

Spatially- and Temporally-Adaptive Communication Protocols for Zero-Maintenance Sensor Networks Relying on Opportunistic Energy Scavenging

Xuejing He
EECS Department
University of Michigan
1301 Beal Ave
Ann Arbor, MI, 48019
xjhe@umich.edu

Robert P. Dick
EECS Department
University of Michigan
1301 Beal Ave
Ann Arbor, MI, 48019
dickrp@eecs.umich.edu

Russ Joseph
EECS Department
Northwestern University
2145 Sheridan Road
Evanston, IL, 60208
rjoseph@eecs.northwestern.edu

ABSTRACT

Wireless sensor networks allow scientists to gather data from remote, difficult to access, and dangerous locations. However, maintenance of aging networks and removal of obsolete or inactive nodes containing toxic materials is expensive and time consuming. Moreover, node lifespan is generally constrained by the reliability of the batteries used in most deployments, especially in the presence of extreme variation in environmental conditions such as temperature and humidity. We consider the problem of designing wireless sensor networks capable of indefinite deployment periods measured in decades, not months. We describe the architectural and capability implications of eliminating batteries from sensor networks and instead relying on opportunistic energy scavenging. Sensor nodes using ambient energy sources become temporarily active at unpredictable but possibly correlated times. In this paper, we use wind power as an example of such a power source, which we model using temporally and spatially correlated random processes. Such models can be built using historical measurements over a geographical range. We describe a method to use energy models in the design of latency-optimized and cost-constrained battery-less wireless sensor networks, and explain the required changes to network architecture, communication protocol, and node hardware. In the context of environmental monitoring applications, we compare the performance of a network designed and managed using our techniques with that of existing design styles.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Distributed networks*; C.3 [Computer Systems Organization]: Special-purpose and application-based systems—*real-time and embedded systems*

Keywords

energy scavenging, routing protocol

This work was supported in part by the NSF under awards CNS-0720820 and CNS-0721978.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CODES+ISSS'12, October 7–12, 2012, Tampere, Finland.
Copyright 2012 ACM 978-1-4503-1426-8/12/09 ...\$10.00.

1. INTRODUCTION

Long unattended lifespans are important for wireless sensor networks because they are often deployed in locations that are difficult to access. Replacing sensor network node components or retrieving nodes can be prohibitively expensive. Therefore, sensor node lifetime is of central importance. Most sensor nodes last a few months; some last a few years. This implies that component or node replacement is necessary in long-term deployments. Our work provides a novel method of designing wireless sensor networks to operate for decades without periodic repair or replacement of sensor nodes.

Several factors constrain sensor node life time. The battery lifespan constraint is typically encountered first. Rechargeable batteries used in many sensor nodes wear out faster than other components. Even when energy constraints are loose, batteries have lifetimes ranging from 1.5 year to 6 years [1, 2]. Battery lifetime is even worse among sensors deployed in outdoor environments because the batteries suffer from ambient temperature and humidity changes, which can degrade the battery lifespans by 75% [2]. Frequent depletion/recharge cycles can result in 20% decrease in battery lifespans. Consequently, sensor nodes must periodically have their batteries replaced, wasting time and effort. Even for sensor nodes that use batteries with lifespans exceeding that of the application, it is still commonly necessary to eventually collect the nodes because many batteries contain toxic materials such as heavy metals [3]. The elimination of batteries opens the possibility of designing wireless sensor nodes that are suitable for long-term, one time deployment. Battery-less wireless sensor nodes using energy scavenging devices as power supplies can meet this requirement.

Eliminating batteries can potentially make long-term deployment maintenance free. However, existing battery-less sensor network designs still have problems. Previous works do not eliminate the need for lifetime-constraining energy storage devices [4, 5]; such nodes use supercapacitors or battery-supercapacitor hybrid systems for energy storage, which still constrain the sensor node lifespan [6, 7]. In contrast, we consider a wireless sensor node design that completely eliminates the need for lifetime-constraining energy storage devices. The sensor nodes use ambient power sources, such as solar power, wind power, and water flow as their energy source, replacing the battery with long-term, stable energy scavenging devices. As a result, the sensor nodes no longer suffer under batter-imposed lifetime constraints.

Eliminating batteries requires changing the sensor node architecture and modifying the sensor network protocol design to specifically adapt to the changes in sensor node activities caused by using ambient power sources: sensor nodes wake up at imperfectly predictable times. Therefore, existing protocols that rely on pre-

scheduled data transmission do not work. Storage and time synchronization constraints in the sensor node architecture design also limit the use of battery-less energy scavenging. We propose a protocol that takes power source availability and remaining memory into consideration.

We consider applications in which the sensing target moves infrequently, as is often the case for long-term environmental monitoring. They fall into one of the two categories: (1) the wireless sensor network only needs to sense when events occur, and these events also provide energy or (2) sensor nodes are deployed in an environment that provides access to time-varying energy sources that are event-independent. These properties are commonly seen in existing distributed sensing applications.

We propose the following modifications to existing wireless sensor network architectures: (1) replace the power supplies of sensor nodes with energy scavenging devices, which may be wind or water turbines, piezoelectric generators, or solar panels, depending on the application and available energy sources; (2) adapt routing decisions based on the spatial and temporal distributions of power availability for nodes; and (3) store intermediate results to non-volatile memory, when appropriate, to compensate for loss of power.

Based on these changes, we describe a new sensor network design that is well suited to energy scavenging. Our major contributions follow.

- We describe a novel routing protocol that works well for sensor networks using battery-less energy scavenging. The routing protocol reacts to imperfectly predictable changes in ambient energy sources. It minimizes the end-to-end latency of packet transmission and achieves $1.3\text{--}3\times$ performance improvement over existing designs for four environmental sensing applications.
- We describe architecture changes to sensor nodes that make them more appropriate for use with intermittent and imperfectly predictable power supplies.
- We categorize commonly used sensor network applications and provide application dependent guidance for designers considering battery-less energy scavenging.

We will discuss the details of our design in later sections and compare it with existing design strategies.

2. RELATED WORK AND MOTIVATION

Large scale, long-term monitoring applications rely on low-maintenance sensor nodes because in-field repairs and replacement are expensive. Energy scavenging nodes offer a compelling solution for reducing maintenance costs; they gather ambient energy from the environment and consequently eliminate the need for battery replacement. However, existing energy scavenging nodes still have practical limitations that either constrain their lifetimes or prevent them from being used to build large-scale network applications. In this section, we point out shortcomings in existing design techniques, summarize our contributions, and argue for new sensor node and network architectures.

2.1 Energy Scavenging with Battery Assistance

Many researchers have proposed using sensor nodes powered by a combination of rechargeable batteries and scavenged energy including power derived from the sun, ambient vibration, wind, water flow, and the motion of animals. Among these, solar-power is most widely used due to its high and stable energy density.

Raghunatha et al. [8] describe a procedure for designing efficient solar-powered sensing systems. Taneja et al. [9] provided network architecture and node design guidelines for micro-solar powered sensor networks. Researchers have also developed routing protocols suitable for solar-powered sensor networks. Voigt et al. [10]

proposed and compared two such protocols. Existing work has used solar-powered sensors for environmental monitoring. Mainwaring et al. [5] developed a wireless sensor network using solar-powered Mica Motes for habitat monitoring on Great Duck Island. They used solar panels that can provide between 60 and 120 Watts in full sunlight. The sensor networks in this paper used rechargeable batteries to store scavenged energy. However, the batteries themselves constrain the sensor node lifetime and hence limit the applicability of these nodes in very long deployments.

2.2 Why Battery-Less Energy Scavenging?

Batteries are the primary energy storage devices in many sensor nodes. However, they typically have short lifespans and frequently limit the lifespan of the whole node. Consequently, researchers have proposed the following methods of eliminating them.

Replace the rapidly degrading rechargeable battery with a supercapacitor. Minami et al. designed Solar Biscuit [4], a battery-less sensor network that only relies on a solar panel and a supercapacitor to power the sensor node. They also developed a routing protocol suited to the long charging time of the supercapacitor. The use of supercapacitors can extend the lifespan of sensor nodes, but only to a point; supercapacitors also degrade. Studies have reported that supercapacitors have 13%–15% capacity degradation and double their internal resistance after one year of power cycle testing [6, 7]. Moreover, the lifetimes of supercapacitors are temperature-dependent. The expected lifetime halves with every 10° Celsius increase of ambient temperature [11, 12]; supercapacitors degrade when deployed outdoor. In addition, supercapacitors are $4\text{--}10\times$ more expensive than rechargeable batteries with the same energy capacities and densities. Supercapacitors remain unsuitable for our goal of long-term deployment, although this may change in the future if their reliable lifespans are increased.

Reduce the number of charge/discharge cycles in the battery. This can be done by attaching a supercapacitor or energy scavenging device to the battery, and only discharging the battery when the other energy supplies fail. Jiang et al. proposed a multi-stage energy transfer system that uses a solar panel together with a supercapacitor as the first stage and a rechargeable battery as the second stage [13]. They argued that, when ambient power is sufficient to power the sensor node, the system can avoid discharging or charging the battery. This approach can increase the sensor node lifetime to 4 years given a 10% duty cycle. However, it does not eliminate batteries, which must eventually be gathered from the environment. We argue that the sensor node lifetime problem should be solved by removing the battery entirely.

A few researchers have proposed battery-less sensor nodes with the goal of increasing sensor node lifetime. Philipose et al. [14] attached an RFID to a battery-less sensor node, powering the sensor node via the RFID reader. Their work completely removed the energy storage device. Vyas et al. [15] and Patel et al. [16] combine a battery-less, wireless tag and a low power sensor node for use in a passive sensor. However, these works require that energy be directed to each active sensor from an external radio frequency energy source. This prevents use in distributed applications. Ng et al. [17] design a near-body network with battery-less wearable biomedical sensors to monitor patient physiological state. This solution is appropriate for body-range transmission and consists of only a few nodes. Our work focuses on applications requiring larger scale distribution of sensors.

2.3 Node Design and Protocol Support

Existing routing protocols and sensor node architectures are not well suited to energy scavenging sensor networks. We propose an

architecture that is based on partial knowledge of the spatial and temporal properties of ambient power sources.

The designers of energy scavenging sensor networks face special difficulties in maintaining functionality and performance. Many environmental power sources, such as solar, wind, and water flow have intermittent availability. This complicates routing protocol and node architecture design. First, the scavenging sensor network is dynamic: its connectivity structure changes dynamically depending on environmental power source status of each node. Routing protocols must adapt to these changes. Second, the wake-up schedule of sensor nodes cannot be controlled by the designer. Algorithms that rely on coordinated activations at pre-determined times to sense, transmit, or receive cannot be used. Third, sensor nodes lose their power sources at imperfectly predictable times, leaving little time for nodes to react by transmitting data or preserving it in non-volatile memory. The designers of energy scavenging wireless sensor networks must consider these domain-specific challenges.

How well would existing sensor node and communication protocol designs fare in a battery-less environment? Prior work has proposed flooding-based routing techniques for energy scavenging sensor networks [4, 8]. These protocols require nodes to wake up at pre-determined times and use redundant transmissions to compensate for the lost messages. They are adequate for small-scale networks. However, they would perform poorly in larger-scale applications such as environmental monitoring. In medium- to large-scale networks, flooding overwhelms the communication channels, resulting in high latencies and data loss rates. Geographic routing is another popular candidate protocol; it is easy to implement and may not require pre-determined schedules for transmissions. However, only using geographic information for routing in energy scavenging sensor network can cause packets to be trapped in inactive nodes (as described in detail in Section 3). The main weakness of existing geographic routing protocols in this application is their failure to account for the fact that nodes frequently become unavailable at imperfectly predictable times and some nodes are available more frequently than others. This limits network scale, and prevents operation when many nodes are frequently inactive. Other candidate protocols require scheduling nodes to transmit at precise times, and therefore cannot be used in battery-less energy scavenging networks.

3. PROBLEM DEFINITION

Our goal is to provide routing protocol and node design techniques suitable for indefinitely deployed sensor networks. We begin by eliminating the use of energy storage devices with highly constrained lifespans. Given temporary losses in node power, we attempt to determine the design techniques yielding the highest end-to-end successful data delivery rate under a (designer-specified and application-specific) constraint on acceptable latency.

We now describe our model of an energy-scavenging sensor network. Ideally, when all nodes have access to sufficient power, they form a connected graph N containing $|N|$ nodes, in which there is a directed edge between two nodes if the first can successfully transmit directly to the second. In this situation, nodes can transmit sensed data to the base station using existing routing protocols. However, the probability of all nodes concurrently having power at any particular time is small. In each time interval, only a subset of sensor nodes have enough scavenged power to operate. The graph of these nodes, N_i , is referred to as the *active subset* for the i th time interval and does not change within the interval. Note that intervals can be defined to end whenever the active subset changes.

We model the network packet transmission latency using active subsets. It is likely that two temporally adjacent active subsets

N_i and N_{i+1} have a non-empty intersection due to the temporal correlation of the power source (described in detail in Section 5). As shown in Figure 3, packets from a faraway node S can travel through multiple active subsets N_{i0} through N_{i3} via their intersections to finally arrive at B.

The packet transmission delay is the sum of $t_{trans,i}$ and $t_{int,i}$ for all active subsets along the path. $t_{trans,i}$ is the transmission latency to populate packets within an individual active subset N_i . $t_{int,i}$ is the time interval between active subset N_i and the next active subset N_{i+1} , which is the time when sensor nodes are inactive.

We define the latency of data transmission in the network to be the time required for all nodes to send their sensed data packets to the base station, i.e., the maximum packet transmission delay. In reality, some packets will be dropped due to channel overuse or collision. Others




Figure 1: Battery-less energy scavenging sensor network. A node S transmits its packets to base station B through temporally intersecting active subsets.

will be trapped in nodes that wake up infrequently. Thus, we define *packet delivery rate* to be the percentage of packets that reach the base station at a particular time. When we compare protocols in later sections, we will compare transmission latencies associated with particular packet delivery rates.

The limited predictability of ambient power sources reduces designer control of wireless sensor networks. It prevents the designer from using pre-computed routing paths and requires routing protocols that adapt to changes in the ambient power source. Pre-computed routing may falsely send packets to nodes that are not in the temporal intersection of two adjacent active subsets, preventing packets from further transmission. As can be seen in Figure 3, node S may send its packet to another node C, which is geographically closer to the base station B but is in a rarely awake, isolated subset. This can delay packet transmission; the 4-hop path via N_{i0} through N_{i3} is faster. Instead, we consider the temporal and spatial statistics of ambient power sources to dynamically change the routing for every active subset and avoid trapping packets.

We will further describe the design and implementation of our battery-less wireless sensor network routing protocol in the following sections.

4. DESIGN AND IMPLEMENTATION

We have developed a routing protocol for sensor nodes that become active at imperfectly predictable times. Sensor node activation events occur at random times prescribed by the temporally and spatially correlated statistical processes used to characterize ambient energy sources. Therefore, nearby nodes have similar behavior. We describe a protocol informed by these properties, and then discuss four variations of it appropriate for a range of sensing applications.

We will first describe a network architecture suitable for energy scavenging environmental monitoring applications. The sensor nodes are widely deployed, and a group of sensor nodes share a base station. The nodes transmit their sensed data to the base station, potentially indirectly via other nodes. The base station uses a high-capacity power supply that will require replacement every

few years at much lower cost than recovering all sensor nodes. In the remainder of this section, we describe how we perform routing given this network structure.

4.1 Sensor Node Architecture

We modify the sensor node architecture to support battery-less energy scavenging and to guarantee that the sensed data contain valid time stamps.

4.1.1 Energy Scavenging Devices

We consider two major modifications to a conventional wireless sensor network node: (1) removing the battery and (2) attaching an energy scavenging device, e.g., a solar panel, wind turbine, water turbine, or piezoelectric device [18]. Our design relies on non-volatile memory to save states between active intervals. Most existing sensor nodes contain non-volatile memory. Depending on the power source distributions, it may also be appropriate to use a higher-performance processor and network interface than is typical, in order to quickly finish processing and communication tasks when power is available.

4.1.2 Time Synchronization

Sensor nodes must generally associate times with samples. The local timer in a sensor node stops working when power is lost, and it needs to be re-synchronized when power is available. Activation events happen at random times, meaning that sensor nodes cannot determine how long their timers have been inactive when reactivated. A node can synchronize with the base station or its neighbors, which have valid local timers. However, this synchronization still has some delay because a node can be synchronized only when it is in the same active subset with other nodes maintaining correct times. The samples gathered during the synchronization delay will have incorrect time stamps.

We now describe a technique to compensate for time uncertainty in intermittently powered sensor nodes. Nodes are desynchronized by power loss. They attempt to resynchronize with neighboring nodes, but until that time they mark data samples with time ranges that are later used by the base station, together with other sample time stamps, to more accurately estimate when the data were gathered. This approach can achieve time stamp errors of less than 80 minutes for 91% packets in a medium scale (500 nodes) sensor network. The environmental monitoring applications we consider usually gather samples several times a day (as described in Table 1). Therefore, samples with time stamp errors of minutes or even hours are acceptable.

Our proposed process works in the following steps.

1. The local timer of a sensor node contains an invalid value at the beginning of deployment. Every sensor node will first try to synchronize with the base station to obtain the correct time.
2. The sensor node refreshes its local timer value (stored in a non-volatile memory) every fixed time interval (e.g., one minute) and at every time it is synchronized to nodes with correct timers. When a sensor node again has access to power, it restarts its local timer using the stored time stamp in the memory.
3. Every node attaches a node identifier and a unique packet identifier maintained in non-volatile memory to each packet, guaranteeing that for a particular node, packets with smaller packet identifiers are always produced earlier than the ones with larger packet identifiers. Packets also carry time stamp upper bounds (initialized to the latency constraint) and lower bounds (initialized to the value of the local timer). A valid bit is also included, which is set to “true” only when the local timer is known to have low error at the time of packet generation.

4. The base station refines the time stamps of packets, working within a fixed time interval. This time interval should be long enough such that most packets generated at a similar time arrive at the base station before the end of the interval. The generation time of a packet (with identifier p) is estimated by examining the packets with the closest smaller and larger packet identifiers (noted as l and u). If these packets have valid time stamps, the upper and lower bounds for the packet of interest are refined as follows:

$$\begin{aligned} packet_{p.lower} &= packet_{l.upper}, \\ packet_{p.upper} &= packet_{u.lower}. \end{aligned} \quad (1)$$

4.2 Precomputed Information

Our routing technique bases decisions on temporal power source distributions. Each node knows the probability of having sufficient power for computation and transmission and its distance to the base station. This information can either be precomputed from the power source distribution or gathered after deployment.

The *activity rate* P_{active} of a sensor node is the probability of it having sufficient power to compute and transmit data [18]. If a sensor node is powered by a wind turbine, it will only be activated when the wind speed exceeds a threshold value providing enough power. The probability of a sensor node being awake is $\int_0^{w_{th}} f(x, \lambda, k)$, where w_{th} is the threshold wind speed and $f(x, \lambda, k)$ is the wind speed distribution at that location. The node activity rate can be computed from historical wind speed distribution data.

The distance d from the sensor node to the base station is the number of wireless communication hops, assuming all nodes are active. This distance can be gathered during network deployment.

4.3 Routing Protocols

In order to determine whether existing protocols are sufficient for indefinitely deployed energy scavenging sensor networks, we make comparisons between several existing routing protocols: simple flooding, geographic routing, buffer size dependent routing, and undirected routing. We also evaluate the Ambient Energy Aware routing protocol we designed specifically for this problem.

Flooding is the most commonly described routing protocol for energy scavenging sensor networks [4, 8]. It is easy to implement in sensor nodes with limited computation power and has adequate performance for small-scale networks in which the data generation rate is low. Multiple nodes keep copies of the same packet; thus, even when some of the nodes lack power and become inactive, the redundant copies of the packet are transmitted by other nodes. In larger networks, flooding faces two problems: (i) limited resources and (ii) unidirectional transmission. Simple flooding creates redundant packets that can exceed sensor node memory capacities resulting in dropped packets. In addition, flooding protocols suffer greatly from limited channel capacity. The network-wide channel capacity is constrained in battery-less energy scavenging sensor networks by the possibly frequent deactivation of nodes that are (temporarily) without power. Flooding creates duplicate packets and easily overwhelms the network. This is especially problematic when nodes wake up infrequently.

Based on the above observations, we now consider routing protocols appropriate for battery-less energy scavenging sensor networks. These protocols aim to avoid the poor performance caused by limited buffer size, packet collision, and the randomness of node active intervals. These protocols have the following characteristics.

- *Acknowledgment.* We enable acknowledgment by both sender and receiver nodes. Receivers acknowledge packet acceptance. When a sender receives the first acknowledgment, it broadcasts a drop request to its neighbors, allowing all but the node that trans-

mitted the first acknowledgment to drop their copies of the packet. This avoids unnecessarily use of memory and communication resources for duplicate packet copies in multiple nodes, while preserving at least one copy of the packets. The acknowledgment delay described later reduces the probability of acknowledgments collision.

- *Directional transmission.* Packets transmit along the path with the smallest expected latency to reach the base station. When receiving a packet, instead of acknowledging immediately, a receiver use a *ranking function* to delay the acknowledgment. The time delay is set to give priority to nodes with higher probability of successfully reaching the base station. The choice of ranking function is a key design feature. Later in this section, we will discuss the selection of the ranking function in greater detail.

- *Random Back-Off.* Nodes perform random retransmission back-off on packet collisions to avoid future collisions.

Given these starting properties, we consider several candidate routing protocols. Three of the protocols (geographic routing, buffer size dependent routing, and undirected) adapt well-known algorithms to the energy scavenging sensor networks considered in this work.

Geographic routing [19]. Nodes always accept packets from the neighbors that are geographically further away from the base station. This increases the probability of the packet reaching the base station. The ranking function is $delay_i = t_{unit} \cdot d_i$, in which $delay_i$ is the delay of the i th node to acknowledge, t_{unit} is a unit time period, and d_i is the number of hops from the i th node to the base station, assuming all nodes are active. As mentioned earlier in this section, this distance d can easily be gathered during node deployment. One significant drawback of this approach is the likelihood of creating holes in the network: some nodes are geographically closer to the base station, but rarely active. Packets will sometimes be transmitted to these nodes shortly before they become active and then remain trapped for a long time.

Buffer size dependent routing. Nodes accept packets from neighbors with less free space in their message buffers. This avoids buffer overflows, which may result in data loss. The ranking function is $delay_i = t_{unit} \cdot (b_{max} - b_i)$, where b_{max} is the maximum buffer size and b_i is the remaining size in the node receiving buffer. Unfortunately, this protocol does not consider the importance of transmission directions. Nodes that frequently wake up or are closer to the base station receive more packets. As a result, the ranking function will assign these nodes longer delays, forcing packets to be forwarded to nodes from which packets are less likely to reach the base station.

Undirected protocol [19]. Nodes are assigned random priorities using ranking function $delay_i = t_{unit} \cdot random(n_i)$, where $random(n_i)$ returns a random number between 1 and n_i , the number of one-hop neighbors of node i . This protocol has the benefit of simplicity but is usually inefficient because random prioritization results in slow, indirect paths for many packets.

The final protocol is a novel approach which is specifically designed to overcome challenges in battery-less sensor networks:

Ambient Energy Aware protocol. Nodes with the highest probabilities to be (possibly indirectly) connected to the base station have the highest probabilities of accepting packets from their neighbors. This protocol makes use of the statistical data on power source availability and the node activity rate (as described in Section 4.2) to compute the ranking function. Our goal is to combine the best attributes from the four protocols described above, while also using available information on the statistical properties of the ambient power source. The ranking function is $delay_i = t_{unit} \cdot (b_{max} - b_i) \cdot d_i / P_{active,i}$, where b_{max} is the maximum buffer

size, b_i and d_i are the remaining buffer size and distance from the base station of node i , and $P_{active,i}$ is the activity rate of node i . The drawback of this protocol is its requirement for additional memory on sensor nodes to store power source statistical properties. Fortunately, this is not a problem in practice because the statistical data can be preprocessed and reduced to a single number: the node activity rate.

The most appropriate protocol depends on application characteristics. Therefore, we will compare the existing and new protocols described above under a variety of operating conditions.

5. EVALUATION

In this section, we provide evaluation results for the protocols discussed in the previous section. We first evaluate the protocols when used in different sensing applications. Then we determine the sensitivity of protocol quality metrics to variations in application characteristics.

5.1 Experiment setup

Our evaluation considers sensor networks that scavenge energy from wind. To model changes in wind conditions, our simulation takes location-dependent time-varying wind traces as input. However, the raw measured wind speed data for large regions and long duration are not publicly available. Therefore, we generate similar wind traces with statistical properties based on recorded wind speed distributions [20]. Wind speed traces are then fed into the discrete-event simulator described later in this section.

Our generated wind speed traces have the following properties. First, the trace for each particular location has particular temporal correlation values. Second, traces at different locations have spatial correlation. We can represent the wind speed traces by a group of correlated Weibull random variables. To generate these random variables, we make use of two sources of information: (1) a regional wind speed atlas [20] and (2) spatial and temporal correlation models.

The World of Wind Atlases houses a publicly available archive of wind data from many regions around the world [20]. For a given region, the atlas logs detailed location-specific information about wind patterns. Specifically, for each location on the map, the atlas records the Weibull parameters (k, λ) that describe variation in wind speed at several altitudes. We select the wind atlas of one island in Denmark as the source of wind data, because its wind atlas is representative of many coastal regions in the world.

We use existing wind speed spatial correlation models [21–23]. The spatial correlation coefficient for wind speed at two different locations is exponentially dependent on the distance between the locations: $c = \exp(-d/d_0)$, where d is the distance between the locations and d_0 is called *correlation distance*. This is the distance at which the correlation between two locations equals $\exp(-1)$. We model the temporal correlation coefficient as exponentially dependent on elapsed time based on the observations of Archer and Jacobson [21].

We use a wind trace duration equal to the maximum tolerable latency for data transmission. In environmental monitoring applications, it is common for the base station to processes or transmit data in a daily or weekly pattern [5, 24]. Some packets will carry incorrect time stamps due to the intermittent power loss of the sensor nodes. Their time stamps will be refined after reaching the base station using the method described in Section 4, with the largest error being the packet transmission latency. Thus, packets arriving at the base station later than one week are likely to have indistinguishable time stamps. They are considered to be invalid. We set our latency constraint to be one week.

We developed a discrete event wireless sensor network simulator in which changes in wind speed are events. We generate wind speed traces based on the parameters described above [20, 21] and provide them to our simulator. The simulator models the activation and deactivation of sensor nodes when the input wind speeds at the sensor locations change, forming different active subsets. Within each active subset, it simulates the behavior of sensing and data transmission of sensor nodes executing any of the four protocols described in Section 4. The network packet delivery rate, node level channel utilization, and transmission latency are recorded.

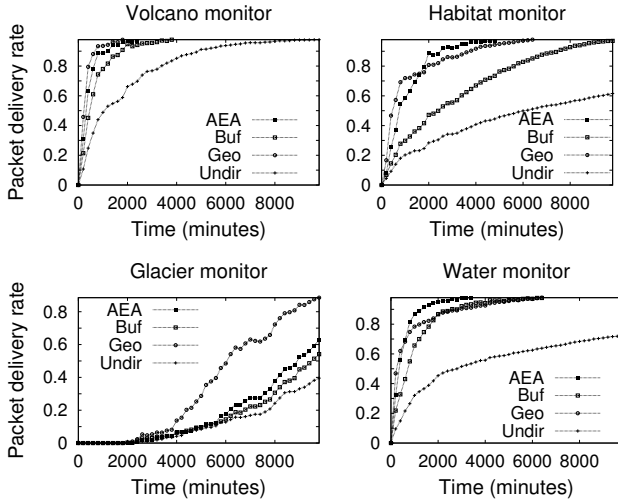


Figure 2: Packet delivery rate comparison for four protocols under different applications.

5.2 Application Based Evaluation

Several sensor network applications are suitable for using energy scavenging techniques. Table 1 lists these applications and their properties [5, 24–30]. In the context of these applications, we evaluate four protocols that were discussed in Section 4: Ambient Energy Aware routing (*AEA*), geographic routing (*Geo*), buffer size dependent routing (*Buf*), and undirected routing (*Undir*). Our results show that no single protocol is best for all applications.

We consider four wireless sensor network applications from Table 1 as our examples for energy scavenging sensor network: habitat monitoring (*habitat*) [5], volcano monitoring (*volcano*) [24, 25], glacier monitoring (*glacsweb*) [26], and meteorology and hydrology monitoring (*water*) [27]. Each application has a set of parameter values that are used in their setup, as described in Table 1. These four applications have distinct parameters. *habitat*, *volcano* are both small- or medium-scale sensor networks (10–100 nodes evenly distributed in an area with 2 km radius) in which sensor nodes require a moderate amount of ambient power (100–200 mW), and only perform data sensing several times a day. *water* is a large-scale network (600 nodes) with moderate power requirement. *glacier* is a medium-scale sensor network but requires high power supply to sensor nodes. We first evaluate our protocols on these four applications with fixed parameter values listed in Table 1.

Figure 2 shows the time-dependent variation in packet delivery rate throughout the whole network for these four applications. Based on the application constraints described in Section 5.1, we set the maximum tolerable packet transmission latency for all packets to be one week. While some applications (*habitat* and *volcano*) have acceptable delivery rates using simpler routing protocols, oth-

ers (*water*) require our proposed Ambient Energy Aware Protocol to achieve satisfactory performance (i.e., packet latencies less than a week). For *glacier*, none of the protocols considered achieve the packet delivery rate within the latency constraint.

This difference in performance of different protocols can be briefly explained by reference to the abstract model described in Section 3. First, let us consider the *habitat* and *volcano* applications. They share common properties: small- or medium-scale and medium power requirement. The first property means that there are always sufficient sensor nodes active at the same time. This results in an active subset N_i that covers a large portion of the network. The second property guarantees that the sensor nodes are likely to have enough memory to store the sensed data without dropping packets. As a result, using simpler approaches such as geographic routing is sufficient. Second, we consider *water*, which has a large-scale sensor network and wider distribution of sensor nodes. This results in a large active subset N_i that has the potential to overwhelm the network with traffic. It is therefore favorable to transmit packets through nodes that are more likely to be active, rather than concentrating the traffic close to the base station. For this application Ambient Energy Aware routing works best. Third, *glacier* has a high node power consumption requirement. Only very high wind speed can provide enough power to activate a node. The resulting active subset N_i is a sparse network, making it less likely to cover the whole network for a given latency constraint. As a result, none of the protocols considered have good performance for this application. To improve the situation, one might switch to a more powerful energy scavenging device, e.g., a larger turbine.

5.3 Delivery rate for varying parameters

No single protocol is best for all application scenarios. To assist application developers to select the most appropriate protocol, we now show the application-dependent parameter ranges for which each protocol is best suited.

From the example applications described above, we observe several variable parameters, which are subject to changes due to special requirements of the application or user preference. These parameters include: (1) network scale, which is the total number of nodes in the network; (2) sample size, which is the size of data gathered at every sensing event; (3) required power consumption, which is the maximum power required for a sensor node to perform data processing and transmission; and (4) maximum direct transmission range.

We evaluate how each parameter affects selection of a protocol using a series of parameter studies. The parameter studies are conducted by varying parameters one at a time while keeping others constant. We select a set of constant values for these parameters: medium scale network (300 nodes distributed in a 3 km×3 km region), small sample size (16 B), medium required power consumption (100 mW), and transmission range for commonly used nodes (400 m).

The properties of the sensor network influence the selection of protocols. On the one hand, in large-scale sensor networks, the edge of the network is far from the base station, requiring more hops to send packets to the destination. Since the wind speed at each sensor node location varies randomly over time, the probability of transmitting a packet from a distant node to the base station depends on the node activity rate along the transmission path. It will often be best to route through a longer path with higher activity rate, rather than the shortest path. Ambient Energy Aware routing uses knowledge of ambient power source statistical parameters, giving it an advantage over other routing techniques in large-scale sensor networks. On the other hand, simpler protocols that

Table 1: Applications And Their Parameters

Application	Sensors	Energy source	Project period	Threshold wind speed (m/s)	Required power (mW)	Scale	Node type	Transmission range (m)	Sample rate	Sample size	Buffer size
CORIE	conductivity, temperature, depth	solar	1–3 years	9.0	N/A	large	N/A	N/A	several / day	N/A	N/A
Pipe network	MEMS based accelerometer	water flow	1 year	N/A	N/A	large	Gopher	300–32000	150/sec	N/A	256 kB or 16 GB
habitat	temperature, humidity, thermopile	solar or wind	9 months	3.0	68	small	mica2 mote	100	6/day	2 B	512 kB
Volcano	microphone, seismometer	solar or wind	5 years	4.0	100	small	TMote sky	200–400	1–2/day	128 B	4 MB
Glacsweb	pressure, temperature, orientation, external conductivity, strain	solar	1 year	9.0	500	medium	mica2 mote	100	6/day	16 B	512 kB
meteorology and hydrology	water level and temperature	water flow	3 years	4.0	100	large	mica2 mote	100–200	6/day	16 B	512 kB

simply send many redundant messages may be sufficient in small-scale networks.

Figure 3 shows transmission latencies for different network scales. When the packet delivery rate requirement is loose, all protocols have similar performance. However, performance differs under stricter packet delivery rate requirements. Geometric routing works well for small- and medium-scale networks, while buffer size dependent routing works well only for large networks. In contrast, Ambient Energy Aware routing works well for small, medium, and large networks.




Figure 3: Transmission latency comparison of four protocols when network scale changes. Arrows indicate that latency exceeded application latency constraint.

5.3.1 Sample size

Sample size is the amount of data gathered by a node per sample. The sample size affects how long a sensor node spends sensing and transmitting data. If the data gathering and transmission time is too long, the sensor node will not have enough time for data transmission, reducing the packet transmission rate. We now discuss the effect of sample size on the network latency.

Figure 4 shows packet finish time as a function of sample size for the four protocols. The packet delivery rates of all protocols stays almost the same when sample size changes. This is because the sensing time takes up only a small portion of the sensor node active time even for the largest sample size that we consider. We conclude that the sample size for similar applications will not significantly affect the packet transmission latency.




Figure 4: Transmission latency comparison of four protocols when sample size changes. Arrows indicate that latency exceeded application latency constraint.

5.3.2 Transmission range

The transmission range of sensor nodes affects routing protocol selection. The larger the transmission range, the more immediate neighbors per sensor node. This is especially important to energy scavenging sensor networks, in which only a subset of neighbors are active at any time.

Figure 5 shows the packet delivery rate as a function of node transmission range for the four routing protocols. The transmission range of a sensor node depends on the radio device used in the sensor node, normally ranging from 100 m to 1 km. When the transmission range is large, simple protocols work as well as Ambient Energy Aware routing. When the transmission range is small, there are so few immediate neighbors per sensor node that selection of neighbors during routing is critical. In this case, Ambient Energy Aware routing protocol outperforms existing alternatives. This result holds for all packet delivery rate requirements.

5.3.3 Power requirement

The power required by a sensor node and an ambient power trace determine when the node will be active. Required power depends on the hardware platform and the software workload. In our example application, ambient power is determined by wind speed. We define the minimum wind speed required by an application as *threshold wind speed*.

The impact of threshold wind speed on packet transmission latency in the sensor network is plotted in Figure 6. When the threshold wind speed is low, all protocols have short packet transmission




Figure 5: Transmission latency comparison of four protocols when transmission range changes. Arrows indicate that latency exceeded application latency constraint.

latencies. For high threshold wind speeds, the Ambient Energy Aware protocol is superior. However, when threshold wind speed is very high, no protocol can finish transmitting enough packets during the active time periods to meet application requirements.




Figure 6: Transmission latency comparison of four protocols when threshold wind speed changes. Arrows indicate that latency exceeded application latency constraint.

6. PROTOCOL SELECTION

This section describes the relationship between protocol performance and both channel utilization and per-node packet delivery rate. It then explains a strategy for selecting an appropriate communication protocol.

6.1 Channel Utilization

Channel utilization is an important additional metric for evaluating our proposed protocols. The channel capacity is limited by the activity rates of nodes in an energy scavenging sensor network. We show the channel utilization for every node in a 500-node sensor network using four protocols. Geographic routing has high channel utilization around the base station, and has very low channel utilization at distant nodes. This increases collision rate near the base station, causing packets to be dropped. Ambient Energy Aware routing has less traffic concentrated at the base station (Figure 8), reducing the packet drop rate. Buffer size dependent routing and undirected routing have higher average channel utilizations compared to the previous two protocols, and their node-level channel utilization distributions follow the node activity rate (Figure 7). This is reasonable because nodes with higher activity rates are available more frequently and therefore may receive more packets. Whether or not this natural bias is helpful depends on the distribution of wind speed. If the wind speed is similar in most sensor locations, the natural bias will lead to spatially balanced channel use. Otherwise, channel traffic will concentrate on active nodes, and is likely to increase packet drop rates around those nodes. In this case, Ambient Energy Aware routing helps by distributing the




Figure 7: The node-level activity rate given by the statistical data.




Figure 8: The node channel utilization for Ambient Energy Aware routing.

channel capacity among nodes with higher probability in delivering packets to the base station.

6.2 Per-Node Packet Delivery Rate

We now consider the *fairness* of the protocols under evaluation, i.e., the variation of packet delivery rates of all nodes in the network. We evaluate the fairness by plotting the location-dependent per-node packet delivery rate across the whole network.

Figure 9 shows the per-node packet delivery rate for a medium-scale sensor network. Ambient Energy Aware routing achieves good fairness among nodes, and results in higher average per-node packet delivery rate, while geographic, buffer size dependent, and undirected routing favor nodes that are closer to the base station. This result is consistent with the ranking function used in each protocol. Geographic routing relies on a ranking function that gives higher priority to nodes closer to the base station. Therefore, it is biased toward causing heavy communication on the nodes around the base station. Ambient Energy Aware routing considers the node distance to the base station as well as its activity rate, and therefore has less severe bias on nodes with different distances to the base station, resulting in better fairness.

6.3 Protocol Selection Strategy

Using the results from this section and Section 5, we are able to provide protocol selection strategies for energy scavenging sensor network designers. We first give the conditions under which the four candidate protocols are appropriate. Undirected routing can only work efficiently in small-scale networks, and when ambient power is sufficient to power sensor nodes for most of the time. Buffer size dependent routing works best for small-to-medium size networks and medium transmission ranges. Geographic routing works best for small scale network and large transmission ranges. Ambient Energy Aware routing works well under most conditions,

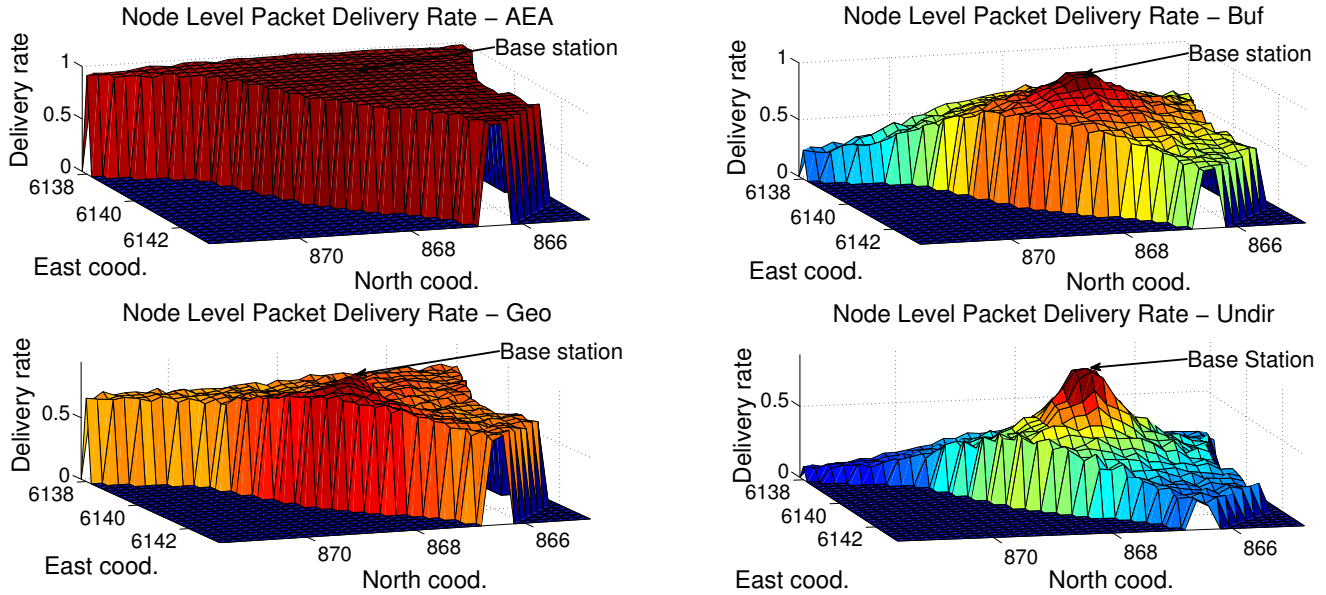


Figure 9: The node level packet delivery rates using four different routing protocols.

and can out-perform others in large-scale networks, even when ambient power is not sufficient to frequently wake up sensor nodes. It adapts better to extreme conditions than other protocols. Based on these working conditions, we provide the designer with several guidelines.

1. Determine required packet delivery rate and latency. The most appropriate protocol depends on these requirements.
2. Select values for sensitive parameters. The values for these parameters should be determined first, since they are very likely to affect the result. These parameters include network scale and sensor node transmission range. The optimal protocol depends strongly on these parameters. On the one hand, Ambient Energy Aware routing outperforms others under “harsh” conditions: when network scale is large or sensor nodes have small transmission range. In addition, it performs well under looser requirements. Designers should choose this protocol when applications require a large number of sensor nodes, or when they are limited by sensor node communication hardware. On the other hand, in a small-scale network using nodes with long transmission ranges, users can instead use simpler protocols.
3. Select values for less sensitive parameters. These parameters either do not greatly affect packet transmission latency, or always result in the same optimal protocol when their values vary. These parameters include sample size and the maximum power consumption of the sensor node. Users are relatively free to select these parameter values; indeed, to improve the overall performance, we recommend selecting the most favorable parameter values for a given budget. For instance, using a larger energy scavenging device can provide more power to the sensor nodes under the same ambient energy conditions; this then allows users to use more powerful sensor nodes that boost the performance.

7. DISCUSSION AND CAVEATS

In this section, we address some of the simplifying assumptions made during our evaluation and discuss the impact that they might have on our reported results. We believe that our experiments capture the most important features of the environment and sensor network well enough to provide a reasonable evaluation of the routing

protocols. However, our evaluation framework does not consider some secondary effects including long-term wind speed variation and adjusting node activity rates online. We now address these issues and consider combining Ambient Energy Aware Design with more sophisticated protocols.

7.1 Long-term Wind Speed Variation

There can be long-term variations in wind speed distribution, on time scales ranging from three months to half a year [21] due to seasonal changes and long-term weather patterns. This means that the wind speed distribution will change multiple times during the lifetime of a long-term deployment. In our previous evaluation, we only use one fixed wind speed distribution at one location. This distribution is used in the ranking function. For a real system deployment, long-term wind speed variation should be considered.

The sensor network design can adapt to this variation by pre-storing multiple wind distributions. The variation of wind speed distribution in one location is periodic, repeating yearly. This periodic distribution is usually available from the wind atlas of local government websites [20, 31]. Therefore, even for long-term operations, only a limited number of wind speed distributions need to be stored in the nodes. We can then program nodes with multiple wind speed distributions and corresponding times at which the wind speed distribution changes. Sensor nodes will know to switch to a new activity rate by monitoring time stamps.

7.2 Online Adjustment Of Node Activity Rate

Sensor nodes can gather information on their activity rates after deployment. These activity rates are representative of the actual power source condition at the node’s location. Therefore, using this value to adjust the pre-stored node activity rates makes them more accurate.

Sensor nodes use timers to gather activity rates. Due to power losses, a node’s timer may not record the correct time. It can nonetheless record how long a sensor node has been active. Every time the node has an opportunity to synchronize its timer, it computes its activity rate by dividing the measured node active time by

the total elapsed time since the last synchronization event and this value is used to update the pre-stored node activity rate.

7.3 Protocol Extensions

Although we mainly compare Ambient Energy Aware routing to geographic routing, other sophisticated geographic routing protocols, such as Greedy Perimeter Stateless Routing (GPSR) [19], also face similar problems when applied to energy scavenging sensor network. GPSR can prevent data from being stuck at the edges of holes in a sensor network. However, it still cannot avoid sending packets to an infrequently active node in an energy scavenging sensor network.

It would be possible to use the concepts described in Section 4 to design an ambient energy aware variation of GPSR by using a similar ranking function for each node. Such a protocol would avoid low activity rate nodes when searching for paths around holes.

8. CONCLUSION

We have described and evaluated a novel class of design techniques for indefinitely deployed sensor networks. To enable this increased lifespan, we proposed eliminating batteries from sensor nodes and introducing a new routing protocol that account for the resulting restrictions in node activity. This protocol can achieve high delivery rate, even when sensor nodes randomly lose their power sources. It uses stochastic models for ambient power sources and takes advantage of spatial and temporal correlation to make routing decisions. We evaluated the newly proposed protocol together with three existing approaches for four commonly used applications. Finally, we provided guidance on selection of routing protocols for specific applications.

9. REFERENCES

- [1] D. Linden, *Handbook of Batteries*, 3rd ed. New York: McGraw-Hill, 2002.
- [2] D. Berndt., *Maintenance-free batteries : lead-acid, nickel/cadmium, nickel/metal hydride : a handbook of battery technology*, 2nd ed. Somerset, England : Research Studies Press, 1997.
- [3] "Health concerns with batteries," http://batteryuniversity.com/learn/article/health_concerns.
- [4] M. Minami, T. Morito, and H. Morikawa, "Solar biscuit: a battery-less wireless sensor network system for environmental monitoring applications," in *Proc. 9th Int. Conf. on Networked Sensing Systems*, 2005.
- [5] A. Mainwaring, et al., "Wireless sensor networks for habitat monitoring," in *Proc. Int. Wkshp. Wireless Sensor Networks and Applications*, Sept. 2002, pp. 88–97.
- [6] E.-H. El Brouji, et al., "Impact of calendar life and cycling ageing on supercapacitor performance," *IEEE Trans. Vehicular Technology*, vol. 58, no. 8, pp. 3917–3929, Oct. 2009.
- [7] "Using a small solar cell and a supercapacitor in a wireless sensor," <http://www.sensorsmag.com>.
- [8] V. Raghunathan, et al., "Design considerations for solar energy harvesting wireless embedded systems," in *Proc. Int. Conf. Information Processing in Sensor Networks*, Apr. 2005, pp. 457–462.
- [9] J. Taneja, J. Jeong, and D. Culler, "Design, modeling, and capacity planning for micro-solar power sensor networks," in *Proc. Int. Conf. Information Processing in Sensor Networks*, Apr. 2008, pp. 407–418.
- [10] T. Voigt, H. Ritter, and J. Schiller, "Utilizing solar power in wireless sensor networks," in *Proc. Int. Conf. Local Computer Networks*, Oct. 2003, pp. 416–422.
- [11] "SMART modular technologies," <http://www.smartm.com>.
- [12] J. Schiffer, D. Linzen, and D. U. Sauer, "Heat generation in double layer capacitors," *J. Power Sources*, vol. 160, no. 1, pp. 765–772, 2006.
- [13] X. Jiang, J. Polastre, and D. Culler, "Perpetual environmentally powered sensor networks," in *Proc. Int. Conf. Information Processing in Sensor Networks*, Apr. 2005, pp. 463–468.
- [14] M. Philipose, et al., "Battery-free wireless identification and sensing," *IEEE Pervasive Computing*, vol. 4, no. 1, pp. 37–45, Jan. 2005.
- [15] R. Vyas, et al., "Design and characterization of a novel battery-less, solar powered wireless tag for enhanced-range remote tracking applications," in *Microwave Conference, European*, Oct. 2009, pp. 169–172.
- [16] A. Patel, et al., "Power harvesting for low power wireless sensor network," in *Antennas Propagation Conference, 2009. LAPC 2009. Loughborough*, Nov. 2009, pp. 633–636.
- [17] J. W. Ng, et al., "Ubiquitous monitoring environment for wearable and implantable sensors (ubimon)," in *Proc. 6th Int. Conf. on Ubiquitous Computing*, Sept. 2004.
- [18] R. Morais, et al., "Sun, wind and water flow as energy supply for small stationary data acquisition platforms," *Computers and Electronics in Agriculture*, vol. 64, pp. 120–132, Dec. 2008.
- [19] F. Zhao and L. J. Guibas, *Wireless Sensor Networks: An Information Processing Approach*. Morgan Kaufmann, 2004.
- [20] "World wind maps – wind atlases of the world," <http://www.mywindpowersystem.com>.
- [21] C. L. Archer and M. Z. Jacobson, "The spatial and temporal distribution of U.S. winds and windpower at 80 m derived from measurements," *American Geophysical Union*, p. C82, Dec. 2002.
- [22] S. Chan, et al., "Operations requirements of utilities with wind power generation," *IEEE Trans. Power Apparatus and Systems*, vol. PAS-102, no. 9, pp. 2850–2860, Sept. 1983.
- [23] I. Damousis, et al., "A fuzzy model for wind speed prediction and power generation in wind parks using spatial correlation," *IEEE Trans. Energy Conversion*, vol. 19, no. 2, pp. 352–361, June 2004.
- [24] G. Werner-Allen, et al., "Deploying a wireless sensor network on an active volcano," *IEEE Internet Computing*, vol. 10, no. 2, pp. 18–25, Mar. 2006.
- [25] G. Werner-Allen, et al., "Fidelity and yield in a volcano monitoring sensor network," in *Proc. Int. Symp. Operating Systems Design and Implementation*, Nov. 2006, pp. 381–396.
- [26] K. Martinez, R. Ong, and J. Hart, "Glacsweb: a sensor network for hostile environments," in *Proc. Conf. Sensor and Ad Hoc Communications and Networks*, Oct. 2004, pp. 81–87.
- [27] J. Lundquist, D. Cayan, and M. Dettinger, "Meteorology and hydrology in Yosemite national park: A sensor network application," in *Proc. Int. Conf. Information Processing in Sensor Networks*, Apr. 2003, pp. 553–553.
- [28] S. N. Simic and S. Sastry, "Distributed environmental monitoring using random sensor networks," in *Proc. Int. Conf. Information Processing in Sensor Networks*, Apr. 2003, pp. 582–592.
- [29] "Center of costal margin observation and prediction," <http://www.ccalmr.ogi.edu/CORIE/>.
- [30] M. Shinozuka, et al., "Nondestructive monitoring of a pipe network using a MEMS-based wireless network," in *Proc. SPIE*, Mar. 2010, p. 7649.
- [31] "Wind energy resource atlas of the United States," http://rredc.nrel.gov/wind/pubs/atlas/atlas_index.html.