such as snow covered area and evaporation. The model subdivides the incoming precipitation into impervious runoff, and pervious runoff, which is further subdivided into an upper (soil) and lower zone (long-term soil moisture and groundwater) to calculate base flows. Each layer of the model separates the water into free water, and water in tension. The total soil moisture budget is given by:

$$\frac{dW(t)}{dt} = P(T) - E(t) - R(t) - G(t) \qquad (5)$$

where $W(t)$ is the soil water content at time t, $P(t)$ is the mean precipitation, $E(t)$ is the mean evaportranspiration, $R(t)$ is the net streamflow divergence, and $G(t)$ is the loss to groundwater. The surface runoff $S(t)$ and base flow component $B(t)$ combine to give the total divergence:

$$R(t) = S(t) + B(t) \qquad (6)$$

Each of the budget components is governed by explicit dynamics set forth by physically guided operational hydrologic forecast procedures. Details on these dynamics, as well as a detailed explanation of model parameters is given in *Huang et al*. (1996).

5.4.5 Flow Routing and Model Implementation

To transform the outputs of SNOW-17 and SAC-SMA into a final hydrograph for the North Fork basin, a number of routing components had to be implemented. A rain-snow elevation component (RSNELEV) was developed to obtain a time dependent rain snow elevation from the mean areal temperature data. Since the upper and lower basins were modeled explicitly, each output was weighed through a separate module (WEIGHT-TS) relative to the area of each

basin before being passed to a unit hydrograph module (UNIT-HG). The unit hydrograph converted the SAC-SMA discharge (height of water) to a mean basin discharge (flow per unit time). The parameters for the unit hydrograph were obtained from operational parameters from NWS through direct correspondence. The resulting instantaneous discharge was then passed through a lag and attenuation model (LAG-K), before being averaged (MEAN-Q) to obtain six-hour mean discharge hydrographs. The final implemented model, along with all model inputs is shown in Figure 5.5.

The model source code was obtained from NWS[3] along with operational parameters for the past five decades. The SNOW-17, SAC-SMA, and UNIT-HG models were implemented in the JAVA programing language. Wrappers were written in the Python language to tie all the components together into one executable procedure.

5.5 Results

5.5.1 Regression Skill Assessment

A skill assessment of streamflow forecasting for the North Fork basin (made on April 1st for the total April-July runoff volume) shows that the DWR forecasts had similar behavior seen in a number of other basins [*Morrill et al.,* 2011] (Figure 5.6, Figure 5.7). On average the mean absolute forecast error for the April-July runoff decreased over the forecast season, improving closer to the final forecast date. The mean error for the February forecasts was 30%, decreasing to an average of 5% in May. Evident in figure 5.6 is the below-zero median of the error distribution, indicating the tendency of the forecast to underestimate the full April-July streamflow.

Of most interest in regard to operational forecasting is the April 1st forecast, showing a mean absolute error of 26.3%. The total runoff was overestimated during dry, below average years, while the method under-predicted the total streamflow during wet years (Figure 5.7). The forecast error was largest during these abnormal wet and dry periods, at times reaching as much as 50-60%.

[3] ftp://hydrology.nws.noaa.gov/pub/CHPS/

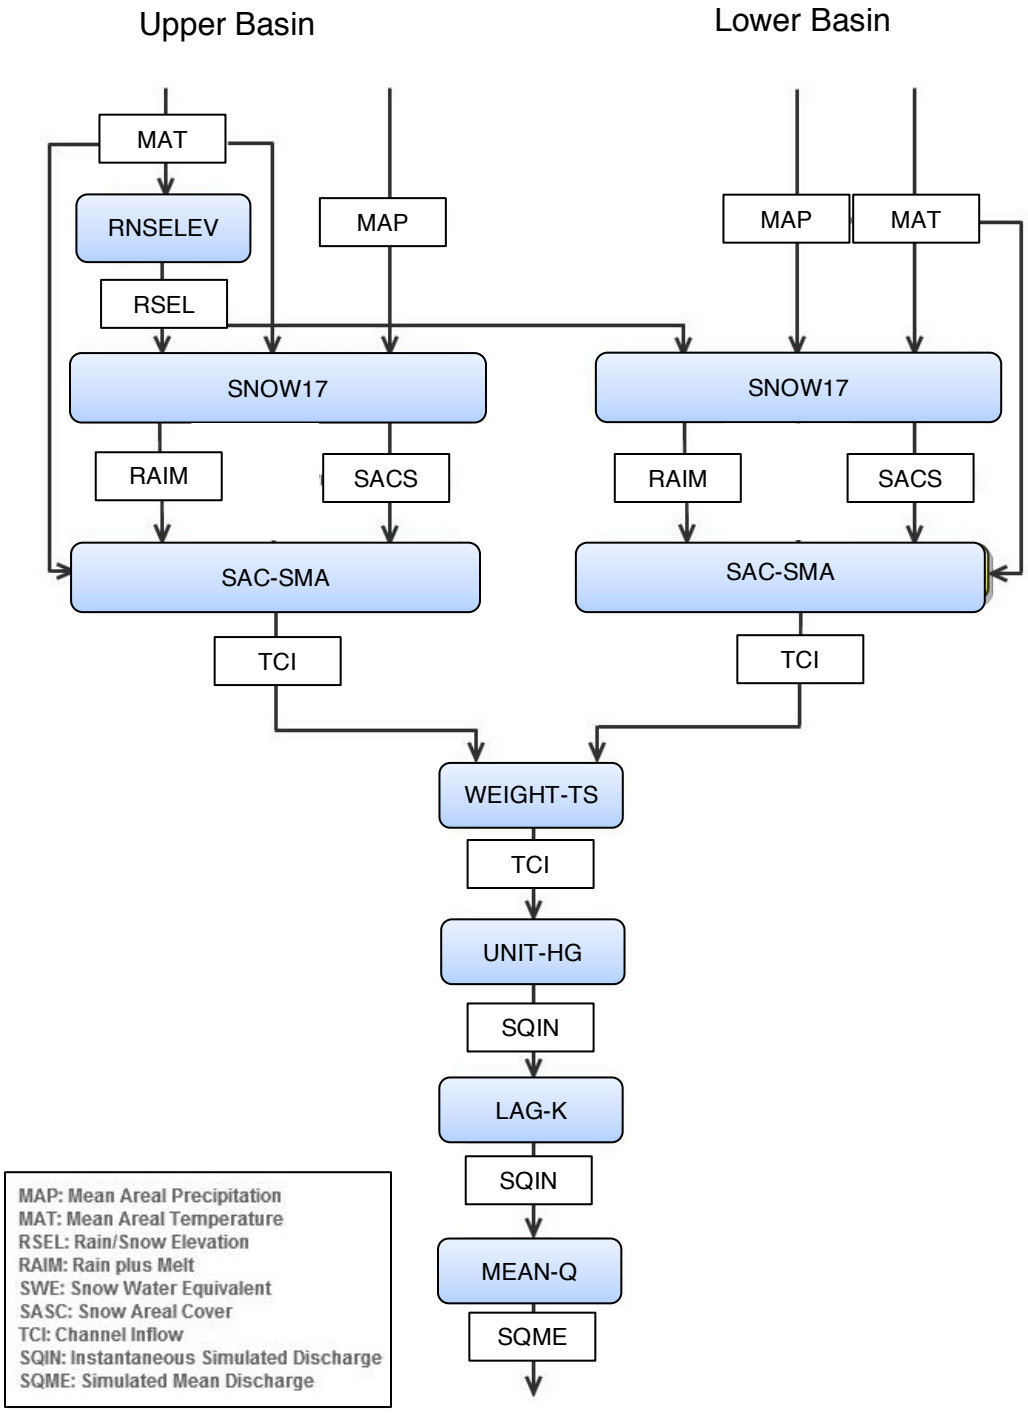

Figure 5.5 Modeling chain for NFDC1 streamflow forecast point.

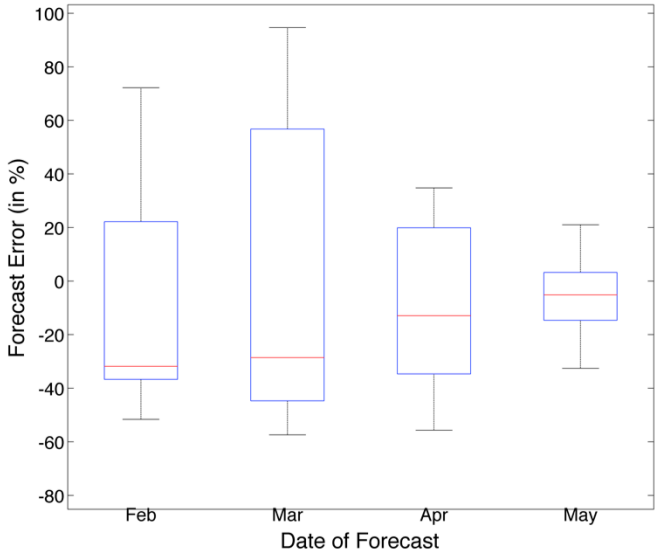

Figure 5.6 Comparison of the February-May forecasts for the April-July runoff using the DWR regression methods (2001-2011). Red lines indicate the median, blue boxes are one standard deviation.

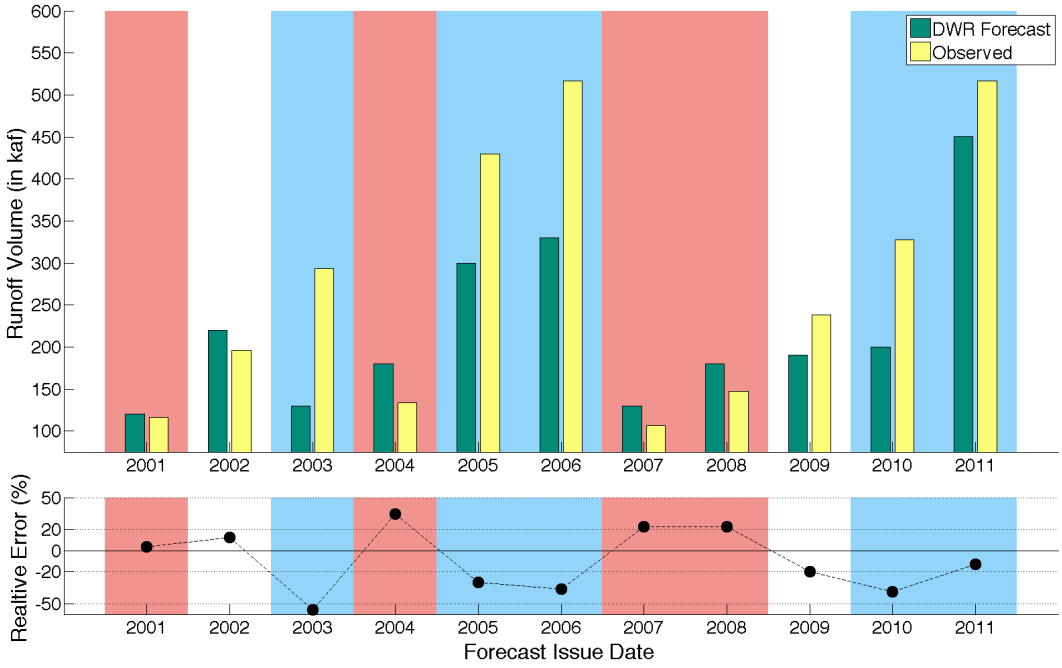

Figure 5.7 Comparison of April-July streamflow observations (yellow) to those derived on April 1st by DWR regression methods for years 2001-2011. Dry, below average, years are denoted with a red background, while wet years are denoted with a blue background.

5.5.2 Effects of Increasing the Number of Sampling Locations on Regression Error

The aim of this part of the study was to implement a complete regression forecast framework and to test the sensitivity these commonly used regression methods to the addition of snowdepth data samples. Stated simply, the aim was to determine if regression methods can improve with the availability of more data (e.g. more WSN deployments). Given the input collinearity discussed earlier, as well as the desire to use a more robust framework than just simple multiple linear regression, principal component regression was implemented.

A PCA analysis carried out on the snowcourse locations (Figure 5.2) revealed that the first principal component accounted for 83% of the variation in the data. The second and third component accounted for 6.3 and 3.7% of the variation of the data, respectively. To ensure that the components were stable in composition, a bootstrap method, as suggested by *Efron and Tibshirani* (1993), was implemented. The available years of data were resampled with replacement and the relative bootstrapping error and loadings of the principal components were computed. It was once again found that only the first principle component had non-significant error, with the mean error of the loadings strictly less than the magnitude of the loadings. The first component was thus the only used in the remainder of the PCR procedure.

To understand the impact of snow course selection, as well as the number of stations required to make robust predictions, we implemented a remove-3 cross validation, consisting of using an exhaustive set of replicates of eight years of training data, and three years of testing data (2001-2011). Exhaustive sets of snow courses readings on April 1st, starting from one course per set, up to 25 courses per set, were then selected. A PCR model was built, and April-July runoff error calculated.

The results show that, on average, the forecast error decreased as more snow courses were added to the PCR model (Figure 5.8). On average, the median error dropped notably after incorporating the first two snow course locations into the model, after which the error leveled off and stayed consistent near 48 kaf. The variance of the error did however decrease notably as more sampling locations were added to the model. Figure 5.9 shows the two stations that consistently provided the best predictions across the bootstrapping procedure. These two snowcourses were some of the most physically distant from the actual streamflow forecast point.

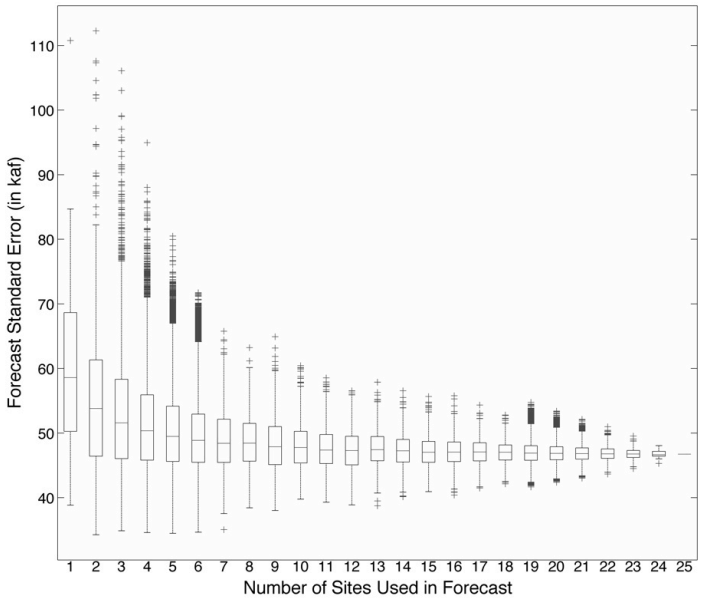

Figure 5.8 The effects of adding more snowcourses on the overall error of a PCR forecast model.

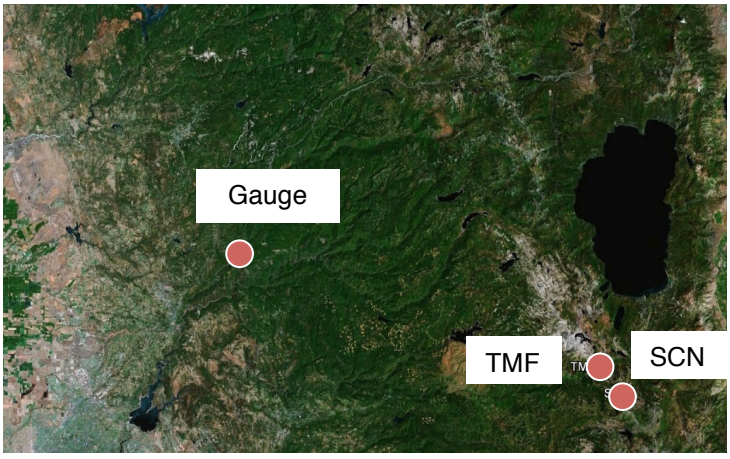

Figure 5.9: Two snowcourse locations which consistently provided the best streamflow predictions through PCR.

5.5.3 Use of Reconstructed SWE Data in Regression Framework

Using the reconstructed SWE data set, a simple linear regression showed a clear relationship between the mean basin-wide SWE on April 1st, and the April-July runoff volume (Figure 5.10). The basin-wide mean SWE was relatively low since the elevations below 1500m, which comprise the majority of the overall basin, typically experience no snow on April 1st.

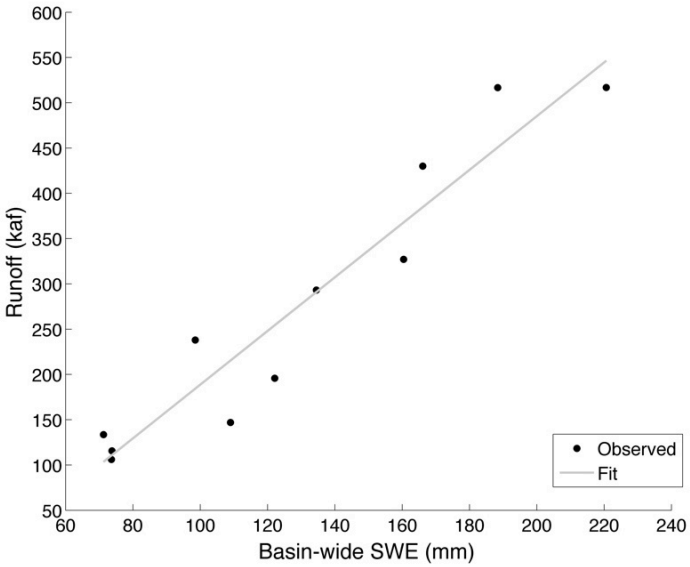

Figure 5.10 Runoff as a linear function of reconstructed (*ground truth*) SWE.

An exhaustive bootstrapping procedure was then carried out to determine the mean absolute error of the April-July forecast when using the linear model based basin-wide mean SWE estimates derived from the reconstructed dataset. The mean absolute error obtained when conducting forecasts using the reconstructed SWE data was 18.5% (Figure 5.11). Compared to the distribution of the DWR forecast errors, the predictions based on the reconstructed SWE data revealed a shift in the median toward zero, indicating less of a tendency to under predict the streamflow. Furthermore the variance of the errors decreased compared to the errors obtained from the DWR forecasts.

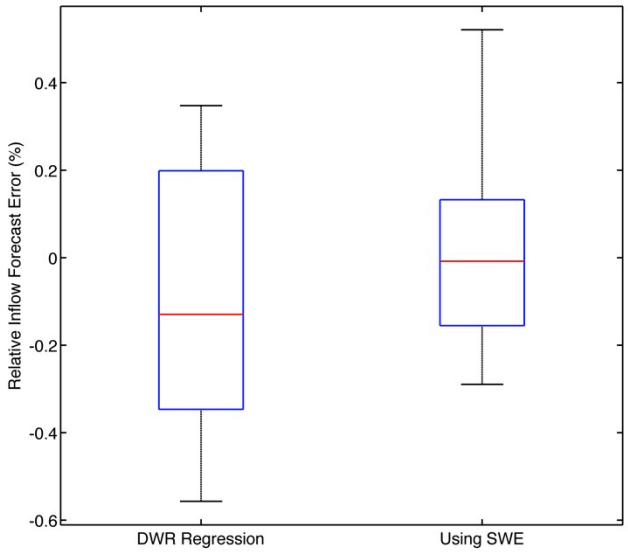

Figure 5.11 Plot of mean absolute error distribution for the April-July runoff forecast, comparing the DWR forecast, to a linear model based on the reconstructed SWE data set. Red lines indicate the median, blue boxes are one standard deviation.

5.5.4 Physical model sensitivity analysis

A sensitivity analysis was carried out to investigate the effects of snowmelt error on SAC-SMA, including the routing components leading up to the modeling of the final basin discharge. The parameterized SNOW-17 model produces as its output a melt time series, which is then routed to the soil accounting model. This time series was taken as the base case for years 2001-2011. The soil moisture accounting model and downstream routing modules were then re-run for a the full 2001-2011 time period, starting at 75% of the total base-case melt, and going up to 125% of the base-case melt in 1% increments. The total change in the predicted April-July basin outflow was computed for each case. While there was a linear relationship between the change in snowmelt and total basin discharge on an annual basis, this relationship varied year to year (Figure 5.12). A 20% difference in the snowmelt estimate corresponded with a difference of 18%-45% for the total streamflow output of the basin. This error range also became more pronounced as the change in snowmelt became larger, as seen by the bounding conical shape in Figure 5.12, showing that the nonlinearities of the model can cause even minor changes in the estimate of SWE to significantly impact the modeled basin discharge.

A similar analysis was carried out higher in the modeling tool chain, varying the modeled base-case of the rain-snow-elevation (RESELEV), showing a nearly 3:2 correspondence between change in outflow against a change in the rain snow elevation. Since the rain-snow-elevation was strictly computed using a point-wise measurement of temperature, the sensitivity of model output to RESELEV effectively reflected both errors in the estimate of the real rain-snow-elevation, as well as the estimate of mean areal temperature.

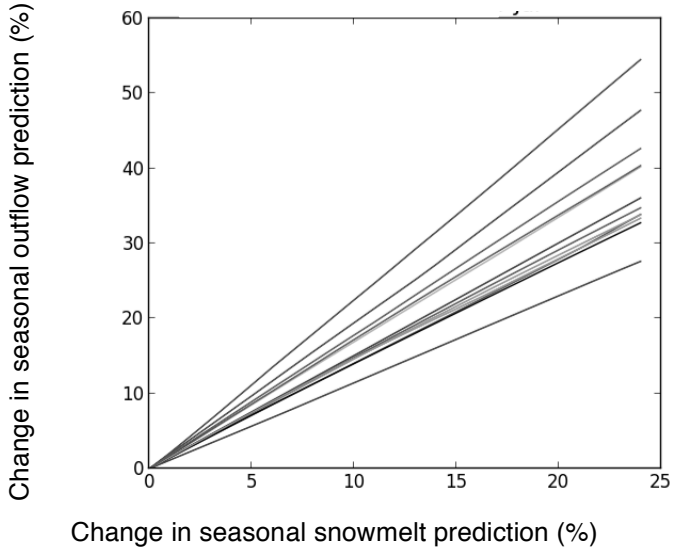

Figure 5.12 The effect of changing the snowmelt quantity (SNOW-17 output) on the effect on overall April-July basin discharge. Each line indicates a model year (2001-2011).

5.5.5 Use of reconstructed SWE data within the physically-based model

Reconstructed runoff, computed from the reconstructed, delineated and averaged SWE for the upper basin, was compared to the upper basin runoff generated by SNOW-17, showing an average of 25% more runoff produced in the April-July season through the reconstructed SWE

data (Figure 5.13). The melt quantities agreed the most in years 2004, and 2006, but no general correlation between the annual precipitation (wet or dry) was evident in the difference between the two runoff estimates. Each model year's relative melt dynamics agreed more in April-May, diverging however significantly in some cases for the May-July period. In particular, there appeared to be significant differences in melt-out timing, with the reconstructed runoff showing, on average, sooner melt-out, and relatively less melt toward the end of the snowmelt season.

Given the results from the sensitivity analysis, particularly those seen in figure 5.12, it was evident that the difference between the reconstructed runoff, and modeled SNOW-17 runoff would produce significant overestimates in the overall basin streamflow when forwarded to SAC-SMA and other routing components. While a thorough comparison of the reconstructed SWE data with snow pillow sites showed no bias in the overall reconstructed data [*Rittger et al*, 2011], it did appear from the presented data that there was a tendency to overestimate SWE values above 1500mm, and under estimate those below. Furthermore, a comparison carried out by the same authors with a separate SWE dataset, interpolated from nearby snow-pillow sites, showed that the snow pillow-based method yielded on average 68% less total water volume when compared to the energy reconstructed SWE data. Since our modeling tool chain was calibrated on nearby snow pillows, this adjustment factor was first applied to the melt output of the reconstructed snowmelt before it was routed to the SAC-SMA.

A 15% mean absolute error was computed for the April-July period for both the modeled basin discharge based on the SNOW-17 output, as well as that modeled discharge based on the adjusted, reconstructed SWE dataset (Figure 5.13). The model based on the reconstructed SWE data set, outperformed the SNOW-17 model in about half the cases, showing especially close correspondence with observed streamflow in 2001 through 2004. The model based on the reconstructed SWE data was also very effective at recreating the correct steam flow conditions during the May-June period, during which on numerous occasions the modeled snowmelt and reconstructed snowmelt showed significantly different melt out dates and quantities. This was also observed by *Rittger et al.*, (2011), who noted distinct melt-out for the reconstructed SWE data set when compared to other reconstruction products.

A breakdown into monthly distributions of error for total flow remaining throughout July, showed that the model based on the SNOW-17 output had a median closer to zero when compared to the model driven by the reconstructed SWE data (Figure 5.14), which showed a consistent median near -12%. The error variance of the SNOW-17 driven model was however significantly greater than that of the model driven by the reconstructed SWE data. Thus, while the reconstructed SWE data model under predicted total seasonal streamflow each month, it did so consistently, and with less than half the error variance.

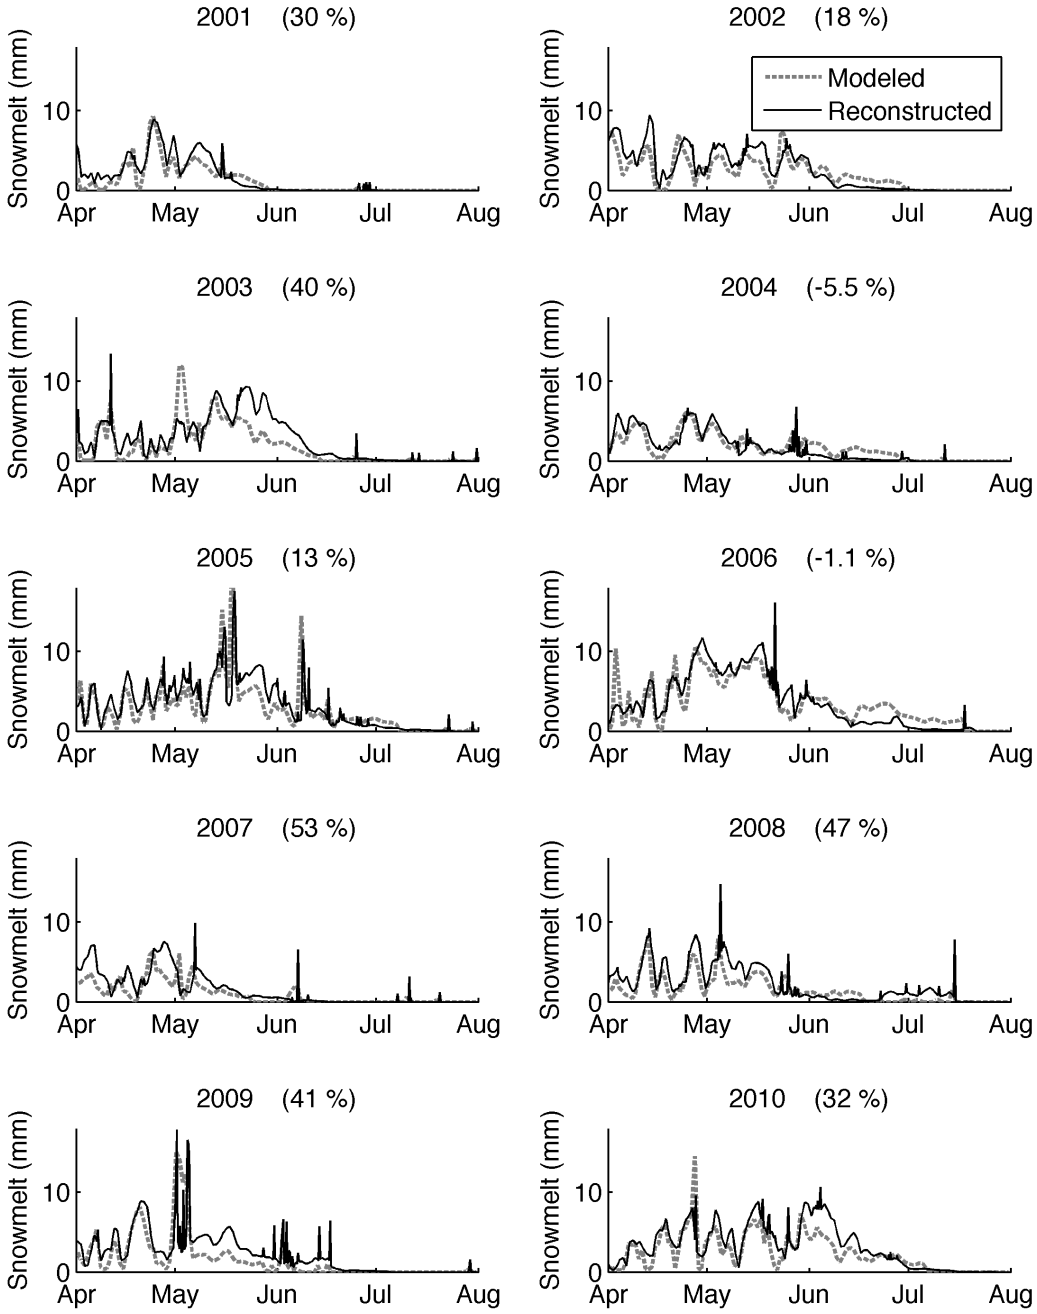

Figure 5.13 Comparison of modeled runoff from snowpack (SNOW-17) against runoff obtained from the reconstructed dataset (2001-2010). The percentages above each plot indicate difference between the reconstructed runoff, and modeled runoff, with a positive percentage indicating more melt produced by the reconstructed data set.

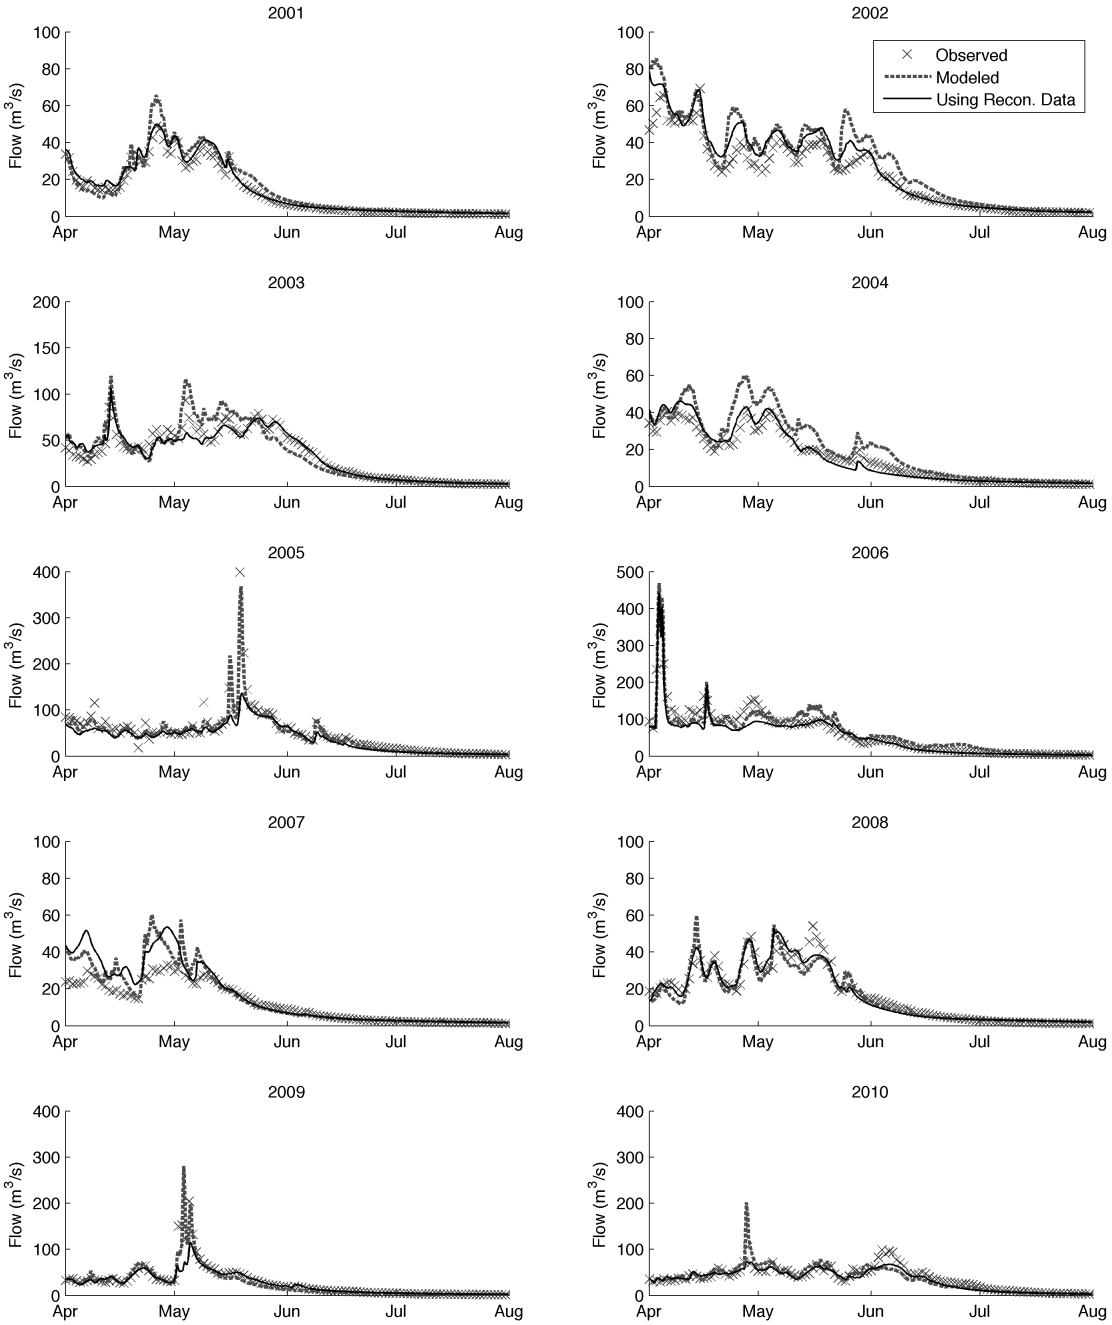

Figure 5.14 Comparison of modeled April-July streamflow against streamflow obtained from the reconstructed dataset (2001-2010).

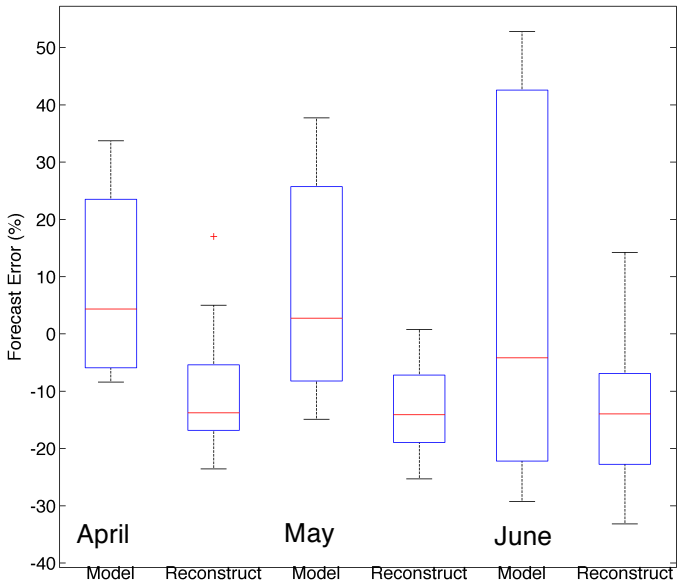

Figure 5.15 Comparison of monthly forecasts of seasonal flow comparing modeled streamflow against streamflow obtained from the reconstructed dataset (2001-2010).

5.6 Discussion

An analysis of historical DWR regression forecasts for the North Fork of the American River further underlines the need to carefully examine the role of snow course data in streamflow forecasts, specifically focusing on the stationarity assumption that is made on an annual basis. As figure 5.7 showed, there is significant tendency to under predict streamflow volumes during wet years, and over predict them during dry years. This behavior is closely tied to the nature of the modeling framework, which tends to drive the predictions closer to a historical mean. Simply increasing the number of predictions does not improve the performance of these models, even when implementing more complex Principal Component Regression. As figure 5.8 showed, adding more snowcourse observations did not noticeably improve the overall prediction. A primary reason for this behavior is given by the co-linearity of the input data set, which features closely correlated snow course observations. These observations, while made at different elevation transects, do not capture the variability of snowdepth that is governed by other physiographic features [*Molotch et al.,* 2006]. This lack of sampling diversity, which was also addressed in chapter three, appears to impose a limit to improvement of these methods. This agrees with some traditional findings, which noted that while these forecasts continue to be important, overall large improvement in their accuracy could not be expected through any further refinement of the technique [Shafer and Huddleston, 1984].

While linear regression methods may posses some inherent limitations due to their simplicity, a large portion of the prediction error is caused by the input dataset. As figures 5.10 and 5.11 showed, a simple, non PCR-based regression relating basin-wide SWE and April-July runoff, was able to reduce the forecast error by nearly 9%, while lowering the prediction variance and reducing the tendency to under-predict seasonal forecasts. While the reconstructed SWE data set may have been biased, as stated previously, the relative relationship between basin-wide SWE and streamflow remained consistent. The improvement of the forecast was a

result of the variability captured by the reconstructed SWE dataset, something that was stressed in Chapter Three, and is currently not captured by the selection of current snow courses.

It was shown that the NWS model is noticeably sensitive to changes in SWE inputs, indicating the SWE estimation error could propagate through the modeling tool chain to induce significant changes to the overall forecast prediction (up to two-fold in some cases, with a 10% snowmelt error causing a 20% forecasting error). This is inherently connected to the non-linearities of the model and underscores the need for establishing consistency in the snowmelt estimates, as large fluctuations in the estimates will result in large variance of the forecasts.

While there was a notable difference between modeled SWE and reconstructed SWE (Figure 5.13), further investigation is needed to determine the ability of either approach to represent true basin-wide snow quantities. A comparison could be drawn with nearby snow pillows, but there are not enough measurements available to ascertain basin-wide variability. Deployments of future WSNs in the regions will determine the efficacy of each approach in the coming years.

Since the NWS model was calibrated to the outputs of SNOW-17, an adjustment-, or scaling-factor, obtained from published records in *Rittger* et al. (2011), had to be applied for meaningful streamflow outputs to be obtained. While this adjusted the total snowmelt, it did so only with regard to total SWE quantity, while retaining the relative variability and melt characteristics of the data set. The final results (Figure 5.15) indicate that while the errors obtained by using modeled snowmelt against reconstructed snowmelt were similar, the total variance of the errors was significantly improved through the use of the reconstructed data set. It is expected that this consistency will offer the opportunity to approach model re-parameterization with a more focused fashion in the future.

5.7 Conclusions

Given the large socioeconomic implication of water resources in the American West, a convincing case of the benefits of improved SWE data has to be made to private and public stakeholder to continue the investment into more advanced snowpack monitoring techniques. This chapter discussed the potential role of improved SWE estimates on the two de-facto water forecasting frameworks currently used in California's Sierra Nevada: DWR regression techniques, and the NWS physically-based streamflow forecast model. Our analysis used the same decade-long SWE dataset that was used in Chapter Three. While this data set may not be the exact ground truth of SWE during this period, it is currently the only data set of its kind, offering truly independent, spatiotemporally dense distributions of SWE. Use of this dataset during the April-July snowmelt season showed that forecasts based on regression techniques could directly benefit from improved SWE data, with a reduction in nearly 9% of mean error, as well as a reduction in error variance. It was also shown that while the output of the physical model was not directly improved through the use of the SWE dataset, the consistency of the error predictions did improve by two-fold, providing a good starting point upon which to base model re-parameterization. Future work will focus on conducting such re-parameterizations, and investigating ensemble estimation frameworks to more optimally employ real-time data from WSNs into the operational model.

Chapter 6: Thesis Conclusions and Future Work

The previous chapters outlined a system-level approach towards developing an improved understanding of snowmelt driven processes in the Sierra Nevada, and outlined methods that can be used to incorporate these estimates into streamflow forecasts. The thesis showed that improvements to large-scale water forecasting begin at the smaller-scale, where better measurements and scientific understating can be expanded to infer system behavior at the larger basin scale. It is expected that the implementation of the above methods could play a significant role in enhancing hydrologic understanding, while improving large-scale water resource management. The major conclusions of this thesis can be summed up as follows:

- **Distributed Wireless Sensing**
  - It was shown that a well-defined deployment strategy can be used to effectively instrument catchment-scale regions using extremely low power wireless sensor network technology.
  - An extension to the proposed method should investigate the feasibility of scaling such deployment strategy to cover significantly larger, basin-scale regions. This is presently being conducted in the American River basin, where thirty smaller-scale networks will be deployed to estimate SWE over the coming years.
- **Variability of catchment-scale hydrologic processes**
  - Given the quantity of data acquired by our WSN deployment, we confirmed the temporal stability of snowdepth variability across the melt season, suggesting that a stratified placement strategy based on evenly instrumenting major physiographic parameters performs well with regard to characterizing the distribution of catchment-wide snow depth at the km^2 scale.
  - An analysis of data acquired by the network over the coming years will solidify our findings on the variability of soil moisture across the snowmelt season. An investigation into variability of evapotranspiration, and groundwater distribution will also be required to begin closing the loop on a truly physically-based water balance estimate.
- **Formal abstractions of natural systems**
  - A computationally tractable, *Hybrid System* model of snowmelt was introduced and shown to accurately reflect the various stages of snowmelt when compared to observations in the Sierra Nevada.
  - This model should be investigated in a larger context of water supply forecasting, by coupling it with WSN data to forecast melt conditions in a large mountain basin.
- **Sensor placement, sampling design, and spatial estimation**
  - A quantitative sensor placement strategy, coupling Statistical Learning with an optimization over *Information Theoretic* variables was developed and shown to improve SWE estimation error by up to 100mm compared to a number other methods.
  - As the proposed method was performed on a historical data set, future work should focus on evaluating its real-world performance by evaluating the ability of the proposed locations to estimate SWE in the American River basin using the WSNs that twill be deployed in the coming years.

- **Effects on improved SWE estimates on basin-scale water resource management**
  - It was shown that present forecasting methods are sensitive to both parameters and inputs. The use of historically reconstructed SWE data showed that streamflow forecasting model precision was improved that total forecast error could be reduced by nearly 10% when using improved SWE estimates.
  - An investigation on the downstream effects of these improved forecasts should be carried out, focusing especially on the benefits to the control and optimization of reservoir releases and hydropower production.

While this thesis advocates a systemic solution to address basin-scale water forecasting, there exist a number of open problems to address how the resulting data and forecasts could be used within a real-time water management framework. Even if present methods did permit for the full state of a water basin to be estimated, an appropriate system-level architecture does not exist to close the control loop between natural system and physical infrastructure. Future water grids will compose an amalgam of natural systems, infrastructure, sensors, actuators, stakeholders, models, control algorithms, and decision makers. An appropriate architecture must be designed to achieve interoperability between these entities, which is essential for scalable and reliable operations, as well as for the overall feasibility of real-time operations. Such an architecture will ensure that addition or removal of any of the other components occurs seamlessly, and without disruption to the performance of the overall system. For example, adding a new sensor to the network must not affect the overall operations of downstream operations, or decision support strategies. Additionally, modularization and decentralized operations will be key, as the change of any of the other components must not affect the operation of the other services. This formal *architecture* will include, but will not be limited to, routing real-time data feeds, automatically registering sensors and clients, forwarding modeling outputs to the control algorithms, and resolving multi-client conflicts. In essence, such a solution will be the *operating system* upon which all of the other components will interact.

A possible template for such an architecture could be inspired by the modularized structure of the *Internet*, an extremely complex system of many components that interact according to a set of closely followed *standards*. Such standards must be developed for water networks, taking into account the physical input-output nature of the system. The result standards will result in a formal architecture for scalable, and reliable information exchange within water systems. Instrumented watersheds will become only one component in this massive *cyber physical architecture*. To connect these natural phenomena into the grid, one must first represented them as formal system components. A step toward this direction is the development of formal models of both natural and man made systems. An example of such a system is the hybrid snowmelt automaton introduced in Chapter Four of this thesis. Such abstractions should be moved to other components in the water grid, focusing initially on dams and reservoirs, to show that it is possible to effectively exchange information between physical infrastructure and natural systems. This will open the door to a suite of never before conceived control and optimization problems.

With all components formally represented in the proposed *architecture*, it will be possible to control the system as a whole, rather than focusing on individual components, as is the case in current water operations. Such an architecture will likely draw on control theory, and dynamical systems to develop policies that can be used to optimally guide reservoirs releases in real-time. Primary control objectives could center on increasing hydropower production, while

mitigating floods, and reducing overall water *spillage*. With an initial focus on upstream operations, such research presents a manageable task, split into two breaths, one focusing on formalizing the *cyber-physical information architecture*, and the other developing real-time hydrologic models and control strategies to guide reservoir releases. It will then be possible to expand the scope of the system to instrument and control down-stream, urban-scale arterial water networks as well, thus opening the door to a truly all-encompassing, *intelligent water grid*.

Bibliography

Al Basset Almamou, A., R. Wrede, P. Kumar, H. Labiod, and J. Schiller (2009*)*, Performance evaluation of routing protocols in a real-world WSN, *Information Infrastructure Symposium GIIS,* 1-5.

Alur, R., C. Belta, F. Ivancic, V. Kumar, M. Mintz, G. J. Pappas, H. Rubin, and J. Schug (2001), Hybrid modeling and simulation of biomolecular networks, *Hybrid Systems Computation And Control*, 2034.

Anderson, E. A. (1973). National Weather Service River Forecast System-Snow Accumulation and Ablation Model. National Oceanographic and Atmospheric Administration, Silver Springs. *Tech. Mem., NWS HYDRO*, *17*, 217.

Anderson, S., R. Bales, and C. Duy (2008), Critical zone observatories: Building a network to advance interdisciplinary study of earth surface processes, *Mineralogical Magazine*, 72(1).

Aspen Environmental Group and M. Cubed, Potential Changes in hydropower production from global climate change in California and the western United States (2005), prepared for the CA Energy Commission. Atkinson, A. C. (1988), Recent developments in the methods of optimum and related experimental designs, *International Statistical Review / Revue Internationale de Statistique,* 56, 2 (August), 99–115.

Association of California Water Agencies (ACWA), California Water (2011), http://www.calwatercrisis.org/program.htm, last accessed November 10, 2012.

Bales, R. C., Hopmans, J., O'Green, A., Meadows, M., Hartsough, P., Kirchner, P., Hunsaker, C., Neaudette, D. (2011), Soil moisture response to snowmelt and rainfall in a Sierra Nevada mixed- conifer forest. *Vadose Zone Journal*.

Bales, R. C., N. P. Molotch, T. H. Painter, M. D. Dettinger, R. Rice, and J. Dozier (2006), Mountain hydrology of the western United States, *Water Resour. Res.*, *42*, W08432, doi:10.1029/2005WR004387.

Balk, B., Elder, K. (2000), Combining binary decision tree and geostatistical methods to estimate snow distribution in a mountain watershed, Water Resources Research 36(1): 13–26.

Barnett, T. P., J. C. Adam, and D. P. Lettenmaier (2005), Potential impacts of a warming climate on water availability in snow-dominated regions, *Nature, 438,* 303-309, doi: 10.1038/nature04141.

Bartelt, P., and M. Lehning (2002). A physical snowpack model for the swiss avalanche warning: Part i: Numerical model, *Cold Regions Science and Technology*, 35(3).

Bartelt, P., and M. Lehning (2002). A physical snowpack model for the swiss avalanche warning: Part iii: Meteorological forcing, thin layer formation and evaluation, *Cold Regions Science and Technology*, 35(3).

Bear, J. (1988), *Dynamics of fluids in porous media*, Dover, NY.

Bell G., and C. Federspiel (2009), Demonstration of datacenter automation software and hardware (dash) at the California franchise tax board, *Internal report prepared for the California Energy Commission*.

Blöschl, G. (1999), Scaling issues in snow hydrology, *Hydrol. Processes*, 13, 2149–2175, doi:10.1002/hyp.847.

Bogena, H.R., M. Herbst, J.A. Huisman, U. Rosenbaum, A. Weuthen, H. Vereecken (2010), Potential of Wireless Sensor Networks for Measuring Soil Water Content Variability Vadose Zone Journal, 2010(9):1002–1013.

Borri, A., M. Domenica, D. Benedotto, and M.-G. D. Benedetto (2010). Hybrid modeling, power management and stabilization of cognitive radio networks, *Hybrid Systems Computation And Control*, San Francisco, CA.

Brubaker, K., A. Rango, and W. Kustas (1996), Incorporating radiation inputs into the snowmelt runoff model, *Hydrol Process*, 10, 1329-1343.

Burnash, R. J. C., Ferral, R. L., & McGuire, R. A. (1973). *A Generalised Streamflow Simulation System–Conceptual Modelling for Digital Computers. Joint Federal and State River Forecast Center*. Sacramento, Technical Report.

Burnash, R. J., Ferral, R. L., & McGuire, R. A. (1973). A generalized streamflow simulation system, conceptual modeling for digital computers.

CADWR (2011). California Data Exchange Center, http://cdec.water.ca.gov/.

California cooperative snow surveys, http://www.water.ca.gov/floodmgmt/hafoo/hb/sss/, last accessed November 10, 2012.

California water plan (2009), http://www.waterplan.water.ca.gov/cwpu2009/index.cfm, last accessed November 10, 2012.

Campbell Scientific (2012), CR1000 Measurement and control datalogger, *http://www.campbellsci.com/cr1000*, January 2012.

Carroll, S.S. (1995), Modeling measurement errors when estimating snow-water equivalent, *Journal of Hydrology*, 172, 247-260.

Carroll, S.S., and N. Cressie (1996), A comparison of geostatistical methodologies used to estimate snow-water equivalent, *Water Resour. Bull.*, 32(2), 267-278.

Cayan, D. R. (1996), Interannual climate variability and snowpack in the western United States, *J. Clim.*, 9, 928–947, doi:10.1175/1520-0442(1996)009<0928:ICVASI>2.0.CO;2.

Chipman, H. A., E. I. George, R. E. McCulloch (2002), Bayesian Treed Models, *Machine Learning,* 48(1), 299-320, doi: 10.1023/A:1013916107446.

Christensen L., C. L. Tague, and J. S. Baron (2008), Spatial patterns of simulated transpiration response to climate variability in a snow dominated mountain ecosystem, *Hydrol. Process.* 22(18): 3576-3588.

Clark, M. P., M. C. Serreze, and G. J. McCabe (2001), Historical effects of El Nino and La Nina events on the seasonal evolution of the montane snowpack in the Columbia and Colorado River Basins, *Water Resour. Res*., 37, 741–757, doi:10.1029/2000WR900305.

Cline, D. W., R. C. Bales, and J. Dozier (1998), Estimating the spatial distribution of snow in mountain basins using remote sensing and energy balance modeling, *Water Resour. Res.,* 34, 1275–1285.

Colbeck, S. (1971) One dimensional water flow through snow, *Res. Rep.,* 296:17.

Colbeck, S. (1972), A theory of water percolation in snow, *J. Glaciology*, 11:369{385.

Colbeck, S. (1972), Water flow through snow overlaying an impermeable boundary, *Water Resources Research*, 10:119{123.

Colbeck, S. (1975), A theory for water flow through a layered snowpack, *Water Resources Research*, 11(2):261{266.

Colbeck, S. (1978), Short-term forecasting of water runoff from snow and ice, *J. Glaciology*, 19:571{588.

Colbeck, S. (1980), *Dynamics of snow and ice masses*, Academic Press Publishers.

Cosgrove, B. A., et al. (2003), Real-time and retrospective forcing in the North American Land Data Assimilation System (NLDAS) project, J. Geophys. Res., 108, 8842, doi:10.1029/2002JD003118.

Crank, J. (1984), *Free and moving boundary value problems*, Oxford University Press.

Cressie, N. (1991), *Statistics for Spatial Data.* Wiley.

Culler, D., L. Doherty, J. Hill, M. Holden, C. Kiers, S. Kumar, J. Kusuma, S. Morris, K. Pister, K. Ramchandran, B. Robbins, J. Scholtz, M. Scott, R. Szewczyk, B. Tang, and A. Woo (2001), 29 palms fixed/mobile experiment: Tracking vehicles with a uav-delivered sensor network, *UC Berkeley Technical Report*.

Curtis, D. C., & Schaake, J. C. (1979). The NWS Extended Streamflow Prediction Technique. In *Proceedings of the Conference on Water Conservation Needs and Implementing Strategies July 9-13, 1979, Franklin Pierce College, Rindge, NH, p 182-195, 1979. 2 Fig, 1 Tab, 17 Ref.*.

Daly, S. F., R. Davis, E. Ochs, and T. Pangburn (2001), An approach to spatiallydistributed snow modeling of the Sacramento and San Joaquin basins, California, *Hydrol. Processing, 14,* 3257–3271, doi:10.1002/ 1099-1085(20001230)14:18<3257::AID-HYP199>3.0.CO;2-Z.

Das, A. and D. Kempe (2008). Algorithms for subset selection in linear regression.*Proc. Annu. ACM Symp. Theory Comput*. ACM, New York, 45–54.

Day, G. N. (1985). Extended streamflow forecasting using NWSRFS. *Journal of Water Resources Planning and Management*, *111*(2), 157-170.

Deems, J.S., S.R. Fassnacht, K.J. Elder (2006), Fractal Distribution of Snow Depth from Lidar Data, *J. Hydrometeor*, 7, 285–297.

Demargne, J., Mullusky, M., Werner, K., Adams, T., Lindsey, S., Schwein, N., Welles, E. (2009). Application of forecast verification science to operational river forecasting in the US National Weather Service. *Bulletin of the American Meteorological Society*, *90*(6), 779-784.

Dozier, J., T. H. Painter, K. Rittger, J. E. Frew (2008), Time–space continuity of daily maps of fractional snow cover and albedo from MODIS, *Advances in Water Resources*, 31(11), 1515-1526, doi:10.1016/j.advwatres.2008.08.011.

Dozier, J., and J. Frew (1990), Rapid calculation of terrain parameters for radiation modeling from digital elevation data, *IEEE Trans. Geosci.Remote Sensing*, 28(5), 963-969.

Dubayah, R. (1992), Estimating net solar radiation using Landsat Thematic Mapper and digital elevation data,*Water Resour.Res.*, 28(9), 2469-2484.

Dust (2006), Technical overview of Time Synchronized Mesh Protocol (TSMP), *whitepaper, Hayward, CA*.

Dust (2011), Embedded wireless sensor networks for monitoring and control, *http://www.dustnetworks.com/,* January 2012.

Efron, B., & Tibshirani, R. J. (1994). *An introduction to the bootstrap* (Vol. 57). Chapman & Hall/CRC.

Ember (2011), ZigBee wireless network systems, *http://www.ember.com/,* January 2012.

EME (2011), Electronically monitored ecosystems, *http://www.emesystems.com/,* January 2012.

Emerson (2010), Emerson Wireless Process Management, *http://www2.emersonprocess.com/en-US/plantweb/wireless/,* January 2012.

Erickson, T. A., M. W. Williams, and A. Winstral (2005), Persistence of topographic controls on the spatial distribution of snow in rugged moun- tain terrain, Colorado, United States, *Water Resour. Res.*, 41, W04014, doi:10.1029/2003WR002973.

Erxleben, J., K. Elder, and R. Davis (2002), Comparison of spatial interpolation methods for estimating snow distribution in the Colorado Rocky Mountains, *Hydrol. Processes*, 16, 3627–3649, doi:10.1002/hyp.1239.

Etzel, B., and H. Braun (2005), High-Performance Wireless Internet Connection to Mount Laguna Observatory, *Bulletin of the American Astronomical Society*, 32, 1428.

Famiglietti, J.S., J.W. Rudnicki, and M. Rodell (1998), Variability in surface moisture content along a transect: Rattlesnake Hill, Texas. *Journal of Hydrology* 210: 259–281.

Faria, D. A., J. W. Pomeroy, R. L. H. Essery (2000), Effect of covariance between ablation and snow water equivalent on depletion of snow-covered area in a forest, Hydrol. Process. 14(15): 2683-2695.

Fassnacht, S. R., K. A. Dressler, and R. C. Bales (2003), Snow water equivalent interpolation for the Colorado River basin from snow telemetry (SNOTEL) data, *Water Resour. Res.,* 39(8), 1208, doi:10.1029/ 2002WR001512.

Fassnacht, S.R., J. E. Derry (2010), Defining similar regions of snow in the Colorado River basin using self organizing maps, *Water Resour. Res*., 46, W04507, doi:10.1029/2009WR007835.

Franz, K. J., Hartmann, H. C., Sorooshian, S., & Bales, R. (2003). Verification of National Weather Service ensemble streamflow predictions for water supply forecasting in the Colorado River basin. *Journal of Hydrometeorology*, *4*(6), 1105-1118.

Friis, H. (1946), A note on a simple transmission formula, *Proceedings of the IRE*, 34, 254-256.

Friis, H. (1971), Introduction to radio and radio antennas, *Spectrum, IEEE*, 8(4), 55-61.

Garen, D. C. (1992). Improved techniques in regression-based streamflow volume forecasting. *Journal of Water Resources Planning and Management*,*118*(6), 654-670.

Garen, D. C., and D. Marks (2005), Spatially distributed energy balance snowmelt modelling in a mountainous river basin: estimation of meteorological inputs and verification of model results, *Journal of Hydrology*, 315(1-4), 126-153.

Ghosh, R., and C. J. Tomlin (2001), Lateral inhibition through delta-notch signaling: A piecewise plane hybrid model, *Hybrid Systems Computation And Control*, 2034/2001:232{246.

Gonzalez-Banos, H.H., and J. Latombe (2001), A randomized art-gallery algorithm for sensor placement, *Proc. 17th ACM Symposium on Computational Geometry,* 232–240.

Grant, L., M. Seyfried, and J. McNamara (2004), Spatial variation and temporal stability of soil water in a snow-dominated, mountain catchment, Hydrol. Process. 18:3493-3511.

Guestrin, C., A. Krause, and A. P. Singh (2005), Near-optimal sensor placements in gaussian processes.*In Proc. Int. Conf. Machine Learning*. ACM, New York, 265–272.

Guo, Q., W. Li, H. Yu, and O. Alvarez, (2010), Effects of Topographic Variability and Lidar Sampling Density on Several DEM Interpolation Methods, *Photogrammetric Engineering & Remote Sensing,* 76(6): 701-712.

Hannah, L. A., and D. B. Dunson (2011), Multivariate convex regression with adaptive partitioning.*ArXiv*. arXiv:1105.1924v2.

Harshburger, B. J., K. S. Humes, V. P. Walden, T. R. Blandford, B. C. Moore, and R. J. Dezzani (2010), Spatial interpolation of snow water equivalency using surface observations and remotely sensed images of snow-covered area, *Hydrol. Processing, 24,* 1285–1295, doi:10.1002/hyp.7590.

Hart, K., and K. Martinez (2006), Environmental sensor networks: A revolution in the earth system science?, *Earth-Science Reviews*, 78(3-4),177 — 191.

Hartmann, H. C., Bales, R., & Sorooshian, S. (2002). Weather, climate, and hydrologic forecasting for the US Southwest: a survey. *Climate Research*,*21*(3), 239-258.

Hasler, A., I. Talzi, J. Beutel, C. Tschudin, and S. Gruber (2008), Wireless sensor networks in permafrost research — concept, requirements, implementation and challenges, *Proceedings of the 9th International Conference on Permafrost*, vol. Fairbanks, Alaska, USA.

He, M., Hogue, T. S., Margulis, S. A., & Franz, K. J. (2012). An integrated uncertainty and ensemble-based data assimilation approach for improved operational streamflow predictions. *Hydrology and Earth System Sciences*,*16*(3), 815-831.

Honeywell (2010), Honeywell Industrial Wireless Mesh Solutions, *http:/www.honeywell.com/ps/wireless*, January 2012.

Huang, J., & van den Dol, H. KP georgakakos. 1996. Analysis of Model-Calculated Soil Moisture over the United States (1931-93) and Application to Long-Range Temperature Forecasts. *Journal of Climate*, *9*(6), 1350-1362.

IEEE (2009), Institute of Electric and Electronics Engineers Standard for Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (LR-WPANs), *Technical Standard IEEE802.15.4-2009*.

Illangasekare, T., R. Walter, M. Meier, and W. Pfeer (1990), Modeling meltwater infiltration in subfreezing snow, *Water Resources Research*, 26(5):1001{1012.

Ingelrest F., G. Barrenetxea, G. Schaefer, M. Vetterli, O. Couach, M. Parlange (2010), SensorScope: Application-Specific Sensor Network for Environmental Monitoring, *ACM Transactions on Sensor Networks*, 6(2), doi: 10.1145/1689239.1689247.

ISA (2009), Wireless systems for industrial automation: Process control and related applications, *International Society of Automation technical standard*, *ISA-100.11a-2009.*

Jennic (2011), *Jennic Wireless Microcontrollers,* http://www.jennic.com, January 2012.

Jordan, R (1990), *A one-dimensional temperature model for a snow-cover*, US Army SNOW Research Center.

Jost G., M. Wirler, D.R. Gluns, Y. Alila (2007), The influence of forest and topography on snow accumulation and melt at the watershed-scale, *J. of Hydrology,* doi: 10.1016/J.Jhydrot.2007.09.006

Karl, H., and A. Willig (2005), *Protocols and Architectures for Wireless Sensor Networks*, John Wiley & Sons, Hoboken, N.J.

Kelley, R., W. Morland, and E. Morris (1986), A three phase mixture for melting snow, *Budapest Symposium on Modeling Snowmelt-Induced Processes*, 155:17{26.

Kemppainen, A., T. Makela, J. Haverinen, and J. Roning (2008), An experimental environment for optimal spatial sampling in a multi-robot system.*The 10th International Conference on Intelligent Autonomous Systems (IAS-10),* ACM, Baden Baden Germany, 265–272.

Kerkez, B., and S. Glaser (2009), Remote sensing and monitoring strategies of large scale natural systems subjected to extreme conditions, *The Fifth International Workshop on Advanced Smart Materials and Smart Structures Technology*, Boston, MA.

Kerkez, B., T. Watteyne, M. Magliocco, Glaser S., and K. Pister (2009), Feasibility analysis of controller design for adaptive channel hopping, *First International Workshop on Performance Methodologies and Tools for Wireless Sensor Networks (WSNPerf),* Pisa Italy.

Kim, S., S. Pakzad, D. Culler, J. Demmel, G. Fenves, S. Glaser, and M. Turon (2007), Health monitoring of civil infrastructures using wireless sensor networks, *Proceedings of the6th international conference on Information processing in sensor networks*, New York, NY, USA, 254—263.

Kraus, J. (1988), *Anetnnas*. McGraw-Hill, New York, NY, USA.

Krause, A. (2010), SFO: A Toolbox for Submodular Function Optimization,*Journal of Machine Learning Research*.

Krause, A., A. Singh, and C. Guestrin. (2007), Near-optimal sensor placements in gaussian processes: Theory, effecient algorithms and empirical studies. Technical report, *Machine Learning Department,Carnegie Mellon Uni.*

Lee, C. K. J. and M. Queyranne (1995), An exact algorithm for maximum entropy sampling. *ACM Trans. Program. Lang. Syst*. 7, 3, 359–379.

Levia, D., D. Carlyle-Moses, and T. Tanaka (2011), *Forest Hydrology and Biogeochemistry*, Springer.

Libelium (2011), Wireless sensor networks, *http://www.libelium.com/,* January 2012.

MacKay, D. J. C. (2003), *Information Theory, Inference, and Learning Algorithms,* 284pp., Cambridge UP.

Mainwaring, A., D. Culler, J. Polastre, R. Szewczyk, and J. Anderson (2002), Wireless sensor networks for habitat monitoring, *WSNA '02: Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*, New York, NY, USA, 88—97.

Marks, D. and J. Dozier (1992), Climate and energy exchange at the snow surface in the alpine region of the Sierra Nevada: 1. Meteorological measurements and monitoring, *Wat. Resour. Res.,* 28:3029{3042.

Marks, D. and J. Dozier (1992), Climate and energy exchange at the snow surface in the alpine region of the Sierra Nevada: 2. Snow cover energy balance, *Wat. Resour. Res.,* 28:3043{3054.

Marks, D., J. Domingo, and J. Frew (1998), Software tools for hydro-climatic modeling and analysis: Image processing workbench, https://www.nmepscor.org/trac/IPW/

Marks, D., J. Kimball, D. Tingey, and T. Link (1998), The sensitivity of snowmelt processes to climate conditions and forest cover during rain-on-snow: A case study of the 1996 Pacific northwest flood, *Hydrological Processes*, 12:1569{1587.

Marks, D., K. Cooley, D. Robertson, and A. Winstral (2001), Long-term snow database, reynolds creek experimental watershed, Idaho, USA, *Water Resources Research*, 37:2835{2838.

Marks, D., T. Link, A. Winstral, and D. Garen (2001), Simulating snowmelt processes during rain-on-snow over a semi-arid mountain basin, *Annals of Glaciology*., 32.

Martinec, J. and A. Rango (1991), Indirect evaluation of snow reserves in mountain basins, *IAHS*, 205:111{119.

McFadden, D. (2011), Water Supply Forecasting : Statistical modeling for forecasting inflow to the Upper American River Project. Internal Document, *SMUD*.

McGinnis, D. L. (1997), Estimating climate‐change impacts on the Colorado Plateau snowpack using downscaling methods, *Prof. Geogr*., 49(1), 117–125, doi:10.1111/0033-0124.00062.

McKay, M. D., R. J. Beckman, and W. J. Conover (1979). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics,* 21(2): 239–245, May.

McManamon, A., Szeliga, T. L., Hartman, R. K., Day, G. N., & Carroll, T. R. (1993). Gridded snow water equivalent estimation using ground-based and airborne snow data. In *Proceedings of the Eastern Snow Conference* (pp. 75-82).

Meadows M.W., Kerkez B., Hartsough P.C., Lucas R.G., Bales R.C., Hopmans J.W., Glaser S.D. (2010), Comparing plot-scale sensor measurements to the watershed level: a comprehensive case study of snow depth and soil moisture in the southern Sierra Nevada, California. Abstract C13C-03. Fall 2010 meeting American Geophysical Union, San Francisco.

Meromy, L., N.P. Molotch, T. Link, S.P. Fassnacht, and R. Rice (2012), Subgrid variability of snow water equivalent at operational snow stations in the western United States, *Hydrol. Processes.*, in press.

Molotch, N. P., and R. C. Bales (2005a), Scaling snow observations from the point to the grid element: implications for observation network design, Water Resour. Res. 41.

Molotch, N. P., and S. A. Margulis (2008), Estimating the distribution of snow water equivalent using remotely sensed snow cover data and a spatially distributed snowmelt model: A multi-resolution, multi-sensor comparison, *Advances in Water Resources*, 31(11): 1503-1514, doi: 10.1016/j.advwatres.2008.07.017.

Molotch, N. P., M. T. Colee, R. C. Bales, and J. Dozier (2005b), Estimating the spatial distribution of snow water equivalent in an alpine basin using binary regression tree models: the impact of digital elevation data and independent variable selection. Hydrological Processes, 19: 1459–1479.doi: 10.1002/hyp.5586.

Moore, R. D., and I. G. McKendry (1996), Spring snowpack anomaly pat- terns and winter climatic variability, British Columbia, Canada, *Water Resour. Res*., 32, 623–632, doi:10.1029/95WR03640.

Morland, L., R. Kelley, and E. Morris (1990), A mixture theory for phase changing snowpacks, *Cold Regions Science and Technology*, 17:271{285.

Morrill, J. C., Hartmann, H. C. and Bales, R. C. (2011), An assessment of Seasonal Water Supply Outlooks in the Colorado River Basin (in progress).

Musselman, K. N., N. P. Molotch and P. D. Brooks (2008), Effects of vegetation on snow accumulation and ablation in a mid-latitude sub-alpine forest, Hydrol. Process. 22(15): 2767-2776.

N. S. I. Datacenter (2009), Snow water equivalent (swe) field measurements, http://nsidc.org/data/swe/, May 2009.

Nash, J., & Sutcliffe, J. V. (1970). River flow forecasting through conceptual models part I—A discussion of principles. *Journal of hydrology*, *10*(3), 282-290.

Nevis (2011), Persistent LAN security and network access control http://www.nevisnetworks.com, January 2012.

Oestges, C., M. Villacieros, and D. Vanhoenacker-Janvier (2009), Radio channel characterization for moderate antenna heights in forest areas, *IEEE Transactions on Vehicular Technology*, 58 (8), 4031-4035, doi:10.1109/TVT.2009.2024947

Pagano, T., Garen, D., & Sorooshian, S. (2004). Evaluation of official western US seasonal water supply outlooks, 1922-2002. *Journal of Hydrometeorology*,*5*(5), 896-909.

Painter, T. H., K. Rittger, C. McKenzie, P. Slaughter, R. E. Davis, and J. Dozier (2009), Retrieval of subpixel snow‐covered area and grain size, and albedo from MODIS, *Remote Sens. Environ.*, 113, 868–879.

Pierce, J.R. (1980), *An Introduction to Information Theory, second edition*, Dover Publications.

Pister K.S.J, L. Doherty (2008), TSMP: Time synchronized mesh protocol, *Parallel and Distributed Computing Systems*, 631.

Porter, J., P. Arzberger, H.-W. Braun, P. Bryant, S. Gage, T. Hansen, P. Hanson, F.-P. Lin, T. Kratz, W. Michener, S. Shapiro, and T. Williams (2005), Wireless sensor networks for ecology, *BioScience*, 55(7), 561—572.

Ramanathan, N., T. Harmon, L. Balzano, D. Estrin, D. Hansen, J. Jay, W. Kaiser, and G. Sukhatme (2006), Designing wireless sensor networks as a shared resource for sustainable development, *Information and Communication Technologies and Development*, 256-265, doi: 10.1109/ICTD.2006.301863.

Rice, J., and B. Spencer (2008), Structural health monitoring sensor development for the imote2 platform, *SPIE Smart Structures Conference/NDE*.

Rice, R., and R. C. Bales (2010), Embedded-sensor network design for snow cover measurements around snow pillow and snow course sites in the Sierra Nevada of California, *Water Resour. Res.*, *46*, W03537, doi:10.1029/2008WR007318.

Rice, R., R.C. Bales, M.W. Meadows, B. Kerkez, S.D. Glaser, M. Anderson, D.G. Marks, A. Mazurkiewicz, J. Dozier, B.J. McGurk (2009), Design and implementation of a snow measurement network using ground-based wireless networks and space-bore measurements in the American River Basin of California, *American Geophysical Union,* Abstract C33A-0504.

Rice, R., R.C. Bales, T.H. Painter, J. Dozier (2011), Snow cover along elevation gradients in the upper Merced and Tuolumne river basins of the Sierra Nevada,*Wat. Resour. Res.*, 47, W08515, doi:10.1029/2010WR009278.

Rittger, K. (2012). Spatial estimates of snow water equivalent in the Sierra Nevada for streamflow foresting. Santa Barbara, CA, University of California. PhD.

Rittger, K., A. Kahl, and J. Dozier (2011), Topographic distribution of snow water equivalent in the Sierra Nevada, *Annual Proceedings of the Western Snow Conference, Lake Tahoe, NV,* 79: 37-46.

Rittger, K., T. H. Painter, and J. Dozier (2012), Assessment of methods for mapping snow cover from MODIS, *Advances in Water Resources*, Elsevier Ltd. doi:10.1016/j.advwatres.2012.03.002

Schaake, J. C., & Peck, E. L. (1985). Analysis of water supply forecast accuracy. In *Proceedings of the 53rd annual Western Snow Conference* (pp. 44-53).

Sellers, S. (2000), Theory of water transport in melting snow with a moving surface, *Cold Regions Science and Technology*, 31:47{57.

Sellers, S. (2001), Water transport in phase-changing snowpacks. *IUTAM Symposium on Theoretical and Numerical Methods in Continuum Mechanics of Porous Materials*, 229{236.

Sentilla (2011), Enterprise management for energy, *http://www.sentilla.com/*, January 2012.

Serreze, M. C., M. P. Clark, R. L. Armstrong, D. A. McGinnis, and R. S. Pulwarty (1999), Characteristics of the western United States snowpack from snowpack telemetry (SNOTEL) data, *Water Resour. Res*., 35, 2145–2160, doi:10.1029/1999WR900090.

Shafer, B. A., and Huddleston, J. M. (1984, Analysis of seasonal volume streamflow forecast errors in the western United States. *Proc. A Critical Assessment of Forecasting in Water Quality Goals in Western Water Resources Management,* Bethesda, MD, American Water Resources Association, 117–126.

Shamir, E., and K. P. Georgakakos (2006), Distributed snow accumulationand ablation modeling in the American River basin, *Adv. Water Res., 29,*558–570, doi:10.1016/j.advwatres.2005.06.010.

Shannon, C.E. (1948) A mathematical theory of communication, *Bell System Technical Journal, 27,* 379–423, 623–656.

Sierra Club, Water issues in California, http://www.sierraclub.org/ca/water/index.asp, last accessed November 10, 2012.

SNOTEL data and products, http://www.wcc.nrcs.usda.gov/snow/, last accessed November 10, 2012.

Song, J., S. Han, A. Mok, D. Chen, M. Lucas, M. Nixon, and W. Pratt (2008), Wirelesshart: Applying wireless technology in real time industrial process control. *Real Time and Embedded Technology and Applications Symposium 2008*, 377–386.

State of California, Department of Water Resources (2009), California cooperative snow surveys, http://cdec.water.ca.gov/snow/, Accessed May 7th, 2009.

Szewczyk, R., E. Osterweil, J. Polastre, M. Hamilton, A. Mainwaring, and D. Estrin (2004), Habitat monitoring with sensor networks, *Commun. ACM*, 47. 34-40.

Tanenbaum S. (2003), *Computer Networks, ed. 4*. Prentice Hall, New Jersey, NJ.

Tanya Hoogerwerf (2012), SAC-SMA, presentation, last accessed December 2012.

Tomlin, C., J. Lygeros, and S. Sastry (2000), A game theoretic approach to controller design for hybrid systems, *Proc. IEEE*, 88(7):949{970.

Trevor J.. Hastie, Tibshirani, R. J., & Friedman, J. H. (2009). *The elements of statistical learning: data mining, inference, and prediction*. Springer.

Trubilowicz, J., K. Cai, and M. Weiler (2009), Viability of motes for hydrological measurement, *Water Resour. Res.*, *45*, W00D22, doi:10.1029/2008WR007046.

Tseng, P., and T. Illangasekare (1994), Modeling of snowmelting and uniform wetting front migration in a layered subfreezing snowpack, *Water Resources Research*, 30.

Turns, S. (2006), *Thermal- Fluid sciences: An integrated approach*, Cambridge University Press.

Varhola A., J. Wawerla, M. Weiler, N. C. Coops, D. Bewley, Y. Alila (2010), A New Low-Cost, Stand-Alone Sensor System for Snow Monitoring. *J. Atmos. Oceanic Technol.*, 27, 1973–1978.

Vereecken, H., T. Kamai, T. Harter, R. Kasteel, J. Hopmans, and J. Vanderborght (2007), Explaining soil moisture variability as a function of mean soil moisture: A stochastic unsaturated flow perspective, *Geophys. Res. Lett.*, *34*, L22402, doi:10.1029/2007GL031813.

Wackerly, D., W. Mendenhall, and R. Scheaffer (2008*), Mathematical Statistics*, ed. 3, Thomson Brooks/Cole.

Watteyne T., X. Vilajosana, B. Kerkez, F. Chraim, K. Weekly, W. Qin, S.D. Glaser, K.S.J. Pister (2012), OpenWSN: A Standards-Based Low-Power Wireless Development Environment, *Transactions on Emerging Telecommunications Technologies*.

Watteyne, T., A. Mehta, K. Pister (2009), Reliability through frequency diversity: why channel hopping makes sense, *PE-WASUN '09: Proceedings of the 6th ACM symposium on Performance evaluation of wireless ad hoc, sensor, and ubiquitous networks*, New York, NY, USA, 116—123.

Watteyne, T., S. Lanzisera, A. Mehta, K. Pister (2010), Mitigating Multipath Fading Through Channel Hopping in Wireless Sensor Networks, *IEEE International Conference on Communications(ICC)*, Cape Town, South Africa.

Western Education Foundation (WEF), A briefing on California Water Issues, whitepaper, http://www.watereducation.org/, last accessed November 10, 2012.

Western, A.W., R.B Grayson, and G. Bloschl (2002), Scaling of soil moisture: a hydrologic perspective, *Annual Review of Earth and Planetary Sciences,* 30: 149–180.

Williams, C. J., J. P McNamara, and D.G. Chandler (2009), Controls on the temporal and spatial variability of soil moisture in a mountainous landscape: the signature of snow and complex terrain, Hydrol. Earth Syst. Sc. 13(7): 1325-1336.

Wu, X., and M. Liu (2012), In-Situ Soil Moisture Sensing: Optimal Sensor Placement and Field Estimation. *Proceedings of the 11th international conference on Information Processing in Sensor Networks (IPSN '12),* ACM, New York, NY, USA, 1-12., doi:10.1145/2185677.2185

ZigBee Alliance (2008), ZigBee Specification, *http://www.zigbee.org/Standards/*, January 2012.