

Designing Wireless Sensor Networks as a Shared Resource for Sustainable Development

Nithya Ramanathan[†] Laura Balzano[‡] Deborah Estrin[†] Mark Hansen[⊕]
Thomas Harmon[∪] Jenny Jay[∏] William Kaiser[‡] Gaurav Sukhatme[∩]

UCLA, Center for Embedded Network Sensing

[†]{nithya,destrin}@cs.ucla.edu [‡]{sunbeam,kaiser}@ee.ucla.edu [⊕]{cocteau}@stat.ucla.edu
[∪]{tharmon}@ucmerced.edu [∏]{jennyayla}@gmail.com [∩]{gaurav}@sagan.usc.edu

Abstract

Wireless sensor networks (WSNs) are a relatively new and rapidly developing technology; they have a wide range of applications including environmental monitoring, agriculture, and public health. Shared technology is a common usage model for technology adoption in developing countries. WSNs have great potential to be utilized as a shared resource due to their on-board processing and ad-hoc networking capabilities, however their deployment as a shared resource requires that the technical community first address several challenges. The main challenges include enabling *sensor portability* – the frequent movement of sensors within and between deployments, and *rapidly deployable* systems – systems that are quick and simple to deploy. We first discuss the feasibility of using sensor networks as a shared resource, and then describe our research in addressing the various technical challenges that arise in enabling such sensor portability and rapid deployment. We also outline our experiences in developing and deploying water quality monitoring wireless sensor networks in Bangladesh and California.

1 Introduction

Wireless Sensor Networks (WSNs), networks of wirelessly connected sensing and computational devices, hold tremendous promise for many areas of development including public health, the environment, and agriculture. A single device has a processor, a radio, and several sensors. When a network of these devices is deployed in a field, the sensing devices measure particular aspects of the environment. The devices then communicate those measurements by radio to one another and to more powerful computers for data analysis. In this way, WSNs can provide detailed observations of various phenomena that occur in the environment.

WSNs are capable of measuring diverse phenomena such as contaminant levels in water, pollutants in the air, and the flow of water for irrigation. As an example of a potential application, consider the recent incident of contamination spilling into the Songhua river in China, the main source of drinking water for many people¹. Determining rate of flow and sometimes direction of the river requires coordination of multiple sampling points. Sensors periodically taking samples at multiple locations along the river could determine the rate, quantity, and direction of contaminant flow using the distributed sensing and

¹<http://www.china.org.cn/english/2005/Dec/150566.htm>

processing of a wireless sensor network.

Unfortunately, the potential of wireless sensor networks for sustainable development² remains largely untapped while they are designed primarily for relatively resource-rich application contexts. The cost of WSNs is one of several major barriers that prevents them from being leveraged for sustainable development applications. Many components of WSNs are becoming cheaper (e.g. computing power), but the sensors themselves remain the most expensive component³. As stated in [5], successful technology-based international development projects rely on shared technology due to excessive cost of personal devices. However, most research on sensor networks is based on long-term deployments owned by a single user, a paradigm not conducive for sharing. The complexity of technology management is another barrier. We use Grameen telecom as a successful model⁴ in which the management and maintenance of shared hardware is centralized. We envision a sensor network much in the same light.

Many sensor network applications are conducive to such a shared model. We base this statement on the observation that sensors may not be required in a single location for extended periods of time for reasons including: (1) a phenomenon of interest may have a slow rate of change, thus a small number of sensors can be moved within a deployment, emulating the density required to sufficiently capture the physical phenomena, (2) the initial deployment

²Sustainable development is defined as a process of developing that “meets the needs of the present without compromising the ability of future generations to meet their own needs”, and whose “interdependent and mutually reinforcing pillars are economic development, social development, and environmental protection.” [12, 1]

³While small temperature sensors are available for less than a dollar, many sensors purchased off the shelf are less common and are significantly more expensive. For example, we needed an ammonium sensor that could be left in the environment for our water quality WSN in Bangladesh. The cheapest acceptable sensor we found cost around \$400 from Sentek (<http://www.sentek.co.uk>)

⁴The model here is one woman who owns a cell phone and re-sells minutes to those who only need the phone for a short period of time. This woman is in charge of the upkeep and management of this piece of technology and has a vested interest in ensuring that the phone continues to work. <http://www.grameen-info.org/grameen/gtelecom> [4]

may have been too dense, thus redundant sensors can be removed, and (3) the duration of the deployment may be short. We discuss these scenarios in more detail in Section 3.

All of the deployment scenarios mentioned above rest on the assumption that sensors can be easily deployed and re-deployed. While WSNs have great potential to be utilized as a shared resource due to their on-board processing and ad-hoc networking capabilities, their deployment as a shared resource requires that the technical community first address several challenges, including enabling *sensor portability* – the frequent movement of sensors within and between deployments, and *rapidly deployable* systems – systems that are quick and simple to deploy. This leads us to our major challenges in Section 4.

Clearly, the primary issues related to successful technology adoption are the social, policy, and logistical questions to be answered in order to enable equitable access and the design of culturally appropriate technology. Our experience, though relevant, is limited to our technical expertise. These challenges and others should be formulated more explicitly with the necessary diverse input from communities, activists, governments and NGOs.

In this paper we focus on justifying the technical feasibility of designing sensor networks as a shared technology (Section 3) and describing the technical challenges that must be addressed to enable WSNs as a shared technology (Section 4). We begin by describing our applications in water quality monitoring in Bangladesh and California (Section 2).

2 WSNs For Water Quality

Wireless sensor networks are made up of small computational devices connected to various sensors and wireless radios. The devices automatically and adaptively form ad-hoc networks (temporary point-to-point networks) over wireless radios to make decisions based on measurements of their environment. The hardware and software are designed to be extremely low power in order to enable long-term in-situ deployments, i.e. undisturbed




Figure 1: Illustration of our deployment in a rice field in Bangladesh (drawing by J.Fisher, UC Merced).

deployments that are left in the environment with minimal human intervention. Device sizes commonly range from that of a quarter to a PDA-like device. In general, resource availability and power consumption are commensurate with size. For example, while it largely depends on the power consumption of the sensors, the lower-power nodes (often called motes) can run for approximately one month on 2 AA batteries.

Sensor networks provide dense spatial and temporal sampling even in remote and hard to reach locations. Thus, they are best applied to applications that need dense sampling in space and/or time. Soil applications are a good example, because the soil environment is heterogeneous across space, requiring dense spatial sampling. Abrupt changes can then be captured with a high temporal sampling rate.

The fact that WSNs are low power and wireless makes them appealing as a technology for developing regions, but in addition the dense sampling is crucial for public health applications. For example, [17] states that while water quality concerns can be extremely critical, “analysis is still primarily conducted in a laborious manner by physical collection of a sample that is analyzed back in a laboratory.” This kind of data collection and analysis is time consuming and mostly undirected, and in many instances misses the toxin events of interest.

We are involved with two ongoing WSN deployments related to groundwater quality: a system to understand the prevalence of arsenic in Bangladesh groundwater, and a system to monitor




Figure 2: Left panel is a photograph of the pylon without the enclosure lid. Right panel is a depiction of the pylon with sensors as they would be embedded in the environment.

nitrate propagation through soils and ground water in California.

Both of our deployments have a similar setup. A *pylon* [10] (Figure 2) consists of an enclosure housing the small wireless devices which connect to groups of sensors embedded at multiple depths in the soil through long wires. Each device can support 7 sensors and there are multiple devices per pylon. Multiple pylons are deployed around the field to attain vertical and horizontal spatial density. The devices wirelessly transmit samples back to a base-station for analysis (Figure 1). The base-station in these deployments was a PDA-class device. It could also be a laptop. It is powered by a car-battery recharged using solar panels. To make data externally accessible, our base-station is connected using WiFi (i.e. 802.11b), or where WiFi is unavailable, using a GPRS (i.e. cellular) network.

Arsenic in Bangladesh In Bangladesh, tens of millions of people in the Ganges Delta drink ground water that is dangerously contaminated with arsenic. If consumption of contaminated water continues, the prevalence of arsenicosis and skin cancer will be approximately 2,000,000 and 100,000 cases per year,




Figure 3: Diurnal variations in Ammonium concentration over a period of days until irrigation occurs at day 8.

respectively, and the incidence of death from cancer induced by arsenic will be approximately 3,000 cases per year [18].

A full understanding of the factors controlling arsenic mobilization to ground water is lacking. A current working hypothesis in some regions is that the influx of dissolved arsenic to ground water is greatly enhanced where irrigation for rice cultivation provides the primary source of aquifer recharge⁵.

In a joint collaboration with scientists at the Bangladesh University of Engineering and Technology and MIT, we deployed a sensor network in January of 2006 in a rice field near Dhaka, Bangladesh in order to aid in validating this hypothesis. A full pylon contains 3 complete suites of sensors (soil moisture, temperature, carbonate, calcium, nitrate, chloride, oxidation-reduction potential, ammonium, and pH), each deployed at a different depth (1, 1.5, and 2 meters below ground), and a pressure transducer at the base to monitor water depth. We could not find an off-the-shelf, in-situ arsenic sensor to include in this deployment. Instead, output from a manual arsenic

⁵An aquifer is a body of geologic material that can supply useful quantities of ground water to natural springs and water wells. Aquifer recharge is the process by which water seeps down through the soil into an underlying aquifer (<http://www.nj.gov/dep/njgs/enviroed/aqfchrg.htm>)

sensor will be combined with the data collected from the sensor network, which is primarily used to get a better understanding of the groundwater chemistry at shallow depths. We deployed one fully-equipped pylon, and two partially equipped pylons (with one and two depths of sensors) for a total of 48 sensors in the field for a period of 10 days. Even with such a short deployment, the sensor network captured some interesting phenomena, as seen in Figure 3.

Ground water Contamination in Palmdale

Water scarcity in arid and semi-arid regions and increasing demand on water supplies has stimulated interest in the reuse of treated wastewater. Despite the many benefits to irrigating with reclaimed water, there remain both real and perceived risks to human health and environmental quality stemming from residuals in the treated wastewater. Proactively addressing these risks requires automating the distributed observation and control of the irrigation water and the trace pollutants that it conveys, including suspended or dissolved solids (TDS), colloidal solids, pharmaceuticals, organic carbon, volatile organic compounds, pathogenic microorganisms, and nutrients such as nitrogen or phosphorus. A water reuse site in Palmdale, California is being used as a testbed for a sensor network with soil moisture, temperature, and nitrate sensors. The network focuses on two things: first, ensuring that environmental regulations are being met, and second, providing feedback to a water control system in order to optimize water flow and minimize chemical penetration into the subsurface. This site is also used to test the software, sensors, and hardware before deploying in Bangladesh.

3 Sensor Sharing Techniques

Sensor sharing will allow many people to benefit from sensor network data collection, even with minimal sensor resources. We believe the following three technical approaches are particularly suited for enabling sensor sharing for sustainable development: (1) moving a smaller number of sensors around in a deployment to *emulate density*, (2) gradually

removing redundant sensors from a deployment to go from *dense to sparse* deployments, and (3) leveraging *shorter deployment cycles* where possible. Here we describe each of these scenarios in greater detail, including a survey of our own and others' work in implementing related or supporting algorithms.

3.1 Emulating Density

Human-enabled mobility can be used to manually emulate the effect of a dense deployment using fewer sensors. People can move a small set of sensors around in a field in order to collect data for a dense spatial map of the field. This technique will be appropriate only for sustainable development applications in which the phenomenon of interest changes very slowly, on the order of days or longer.

Here we describe two existing systems that emulate density. In each of these systems, robotic mobility of one node enables dense mapping of a large space that would be extremely expensive with a static deployment of sensors. Decisions in these systems of when and where to move the mobile WSN devices are derived based on ongoing sensor data collection. These algorithms can be altered to direct a human placing a sensor instead of directing a robot's movement. This framework must be flexible to account for the error of human placement, but will benefit from an interactive human audit.

Infrastructure-based Robotics The Networked Infomechanical Systems (NIMS) [2, 13, 14] project is an infrastructure-based robotics system which seeks to allow autonomous high-precision control of sensor position to attain a dense map of large three-dimensional spaces. The setup of NIMS involves a suspension cable system that allows coordination of movement of attached sensors. This suspension system includes lightweight cables that provide both horizontal traction force and vertical elevation tension for translation of the mobile device. NIMS is an example of a system which emulates density. Instead of deploying multiple static sensors in an area, we deploy a single NIMS node with a set of sensors and move this node to obtain a map of the area. NIMS is also rapidly deployable (Section 4).




Figure 4: The upper panel shows a deployment of the NIMS node (pointed to by the arrow), which moves along a cable run across the San Joaquin River downstream of a confluence. The lower panel shows contributions of contamination from the two merging rivers.

One NIMS system is currently being used to resolve the spatial distribution in contaminants that results from the mixing of the San Joaquin and Merced rivers [9]. The NIMS node is deployed on a cable attached to two anchor points on either side of the river. The deployment is downstream from the confluence. We have been able to see distinct gradients in redox active species, such as ammonium, in channel cross-sections within the confluence zone (Figure 4). The lower panel is a spatial map of ammonium concentration obtained using NIMS. The combination of flow velocity and concentration data provided by the NIMS node permits a more accurate estimation of the total mass flow of contaminants than was previously possible.

Adaptive sampling algorithms [3, 14] have been developed for the NIMS node to control the direction and speed of movement of the node. The algorithms assume the phenomenon of interest does not change more rapidly than it takes for the node to cover the transect. As the node moves across the transect and data is collected, the sampling pattern of the node is adapted to sample more in areas where there is high




Figure 5: Reconstructions of chlorophyll concentration at the surface of Lake Fulmor in Southern California, based on measurements made in October 2005, with 5 (top) and 9 (bottom) sensor buoys. The denser sampling reveals a finer structure.

spatial variability.

Acquatic Actuated Sampling The Networked Aquatic Microbial Observation System (NAMOS) project seeks to establish a continuous sensing presence for the analysis of data related to chemical, physical and biologically pertinent phenomena in aquatic environments [7, 19]. The system operates at scales relevant to the study of micro-organism abundance ($10^2\text{m} - 10^4\text{m}$), and at these scales it is infeasible to deploy a dense set of static monitoring stations over a water body for continuous monitoring. NAMOS has deployed a network with both static and mobile components to simultaneously provide flexibility in both spatial and temporal sampling.

The current NAMOS system consists of a single mobile boat and 10 static buoys, each equipped with a fluorometer and thermistors. By controlling the deployment density of the static nodes and ‘filling in’ with measurements using the boat, it is possible to reconstruct the temperature and chlorophyll fields of moderately sized water bodies at multiple scales (Figure 5); the denser sampling reveals a finer structure.

The NIMS and NAMOS systems are not yet designed for sustainable development. If such deployments are to be effective in resource-constrained situations, the long-term sustainable strategy is to remove robotics and put a human in the loop for the control of a hand-cranked NIMS node or NAMOS rowboat. The computer algorithms collecting and analyzing data will direct the human user to navigate the node to the appropriate position. The issues which arise due to latency and accuracy of human mobility are discussed in Section 4.

3.2 Dense to Sparse Deployments

Some sensor network applications require a dense mapping of the environment. Once sensors are densely deployed and details of the phenomenon are revealed, we may see it is possible to capture sufficient information with fewer sensors, freeing sensors for deployment elsewhere. Here we describe applicable work which is ongoing in the sensor network community.

In [16], a technique called Backcasting is described to identify unnecessary sensors. This work assumes that the field is densely deployed, and the algorithm turns off as many sensors as possible to maintain a certain level of fidelity of sensing. Their method uses measurements from the dense deployment to estimate the spatial nyquist frequency throughout the field. Where the frequency is low, sensors can be shut down to conserve energy. This can be adapted to direct a human user to remove these unnecessary sensors from the field.

Another approach views the sensor network as a database query-response system: Model-based Data Acquisition [6]. A user issues a query to a densely deployed network. Using a gaussian process (GP) model built from past data, the system chooses only a select number of sensors that must be queried to get the appropriate response. This can be extended to applications we have described where there is a single query for the duration of the deployment. Thus, the algorithm will find a particular subset of sensors that is useful to answer that query and can direct a human user to remove the other sensors.

A third applicable work on optimizing sensor

placement [11] takes data from an initial dense deployment and uses it to redeploy sensors to near-optimal locations. GP models are built from the initial data. The optimization program then uses these models to find near-optimal redeployment locations for the sensors in order to minimize communication cost and maximize sensing information. We could slightly change the optimization problem to minimize the number of sensors given a required overall sensing fidelity. Then, upon redeployment, some of the sensors in the first deployment would be unnecessary and thus would be removed, and the remaining sensors would be deployed in new, near-optimal locations.

In future deployments in Bangladesh, one of our goals is to densely deploy a field in order to determine indicators for the presence of arsenic and arsenic's temporal variation in a certain region. Once dense sampling has provided a map of the field, unnecessary sensors can be removed from the deployment. Families in this region can maintain a smaller and simpler sensor network in a few crucial locations in order to monitor for the presence of arsenic. If arsenic is increasing, a more sophisticated sensor network could be brought in to again densely sample the area.

3.3 Short Deployment Cycles

Some applications only require short-duration deployments and therefore are ideal for sensor sharing. Our deployment in Bangladesh is an example of an application with a short deployment cycle. We wanted to collect data to validate a hypothesis about diurnal variations, and so we wanted several days of data for analysis.

Another scenario in which short deployment cycles are appropriate is in a trigger-response sensor network usage model. Individuals own simple, inexpensive sensors for a particular contaminant, which communicate their measurements through a cellular network. Upon detection of unusual phenomena in an area, an NGO could bring a more sophisticated sensor network for a short-duration, detailed analysis of contamination transport.

A trigger-based usage model such as this can

build on systems like one which is being developed at Columbia University. A group there plans to distribute needle-sampler arsenic sensors connected to cell phones to communities in Bangladesh⁶. A phone-in data system uses statistics to estimate the probability that arsenic is present at a given depth in the ground.

A trigger system designed in this way would require the application of certain principles from robust statistics. No one sensor measurement should be the trigger for a major movement of resources. We are working on algorithms appropriate for sensor data to both filter out unreliable measurements and identify and validate data of interest. Multiple sensor triggers along with triggers from local health information will help to certify interesting data. The probability of event detection, false positives and event misses will vary depending on the number and type of available triggers in a particular region.

An additional benefit from short deployment cycles is that a user is readily available throughout the deployment to ensure that the data collection is successful. For example, in our Bangladesh deployment, when issues came up in the deployment we could address them immediately, thus maximizing the usable data retrieved from the deployment.

4 Challenges

Numerous technical challenges arise in order to be able to quickly deploy and move sensors, primarily because the work to date has largely focused on static, long-running deployments.

Given that we have the goals to emulate density, reduce dense deployments to sparse ones, and leverage short deployments cycles, we find the following three challenges to be the most pertinent. Algorithms must be interactive and robust to human error. Faults in the system must be quickly identified to maximize the amount of good data received. Finally, systems must be made to be rapidly deployable. In this section we discuss our research in these three areas.

⁶<http://www.earthinstitute.columbia.edu/news/2005/story01-05-05.html>

4.1 Algorithm Issues

Sensor portability introduces challenges and new requirements in algorithm design.

Robust and Interactive Deployment Algorithms Using human-enabled mobility to move sensors in a deployment can be cheaper than robotics, depending on labor costs. Of course, human-enabled mobility is neither as accurate nor as sensitive to latency as robotic mobility, and deployment algorithms must take this fact into account.

Thus, in order to guide a person to move sensors in a deployment, our algorithms must do two things. First, algorithms must employ some very basic audio or visual cues that provide feedback for a user. For instance, a green light might turn on when the node is in the correct location for deployment, or a red light might turn on once a sensor can be removed. Mechanical mechanisms could also be built into the hardware; a retractable measurement device could be attached to a pylon so that no two pylons are deployed within a certain distance of each other.

Second, algorithms must be robust to human error and avoid frustrating the user. Requiring a user to place a sensor within a very small area or within a very small amount of time is unreasonable. In order to appropriately utilize human mobility, the algorithm must be able to tolerate placement errors and latency.

In many cases, the *human-in-the-loop* interactivity actually aids in designing more fault-tolerant algorithms [5]. Instead of treating human interactivity as an after-thought, algorithms that rely on human intervention can be more robust to faults. For example, a user pushing a button as a way of saying “I have now deployed this node” may be more reliable than robotic-enabled mobility. Interaction between people and technology has been intensely studied (e.g. human/computer interaction or HCI), thus we leave a discussion of the ideal sensor interface for future work.

4.2 Detecting Data Disruption and Faults

The primary cause of faults and data disruption in a wireless sensor network is a failure in the communication or a failure in the sensor. Finding and fixing failures in a sensor network is a difficult problem because the devices are low-power and relatively cheap. Failures in the network can occur for a variety of reasons, including bad wiring, faulty hardware, uncalibrated sensors, buggy software, badly placed sensors, or bad communication between nodes due to physical obstructions or distance.

The prevalence of failures increases the importance of detecting and fixing them immediately in order to enable sensors as a shared resource. However, the amount of system and sensor data can often be overwhelming, and manual management and analysis can quickly become intractable. Thus we need tools that can be used in the field for monitoring network health, for validating data, and for knowing when to calibrate, fix, or replace sensors as the data is being collected.

Monitoring Network Health In order to aid users in finding and fixing network failures, we designed a tool called Sympathy [15]. Sympathy highlights anomalous network behavior based on the quantity of data expected at the base-station from each node in a network. For example, if the base-station expects to receive a sensor measurement from every node in the network once every five minutes, Sympathy identifies nodes that are not transmitting these measurements. Sympathy uses information periodically collected from the network and a decision tree derived from the data flow model in order to identify a potential cause for every failure, such as a lossy communication link or dead hardware.

Sympathy also aims to minimize the number of failures the user needs to fix to get the network up and running. Once Sympathy identifies a cause for every observed failure in the network, it then goes back and groups failures based on a common cause in order to reduce the number of failure reports to the user. This helps the user focus their efforts on only fixing critical problems in the network.

Detecting Sensor Faults Data integrity in short-term deployments is a critical issue. In our Bangladesh deployment, we saw faulty measurements for issues including broken sensors and shorted circuits. Often sensor measurements were indecipherable due to excessive faults in the data.

We are working to develop a toolbox for detecting these problems real-time during deployment of our system. Our approach is to identify patterns in the data that indicate sensor failure. These fault patterns will be associated with particular causes for the user to fix. In this way, the user can address issues immediately in order to maximize usable data collected by the network.

4.3 Rapidly Deployable Systems

In order to frequently move sensors, they must be extremely easy to deploy and re-deploy.

NIMS systems may be deployed rapidly in environments by simply attaching the NIMS cable system between two fixed points and attaching a WSN node control device for cable actuation. NIMS rapidly deployable systems have been developed for river and stream monitoring. For example, an investigation of the spatiotemporal distribution of nitrate concentration and other variables in an urban stream of Los Angeles, California is performed monthly. The current time for deployment is only two hours, and the system operates over 24 hours.

In soil applications, a challenge is to minimize the disturbance due to placing sensors in the environment. Depending on soil type and moisture conditions, disturbed soil can require days to months to recover from intrusions made for the sake of sensor placement. In the pylon unit described in Section 2, sensors extend from the conduit to achieve intimate contact with the surrounding soil. Deployment strategies causing less disturbance would be highly desirable. We are developing conduits called javelins [8] for this purpose. For aquatic chemical sensors, the javelin requires water-saturated soil conditions, because the target chemicals must be transported through openings in the conduit to the sensor. Additionally, soil is a harsh environment

for sensors, so such conduits would allow sensor withdrawal for cleaning and maintenance.

5 Conclusion

Wireless sensor networks have the potential to be a useful tool for sustainable development. This can be facilitated by the technical community if we focus on issues with developing wireless sensor networks as a shared technology. In order to implement WSNs as a shared resource, we identified three promising technical approaches: emulating density, moving from dense to sparse deployments, and implementing short deployment cycles. We discussed our work on deployments that have demonstrated these techniques and described our past and ongoing work to address the major challenges which arise.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. CCF-0120778, 0331481, and BES-0605515.

References

- [1] The General Assembly. *2005 World Outcome Document ref: Resolution adopted by the General Assembly 60/1. 2005 World Summit Outcome*. United Nations, October 2005.
- [2] M. Batalin, W. Kaiser, R. Pon, G. S. Sukhatme, G. Pottie, Y. Yu, J. Gordon, M. H. Rahimi, and D. Estrin. Task allocation for event-aware spatiotemporal sampling of environmental variables. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, August 2005.
- [3] M. Batalin, M. H. Rahimi, Y. Yu, D. Liu, A. Kansal, G. S. Sukhatme, W. Kaiser, M. Hansen, G. Pottie, M. Srivastava, and D. Estrin. Call and response: Experiments in sampling the environment. In *Procs. of ACM SenSys*, pp. 25–38, Nov 2004.

- [4] D. Bornstein. *The Price of a Dream*. The University of Chicago Press, 1997.
- [5] E. Brewer, M. Demmer, B. Du, M. Ho, M. Kam, S. Nedeveschi, J. Pal, R. Patra, S. Surana, and K. Fall. The case for technology in developing regions. In *IEEE Computer*, June 2005.
- [6] A. Deshpande, C. Guestrin, S. Madden, and J. Hellerstein. Model-driven data acquisition in sensor networks. In *Proceedings of VLDB*, 2004.
- [7] A. Dhariwal, B. Zhang, B. Stauffer, C. Oberg, G. S. Sukhatme, D. A. Caron, and A. A. Requicha. Networked aquatic microbial observing system. In *IEEE Intl. Conf. on Robotics and Automation*, May 2006.
- [8] Thomas C. Harmon. Microsensors to the model forecasts: Multiscale embedded networked sensing of nutrients in the watershed. *Eos Trans. American Geophysical Union, Fall Meet. Suppl.*, 86(52):H23G–08 (Abstract No), 2005.
- [9] Thomas C. Harmon, Richard F. Ambrose, Robert M. Gilbert, Jason C. Fisher, Michael Stealey, and William J. Kaiser. High resolution river hydraulic and water quality characterization using rapidly deployable networked infomechanical systems (nims rd). *CENS Technical Report No. 60*, 2006.
- [10] E. Haux, T.C. Harmon, N.D. Busek, T. Schoellhammer, J. Saez, J. Kim, Y. Park, and D. Estrin. Embedded network sensing of moisture and nitrate propagation during irrigation with reclaimed wastewater. In *Proceedings of the Annual International Meeting of the American Society of Agricultural Engineers*, August 2004.
- [11] A. Krause, C. Guestrin, A. Gupta, and J. Kleinberg. Near-optimal sensor placements: Maximizing information while minimizing communication cost. In *Proceedings of Information Processing in Sensor Networks*, 2006.
- [12] United Nations. *WCED, Our Common Future*. Oxford University Press, 1987.
- [13] M. H. Rahimi, M. Hansen, W. Kaiser, G. S. Sukhatme, and D. Estrin. Adaptive sampling for environmental field estimation using robotic sensors. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pp. 747–753, August 2005.
- [14] M. H. Rahimi, R. Pon, W. Kaiser, G. S. Sukhatme, D. Estrin, and M. Srivastava. Adaptive sampling for environmental robotics. In *IEEE Intl. Conf. on Robotics and Automation*, pp. 3537–3544, 2004.
- [15] N. Ramanathan, K. Chang, R. Kapur, L. Girod, E. Kohler, and D. Estrin. Sympathy for the sensor network debugger. In *Proc. of Sensys*, 2005.
- [16] R. Willett, A. Martin, and R. Nowak. Backcasting: adaptive sampling for sensor networks. In *Proceedings of Information Processing in Sensor Networks*, 2004.
- [17] X. Yang, K.G. Ong, W.R. Dreschel, K. Zeng, C.S. Mungle, and C.A. Grimes. Design of a wireless sensor network for long-term, in-situ monitoring of an aqueous environment. *Sensors*, 2:455–472, 2002.
- [18] W. Yu, C. M. Harvey, and C. F. Harvey. Arsenic in the groundwater in bangladesh: A geostatistical and epidemiological framework for estimating health effects and evaluating remedies. *Water Resources Research*, 39(6), 2003.
- [19] B. Zhang, G. S. Sukhatme, and A. A. Requicha. Adaptive sampling for marine microorganism monitoring. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pp. 1115 – 1122, 2004.