Geometric Approaches to Ad Hoc and Sensor Networks: Full Report of the 2006 NSF Workshop

Subhash Suri
Department of Computer Science
University of California
Santa Barbara, CA 93106

Leonidas J. Guibas
Computer Science Department
Stanford University
Stanford, CA 94305

Alon Efrat
Department of Computer Science
University of Arizona
Tuscon, AZ 85721

December 3, 2006

Anish Arora (Ohio State U.), Sandor P. Fekete (U. Braunschweig, Germany), Jie Gao (SUNY Stonybrook), Robert Ghrist (U. Illinois), Carlos Guestrin (CMU), Thomas C. Henderson (U. Utah), Brad Karp (U. College London), Thomas J. Peters (U. Connecticut), Sylvia Ratnasamy (Intel Research, Berkeley), Andrea Richa (Arizona State U.), Gaurav Sukhatme (USC), Andreas Terzis (Johns Hopkins U.), RogerWattenhofer (ETH Zurich), Feng Zhao (Microsoft)

Abstract

Embedded networked sensing devices are becoming ubiquitous across many activities that are important to our economy and life, from manufacturing and industrial sensing, to agriculture and environmental monitoring, to hospital operations and patient observation, to battlefield awareness and other military applications. In each such deployment modest to large numbers of simple devices that possess sensing and processing capabilities are networked together to form a sensor network. The fact that nodes in these networks are embedded in the physical world and their sensed data is highly correlated with their physical locations imparts a uniquely geometric character to these systems. The geometry and topology of both the sensor field layout, as well as that of the signal landscapes studied, greatly affects issues such as routing, data aggregation and information brokerage, outlier detection and other statistical processing, and so on. This report brings the key findings about the opportunities presented by the exploitation of geometric methods in ad hoc and sensor networks, based on a two-day NSF-sponsored workshop held at the University of California at Santa Barbara during June 12-13, 2006.

Contents

1 Introduction **3**

2 Sensor Networks and the Role of Geometry **4**

3 Information Dissemination, Aggregation, Brokerage, and Summarization **4**
3.1 Dissemination . . . 4
3.2 Aggregation . . . 5
3.3 Brokerage . . . 5
3.4 Summarization . . . 6

4 Structure Discovery and Self-Organization **7**
4.1 Node localization . . . 8
4.2 Virtual coordinates . . . 8
4.3 Sensor field morphology . . . 8
4.4 Clustering and other hierarchical structures . . . 9

5 Sensor placement and coverage **9**
5.1 Motivations and Impact . . . 9
5.1.1 Algae Biomass Monitoring . . . 9
5.1.2 Sensing Light Intensity: understanding photosynthesis . . . 10
5.1.3 Water Distribution Network: detecting dangerous pathogens . . . 10
5.1.4 Surveillance and Search . . . 10
5.1.5 Sensor placements and Autonomic Computing . . . 10
5.2 Fundamental Problems and Challenges . . . 11
5.2.1 Stationary Coverage . . . 11
5.2.2 Noisy Sensors and Uncertainty . . . 11
5.2.3 Active learning . . . 11
5.2.4 Robustness to failures and adversaries . . . 11
5.2.5 Communication Constraints . . . 11
5.2.6 Mobile Coverage . . . 12
5.2.7 Power Constraints and Adaptive Allocation . . . 12
5.3 Mathematical Tools . . . 12
5.3.1 Computational Geometry . . . 12
5.3.2 Algebraic Topology . . . 12
5.3.3 Statistical Models . . . 13
5.3.4 Optimization Theory . . . 13

6 Naming and Routing **14**
6.1 Location-based naming and routing . . . 15
6.2 Virtual coordinates and routing . . . 15
6.3 Challenges . . . 16

7 Sensor Networks and Mobility **16**

8 Wireless network models **18**
8.1 Introduction . . . 18
8.2 Models . . . 18
8.2.1 Connectivity . . . 18
8.2.2 Interference . . . 21
8.2.3 Power Consumption . . . 22

9 Opportunities and Impact **22**

1 Introduction

Enabled by recent advances in micro-electronics and fabrication, a new generation of integrated embedded devices, called *smart sensors*, has emerged that seems capable of realizing the long-cherished vision of *sensory omnipresence* or *ubiquitous awareness*. Through collaboration and *ad hoc* wireless networking, a collection of such devices can provide real-time, fine-grained sensing, monitoring, and actuation across large geographical areas. Because of their small form factors and ability to operate in an untethered mode, these *sensor networks* can achieve an unprecedented level of universality: they can be deployed almost anywhere (even air dropped), are able to organize themselves into a network through self-localization and ad hoc wireless communication, and function unattended for long durations. Building on these ideas, a number of exciting research prototypes have already been proposed and implemented as proofs of concept during the last few years, with varying goals of habitat monitoring, wildlife dynamics, aquatic observations, surveillance, structural monitoring, as well as global research initiatives like the International Polar Year (IPY). Given the scientific and engineering boldness of this vision and the enormous potential benefit to the society at large, it is not surprising at all that sensor networks have attracted strong interest from both academia and the industry.

Within academia, sensor networks have elicited interest from an unusually broad spectrum: from device designers to domain scientists, from computer architecture to operating systems, from programming languages to database systems, from signal processing to information theory, from physical layer medium to wireless networking protocols, from algorithms to computational geometry, from graph theory to topology, and so on. In industry, sensor networks have drawn the interest and investment from tech giants such as Intel and Microsoft to many startups such as Crossbow, Ember, MillenialNet, and Dust Inc, among others. This level and breadth of activity underscores both the *intellectual* richness and complexity of these systems as well as their commercial potential.

Today, with nearly a decade of research and development behind it, the field of sensor networks is at a formative stage where the scope and complexity of many basic issues is well-established and early research prototypes have either validated some of the design principles or shown their limitations. These early efforts have also highlighted the need for inter-disciplinary research because many of the fundamental issues in sensor networks span across multiple areas that have traditionally had only limited collaboration and interaction.

With this backdrop and motivation, a workshop on Geometric Approaches to Ad Hoc and Sensor Networks was organized at University of California, Santa Barbara, during June 12-13, 2006, under the auspices of the National Science Foundation. The workshops goal was to provide a forum for some of the leading experts in areas most relevant to algorithmic and geometric foundations of sensor networks, to discuss their research, help identify important research challenges, and formulate a set of recommendations that can help advance the field and accelerate the adoption and deployment of sensor networks. The following report brings what the workshop participants believe are the most significant (algorithmic) research challenges in this endeavor. Based on this report, we created an executive summary, which also includes recommendations that, with the help of NSF initiatives, could provide a significant boost to the future scale, scope and adoption of sensor networks. The executive report, and the participants' position papers and slides of their talks, can be found in the workshop webpage at *http://www.cs.ucsb.edu/~suri/Workshop06/workshop*

2 Sensor Networks and the Role of Geometry

A key fact distinguishing sensor networks from other networked systems is that sensor nodes are deeply attached to the physical environment in which they function–*they are embedded systems*. As a result, geometry plays a fundamental and crucial role in all aspects of the sensor network, including their design and operation. In particular, the physical layout of the network deeply affects issues such as routing and information discovery; communication depends on node proximity, and node proximity in turn determines the correlation between sensed values and affects what information is necessary to transmit; the global structure of signal landscapes determines whether local greedy methods get stuck locally or can reach the desired global optimum; and tracking mobile phenomena requires migrating processes in the network and robust end-to-end connections between them. Unlike more traditional networks such as the Internet or the phone network, communication in sensor networks is dictated less by the desires of the end nodes and more by the geography of the sensor field and the associated signal landscapes, as well as the overall network task.

At the same time, the geometry of ad hoc and sensor networks is not as explicit as the geometry traditionally studied in computational geometry and related disciplines. Node locations are not always available, proximity does not always imply connectivity, and wireless link variability creates volatile connectivity graphs. There is a sense in the community that there is geometry in sensor networks "in the large"–though perhaps not at the scale of an individual node or two.

Motivated by these observations, the discussion during the workshop focused largely on techniques of a geometric or topological character that are particularly relevant to sensor networks. The following sections discuss the six areas of technical problems whose solution, we feel, can significantly advance the state of sensor networks.

3 Information Dissemination, Aggregation, Brokerage, and Summarization

Most early sensor networks were distributed data collection systems, where each node samples the environment and sends the signal(s) detected to a central base station. Communication bandwidth and energy limitations in the nodes have generated a lot of interesting research in distributed signal compression, in-network processing for data aggregation, and so on, so as to economize on scarce network resources while accomplishing this task. As sensor networks grow larger in size, however, and especially as we envision systems consisting of interconnected sensor networks spanning large regions or even the entire country, this notion of accumulating data at a single central station does not scale, nor can it provide the low-latency access to current data that the multiple users of such inter-networks may want or need. In such settings there will be multiple agents in the network, static or mobile, that seek actionable information in a dynamic environment. Each sensor node will face the decision of when and what to sense, and where to store past observations, so as to best serve the current network users — while at the same time conserving resources in order to maintain network coverage, connectivity, and longevity for future use.

3.1 Dissemination

Flooding or blind gossiping are expensive ways to disseminate information, as the same information is routed to the same node by multiple paths, creating medium access contention and wasting valuable energy through needless repetition. There are many opportunities to exploit geometric and topological knowledge about the network to perform such dissemination tasks more economically. A possibility is to use a communication pattern we call a *sweep* of the network, motivated by sweep methods in Computational Geometry. A sweep is akin to a wavefront that moves over the network and is guaranteed to pass over each node exactly once (as long as the network is connected), regardless of redundancy in the topology. The sweep is implemented by a narrow connected band of nodes that 'moves' over the network by issuing invitations to new nodes to join the band, and dropping nodes that no longer serve an essential purpose. The sweep can be designed so as to keep communication parsimonious and well-scheduled, avoiding MAC-layer collisions.

3.2 Aggregation

A sweep is also very appealing for data aggregation, as it does not depend on a fixed aggregation tree, like most extant database-inspired methods, such as TinyDB. Such methods are usually coupled with scheduling protocols that control the data flow so as the desired operation can be performed in an orderly manner. However, tree structures are not robust: any single node or link failure can disconnect the tree. Obvious approaches to this problem introduce problems of their own. Retransmissions over failed links can introduce delays, while the use of more richly connected routing structures, such as DAGs, raises issues of duplicate suppression and information over-counting. This can be ameliorated using duplicate-insensitive encodings and approximations, but the design of such methods is itself a challenging problem and must be thought through anew for each type of aggregation desired. In contrast, with a sweep there is no long-lasting data structure needed for a particular data collection/dissemination operation in a well designed topological sweep algorithm. Data will be collected/distributed correctly, even under modest instability in the wireless links.

A second issue about aggregation that needs more examination is how it can be made more resource-adaptive, in an application-sensitive way. Simply optimizing low-level communication primitives is not enough. Because transporting data from other nodes is both expensive and unreliable, each application must balance the cost of obtaining non-local data against their benefit. In particular, the algorithms must decide what data to fetch, how much of it, and from where. We believe that the ability to trade-off *computation for communication* at the application level can lead to dramatic performance gains.

3.3 Brokerage

In future networks most nodes, at different times, will perform all three of these functions: (1) data acquisition and local processing, event detection, etc., (2) query injection into the network for certain types of data, and (3) information routing and/or aggregation to aid in matching the requests to sensed data. We must provide mechanisms for information *consumers*, the nodes that pose queries and want information (also known as sinks), to be given access to relevant data collected by other nodes, the information *producers*, who have measurements and/or detections (also known as sources). This is the central problem of *information discovery, dissemination, and brokerage*. Though some stable data extracted from the sensor network could be stored externally and accessed through traditional database mechanisms, our emphasis here is on how to access up-to-the-minute information that is still inside the network, by users who may themselves be embedded and moving in the same physical space where the network is operating. This demanding scenario is made more realistic by technological advances, such as the advent of inexpensive large flash memories that allow significant permanent storage in each network node, as well as emerging sensor network applications that integrate sensors with actuators and human resources to provide intelligent monitoring, reasoning, and control over the environment.

Future research needs to address a number of fundamental questions about providing the diffusion of data hints and indices that facilitate information access as well as data replication for robustness, while still trying to conserve memory and network bandwidth as much as possible. Techniques must be developed that can be used within a single sensor network application, but the real pay-off comes from enabling data discovery and brokerage services that will allow many different applications to interoperate in ways that were not planned beforehand. The issue of data robustness and efficient access is central to the usefulness of sensor networks and deserves to be solved well. Although clearly there are application-dependent aspects to this issue, we believe that a small number of information access patterns can cover the vast majority of cases. By 'outsourcing' its data management tasks to a data management service, a sensor network application can get certain efficiencies that are not possible in solving this problem separately in each case, including the sharing of data across multiple applications, low-cost data recovery in case of faults, etc.

The key technical problem to be addressed is how to be highly selective in matching information providers and information seekers — but, unlike say the situation in peer-to-peer systems, in ways that respect the true communication costs of the network. In this we can heavily exploit the fact that a sensor network is embedded in physical space and employ techniques motivated by geometric, topological, or physical analogies. As a physical example, we are all familiar with following a sound gradient to arrive at a sound source, such as a water fall — and such methods have been used in sensor networks for access to non-physical quantities such as information. Indeed, these gradient methods for locating information have nice robustness properties. But imagine now there are a myriad of sounds (really, information) sources in the environment — can we follow the chirping of a particular bird in the cacophony of a rain forest? A variety of information coding techniques can be brought to bear so as to allow us to selectively follow only the gradients of information that we are truly interested in. Such coding techniques also add some measure of information security to the system. As another geometric example, 'road-systems' can be established in a network that naturally create rendez-vous points for information producers and information consumers, following the natural morphological features of the network.

3.4 Summarization

We propose that new summary structures be developed following a "geometric" approach that treats the sensed data as a "signal landscape" (a surface) and explicitly attempts to summarize its significant features (such as peaks, valleys, and areas of steep gradient). As sensor networks mature and applications become more sophisticated, adaptive geometric summaries that allow users finer control over the data views would be highly valuable. Such an approach has several advantages:

1. A "surface," modeled as a bivariate function, is a powerful abstraction for many different types of data that are likely to be focus of sensor-based remote observation.

2. Explicitly taking the geometry of the signal into account let the users exercise greater control over the features that they want to monitor.

3. There already exist a well-developed theory and set of techniques for approximating well-behaved surfaces that can be leveraged.

Broadly speaking, given a sensor network that samples a smooth physical function, we want to extract and maintain (as this function evolves over time) a sparse subset of sensor readings from which this function can be well approximated. Clearly there is a trade-off between the size of this subset and the quality of the approximation that can be obtained. From the monitoring point of view this problem becomes interesting when we can get a lot of information about the shape of our function from a small set of sensor readings. Since physical phenomena often exhibit a high degree of coherence and signal compression techniques have been widely and successfully deployed, we can hope that a sparse but informative samples will exist for many problems of interest. Although signal information can be compressed in a wide variety of ways, a sampled representation has several advantages: (1) the recorded values are directly meaningful to the application, (2) no additional distributed signal processing is necessary, and (3) drilling down to the network region where something of interest is going on can be supported.

One could attempt to use the "surface approximation" methods to build compact representations of the signal landscape, for instance, as a *piecewise linear function*. When the surface is convex, the problem is well-understood in the approximation theory and computational geometry. However, for non-convex surfaces, the problem becomes more challenging, but there is a substantial literature on Delaunay-based refinement and related techniques, which could be a good starting point for distributed and lightweight approximations in sensor settings.

4 Structure Discovery and Self-Organization

A sensor network is a decentralized computing engine that must first discover its own network architecture when initially powered on. Unlike most conventional computing systems, even this basic initialization process can be quite challenging for a sensor network due to the lightweight nature of its components and the variability present in node placement and network links. A proper understanding of this architecture, including network routing and information discovery, is essential for an efficient implementation of the higher level abstractions needed for the applications that use this computing engine. The architecture of sensor networks is also programmable to a large extent, and the system must adapt its structure to best deal with the data it is likely to sense and the functions it has to perform.

This architecture discovery and adaptation must occur at multiple levels of abstraction. The first level detects the *network topology*, which determines the basic communication architecture of the system. The second level adapts to the *signal landscape*, which deals with the structure of the raw data on which a sensor computer must operate and the sampling patterns most appropriate to the particular signal field. The third and final level is *distributed storage and information brokerage*, which deals with issues of managing information in the network and providing mechanisms by which information seekers can locate information providers.

The goal is to understand certain global aspects of the network structure, or the signal landscapes it monitors, and adapt or exploit those constraints. Examples of the former are the detection of network holes, narrow bridges or passageways, and other sensor layout features that affect the quality and load-balance of information delivery in the network. We may also be interested in estimating the large scale characteristics of relevant functions over the sensor network, whether these are network health related (such as node energy reserves or processing load), or attributes of the signal landscape observed by the network (number of temperature peaks, for instance). The challenge in these global estimation questions is that each node has only a local view of the state of the network and wide-area collaboration will be necessary to obtain meaningful and robust answers.

4.1 Node localization

The network structure is intimately related to the positions of the nodes. Furthermore, knowledge of node positions enables geographic routing and other efficient algorithms for routing and information discovery.

Self-localization is important in sensor networks because manual calibration is rarely feasible for large networks, and GPS hardware at every node is not practical due to cost, power consumption, or large form factor. In self-localization, nodes use inexpensive ranging devices to estimate distances (or angles) among neighbors, and then deduce global coordinates from this partial distance matrix. This is the classical problem of *graph embedding*, with a long history in graph theory, rigidity theory, distance geometry and topology. Not surprisingly, the general problem is intractable and, indeed, it has been shown that even with the unit-disk model, the self-localization problem in sensor networks is also NP-complete. This inherent complexity is not merely academic: all current localization methods fail to localize parts of the network, even for relatively small network sizes and well-distributed nodes.

Despite the fundamental nature of the localization, the problem still lacks a satisfactory solution. We believe a closer collaboration between geometers, mathematicians, and networking researchers is needed to accelerate progress on this key problem. When confronted with a difficult problem, engineers are especially good at "changing" the problem. In sensor localization, this can occur through the use of additional machinery, such as beacons, mobile nodes, or through incremental localization. Research on these "relaxed" forms of embedding questions is only recently begun, and new theoretical as well practical questions are emerging. For instance, how best to place the beacons; how best to use mobile nodes; how to minimize additional measurements in incremental localization, and so on.

4.2 Virtual coordinates

In other scenarios, nodes do not even have a way of measuring coordinates, geometric distances to other nodes, or their direction. Their only way of interacting with other nodes is to send or to receive messages from any node that is within communication range. It is important to note that for many of these scenarios, computing actual node coordinates may not be inherently vital, and can indeed turn out to be a red-herring chase. Instead, *virtual coordinates* may be much simpler and even better suited for the tasks at hand. Although extensive work has been done in obtaining virtual coordinates, current methods are neither efficient nor robust. Effectively, virtual coordinates require a global embedding of the communication graph in the plane and such algorithms work well only for dense and uniform networks where communication graph distances are good estimators of the underlying (but unknown) Euclidean distances. More approaches need to be explored, including collections of local embeddings, coupled with alternative, high-level routing mechanisms.

4.3 Sensor field morphology

By sensor field morphology we mean an understanding of the sensor field layout, including the identification of boundaries, both external and those of network holes, the detection of narrow passages and other communication bottlenecks, etc. Such analyses can greatly benefit from knowledge of the node positions, but recent research has shown that, surprisingly, they are also quite viable in settings where only the node communication graph is known. Morphological understanding is key to performing load-balanced routing and may other essential tasks of the network. The presence of a hole can simply indicate the presence of some physical obstruction where nodes could not be placed (e.g., a pond) —

or it could signal a region where something bad is going on, such as an area where nodes have depleted their resources or were destroyed. We believe that topological methods (in the sense of the so-named part of mathematics) will play a fundamental role here, as they are by nature both global and robust. Ultimately, a proper morphological understanding can lead to a segmentation of the sensor field into regions of simple shape and connectivity, in which simpler protocols can work effectively. A second interesting outcome of such topological analyses is that each node can have a little knowledge about how its position or role in the network is special (i.e., whether it is near a boundary, in a narrow bridge, and so on). This in turn can allow each node to make better local decisions in executing its functions. This area needs additional research, so as to couple simple reactive local protocols that require nice layouts within certain regions, with a higher-level global understanding that can aid the overall planning of resource usage in the network.

4.4 Clustering and other hierarchical structures

Related to the above is the development of hierarchical representations of the network, such as cluster hierarchies, that can be exploited for node naming, routing, information discovery, and so on. How to make such hierarchies robust to the underlying wireless link volatility remains a major challenge.

5 Sensor placement and coverage

At the interface of problems involving sensor capabilities [a problem of engineering] and network communication protocols [a problem of information] lies the problem of sensor placement and coverage. Such problems, which range from determining optimal placement and allocation of sensor nodes to adapting a given distribution of nodes to the desired task, are uniquely geometric in nature. One cannot decouple sensing or communication issues from the geometry of the embedded network.

5.1 Motivations and Impact

Numerous and diverse real-world applications will benefit from algorithmic and mathematical developments in sensor placement, coordination, and coverage. A few such hi-impact scenarios are as follow:

5.1.1 Algae Biomass Monitoring

The study of aquatic ecosystems has been of great interest amongst Biologists. Of particular interest is the study of biomass, including the algal blooms. Abundance of algae in aquatic systems can block sunlight to underwater vegetation, consume oxygen in the water, and produce surface scum and odors. Hence this nuisance algal bloom could impair the beneficial use of aquatic system. Recently, some studies have been performed [10] to characterize the algae content in Lake Fulmor at James Reserve [23]. These studies employed Networked Aquatic Microbial Observing System (NAMOS) [7] that consist of a small set of boats and buoys equipped with Fluorometer and temperature sensor. These deployments examined the spatial heterogeneity of temperature and chlorophyll content, which can be used to predict algae biomass. Impending deployments will include a larger number of sensors, both static (buoys) and mobile (boats), along with robotic sensors that traverse a cable and can descend to different depths in the lake. In this larger deployment, the positioning of the buoys, paths of the boats and motion of the NIMS must be jointly optimized.

5.1.2 Sensing Light Intensity: understanding photosynthesis

Understanding the photosynthesis of understorey plants in dense forests is another application of interest to biologists [49]. Several studies show that sunfleck (a region covered with sunlight with area around covered with shadow) frequency and duration affects the growth rate of such plants [47]. Existing deployments have provided coarse data from only a handful of sensed locations. Due to a dense canopy, the sunfleck distribution exhibits high spatial and temporal variation which makes it complex to study using a few static sensors. It is therefore necessary to determine algorithms for mobile sensor placement and coverage.

5.1.3 Water Distribution Network: detecting dangerous pathogens

Water distribution networks, the network of pipes that bring water to our taps, are highly susceptible to the development of dangerous pathogens and to malicious intrusions. Since pathogens can be introduced at arbitrary locations in the network and at arbitrary points in time, there is a tremendous number of possible attack scenarios. Fortunately, the tampering with water distribution can be detected using sensors placed in the distribution network. Since every sensor is expensive (roughly $ 5,000 per sensor), the strategic choice of locations is of utmost importance. Optimizing sensor placements in such scenarios poses at least two fundamental challenges: the evaluation of the quality of any given sensor placement, and the optimization over a combinatorial number of such placements. The quality of a sensor placement is based on various performance criteria, such as average time to detection, population affected by an intrusion, consumption of contaminated water and the fraction of detected scenarios (cf., [39]).

5.1.4 Surveillance and Search

Though environmental surveillance has been the most active domain of sensor network implementation, there are increasing opportunities and demands for sensor networks in surveillance for purposes of homeland security, national defense, and search/rescue. These problems range in scope and modality from coordinating [mobile] UAV's to placing [stationary] acoustic sensors for detection. The scale of such problems spans from very localized (e.g., a search and rescue operation in a mine) to very global (e.g., sensor deployment across national borders). The impetus to develop rigorous data about coverage and sensitivity of a network is especially severe in these settings.

5.1.5 Sensor placements and Autonomic Computing

Autonomic computing seeks to provide the ability to automatically configure, adjust, repair and optimize large-scale complex computer systems. This in turn requires effective, cost-efficient instrumentation of these large-scale systems. Unfortunately, the amount of data that can be extracted from complex systems can overwhelm the ability of state-of-art machine learning algorithms to learn effective management policies. Consider, e.g., the data that can be collected in a large-scale Internet Data Center: this may range from very fine-grain statistics on individual requests or individual Java threads, to very coarse aggregate performance statistics over entire web applications. Moreover, due to highly dynamic traffic and system behavior, monitoring of such statistics may be required on a multiplicity of time scales. Additionally, in such Data Centers it is critically important to monitor and control power consumption and heat generation on fine-grained spatial and temporal scales. Sensor placement and selection methods can also provide solutions for these challenges.

5.2 Fundamental Problems and Challenges

The following general problems are distillations of the most critical coverage and placement problems into forms most amenable to mathematical or statistical analysis.

5.2.1 Stationary Coverage

Assuming each sensor node has some fixed domain of sensing or radio coverage, how can one tell if the union of these domains over all active nodes contains the entire region to be monitored [the 'blanket' coverage problem] or surrounds/separates the appropriate region [the 'barrier' coverage problem]? What is the optimal placement to maximize available coverage?

5.2.2 Noisy Sensors and Uncertainty

In real settings, sensors make noisy, uncertain observations. This uncertainty fundamentally affects the problem: We cannot view sensors as making exact observations in a region, but rather as providing noisy predictions. How can we exploit statistical models of the sensor noise and of the spatial distribution of the phenomena to find a placement that maximizes the probability of detection, or the amount of "information" provided by the sensors?

5.2.3 Active learning

Consider, for example, the problem of understanding algal blooms in a new lake. When we start the deployment we know very little or nothing about the structure of the underlying domain. In this setting, as we place sensors, we are learn something about the current state of the lake, and about the statistical distribution of algal blooms in this space. How can we optimize sensor placements when we know very little or nothing about an environment, but we can incrementally introduce new sensors as we learn the properties of this environment? This general, very challenging version of sensor placement is a core long-term problem that we must address.

5.2.4 Robustness to failures and adversaries

Low-cost sensor nodes are prone to failures. In deliberate attacks to our water networks, an adversary may select attack locations that are difficult to sense. How can we guarantee that our placements are robust to these situations?

5.2.5 Communication Constraints

Thus far, the discussion in this section focuses only on selecting the location of sensors in order to maximize the amount of information collected or coverage. When using wireless sensors, however, not only should the sensors be informative, but they should also be able to communicate efficiently. Sensor with high coverage tend to be spread apart, leading to poor communication. On the other hand, sensors that are close together provide better communication, but poor coverage. How do we design algorithms with provable quality guarantees that optimizes this tradeoff?

5.2.6 Mobile Coverage

Assuming that nodes are in motion as a function of time (and perhaps including going off-line and coming back on-line), can one determine time-dependent coverage properties? Do the coverage domains properly sweep the region? What is the appropriate strategy for mobile nodes with intrinsic dynamic and/or kinematic constraints in order to maximize coverage?

5.2.7 Power Constraints and Adaptive Allocation

Given the limited battery life of the typical ad hoc network node, what is the appropriate sleep/wake schedule for maintaining or optimizing coverage? Given the ability to add nodes to an existing network, which locations will best satisfy both coverage and power constraints? Where are the 'holes' in network coverage to which nodes are most critically needed? Where is there an excess of nodes?

5.3 Mathematical Tools

The mathematical tools available for solving problems of coverage and placement range across many disciplines, not all of which are currently emphasized in the wireless ad hoc networks community.

5.3.1 Computational Geometry

As with problems involving structure discovery and localization, the most natural approach to solving coverage and placement problems involves a recourse to computational geometry. The classical Art Gallery Problem [38] is the abstract forerunner of current problems in optimal placement. Many of the standard geometric methods (Voronoi decompositions, Delaunay triangulations, etc.) for coverage estimation and verification are abundantly present in simplified network models [34, 21]. In particular, the combinatorial and geometric properties of unit disc graphs (graphs whose vertices correspond to sensor nodes and whose edges correspond to nodes within a fixed broadcast distance) provide the most basic model for capturing communication range constraints. Although solving many localization problems within this category of graphs has high computational complexity in general [4], progress appears possible with a mild strengthening of the data structure [5, 15].

5.3.2 Algebraic Topology

Most current research in sensor networks requires having geometric data about the node positions — either absolute coordinates, or relative distance measurements. Results from algebraic topology suggest that a lack of relative or global localization in a network is not an obstruction to determining global network features. Recent results [46, 45] indicate that homology theory provides a powerful yet computable set of criteria for coverage in idealized networks. It is possible that many of the tools developed by topologists over the past century for passing from local combinatorial data to global topological data (e.g., the Mayer-Vietoris principle, the Excision principle, spectral sequences, etc.) may provide valuable insights in extracting a global environmental picture from nodes with local communication links.

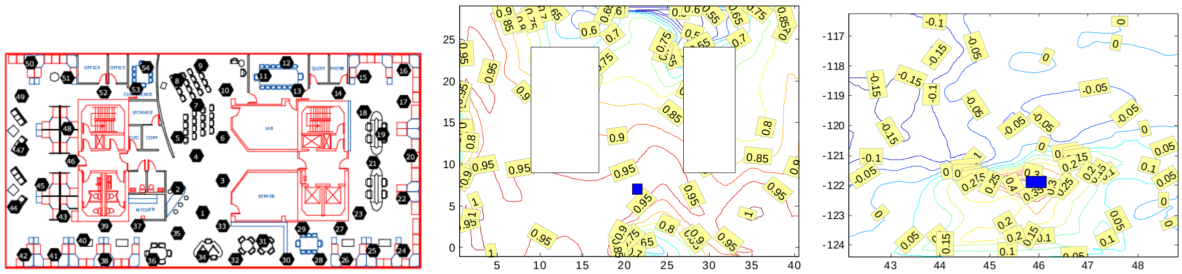

(a) *54 node sensor network deployment* (b) *Temperature correlations* (c) *Precipitation correlations*

Figure 1: *Correlation between a sensor placed on the blue square and other possible locations for: (b) temperature data from the sensor network deployment in Figure 1(a); (c) precipitation data from measurements made across the Pacific Northwest.*

5.3.3 Statistical Models

One approach for sensor placement is to solve the task as an instance of the art-gallery problem (*c.f.*, [20, 16, 11]). In practice, however, this assumption is too strong; sensors make noisy measurements about the nearby environment, and this "sensing area" is not usually characterized by a regular disk, as illustrated by the temperature correlations in Figure 1(b). Furthermore, note that correlations can be both positive and negative, as shown in Figure 1(c). Often, correlations may be too weak to enable prediction from a single sensor, suggesting the need for combining data from multiple sensors to obtain accurate predictions.

An alternative approach from spatial statistics [9], making weaker assumptions, is to use a pilot deployment or expert knowledge to learn a *Gaussian process* (GP) model for the phenomena, a non-parametric generalization of linear regression that allows to represent uncertainty about the sensed field. For example, using a GP, from the 54 temperature sensors in Figure 1(a), we obtain a prediction of the temperature field throughout the environment (even where there are no sensors), Figure 2(a), and an estimate of the uncertainty in this prediction, Figure 2(b).[1]

As an example of how statistical models can be used for optimizing sensor placement, consider the recent work of Guestrin et al. [18] who proposed an algorithm with formal guarantees for finding sensor placements based on GPs, by finding positions that are most informative about unsensed locations. Figure 2(c) shows that this criterion leads to better placements than disk models, see [18] for several other comparisons.

5.3.4 Optimization Theory

New developments in optimization theory can also provide tools for addressing sensor placement. For example, *semi-definite programming* has been applied to several placement problems [3]. [18, 27] use combinatorial optimization techniques for optimizing placements, for several versions of the problem, including coverage, communication constraints and mobile sensors. In particular, they not that many sensor placement problems require the maximization of a *submodular* function [36]. Submodularity is a property of set functions that intuitively represents "diminishing returns": adding a sensor y when we only have a small set of sensors $\mathcal{A}$ gives us more advantage than adding y to a larger set of sensors $\mathcal{A}'$.

These two examples are just some of the possible connections between optimization theory, geom-

[1]Note, for example, that in the central areas of the lab where no sensors where placed, our estimate in Figure 2(b) has high uncertainty.

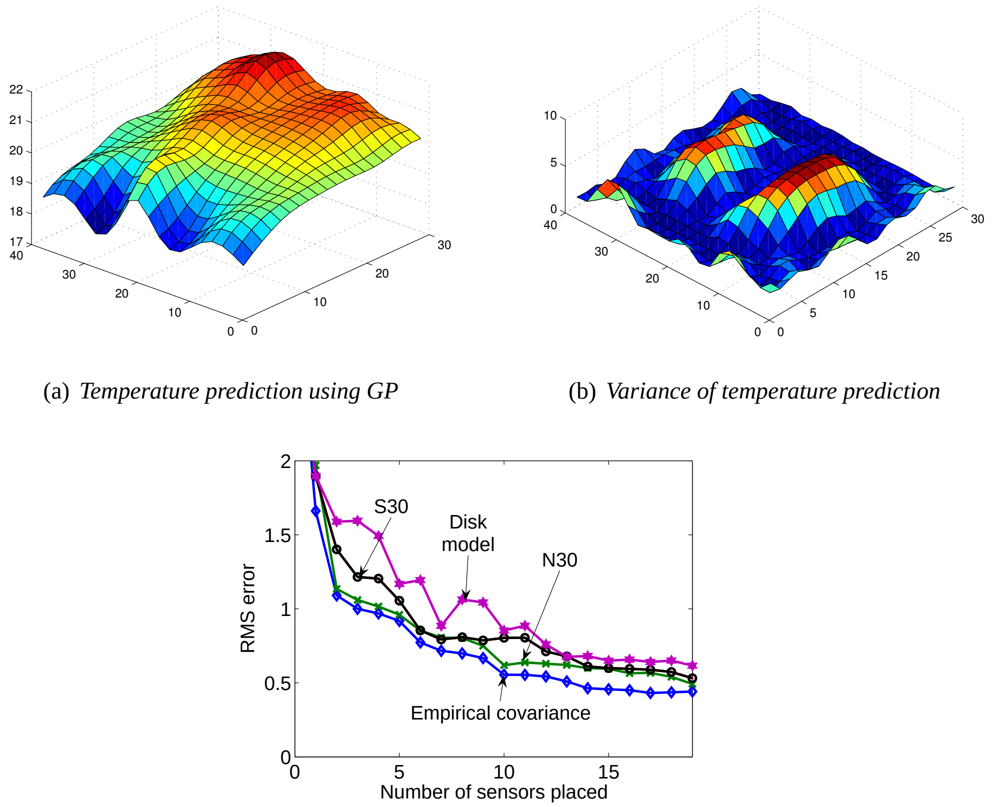

(a) *Temperature prediction using GP*

(b) *Variance of temperature prediction*

(c) *Comparison with disk model*

Figure 2: *Posterior mean and variance of the temperature GP estimated using all sensors. (a) Predicted temperature (b) predicted variance. (c) Prediction accuracy for sensor placements selected according to disk model, and mutual information criterion with different GP models: S30 – a stationary GP learned from 30 sensor locations, N30 – a non-stationary GP learned from 30 sensor locations, Empirical Covariance – learned Gaussian model over test locations (Disk Model used empirical covariance for prediction).*

etry and sensor placement problems that should be explored further.

6 Naming and Routing

A key design issue in most networks is that of node *addressing* – assigning nodes with labels that designate their position or location within the network. Telephone numbers, IP addresses, postal addresses, *etc.*, all represent labeling schemes that help identify, locate and establish communication with entities within a network. The unique characteristics of sensor networks however requires rethinking the traditional role of (and indeed, even the need for) addressing in these networks. For example, sensor networks often require easy network deployment and initialization and are hence less amenable to carefully configured addresses (as typical with IP or telephony networks) while the data-centric nature of sensornet applications leads to addressing schemes that primarily assist data collection and processing rather than the pairwise communication that traditional Internet addressing focuses on.

A perusal of the sensornet research literature to date reveals three broad uses for global addressing. The first is to support routing or communication between nodes. For example, gradient or tree-based addressing has been used for data collection protocols such as Directed Diffusion [22] and TAG [33] while landmark or geographic addressing provides routing support for data-centric storage [44]. A more esoteric application uses the node address space as a global frame of reference by which to decompose a network-wide task into multiple, localized tasks. For example, various algorithms for spatio-temporal[17] or multi-resolution indices[14], multi-dimensional querying[32] and distributed leader election[43] all employ a hierarchical decomposition of the network's geographic extent as a means of subdividing a global task. Finally, addresses can serve to establish a node's position in the physical world which is useful in identifying the physical source of sensed data when sensing nodes are mobile or the placement of nodes is unplanned.

6.1 Location-based naming and routing

At first glance, physical geography offers an elegant solution to the above addressing needs — using GPS or other geo-positioning devices, nodes can independently discover their position/address in the physical world, coordinate routing algorithms offer simple communication solutions (*e.g.*, GFG[2], GPSR[24], GOAFR[30]) and, a geographic address space offers a well-defined frame of reference by which to decompose tasks (as in DIM[32], GHT[43], Dimensions[14], *etc.*). Despite the elegance of the fit, the adoption of geographic techniques on existing hardware platforms has been hampered by concerns on two fronts. First, the cost, power consumption, and usability in mobile or indoor environments of geo-positioning devices appears to have raised the bar to their use in real deployments. More unexpectedly, empirical studies [13, 25, 50] have repeatedly shown that physical proximity and robust wireless connectivity are not always congruent which violates a necessary assumption underlying traditional geographic routing algorithms.

Ongoing research aims to address these concerns. Schemes such as CLDP[25] and GDSTR[31] continue to rely on the use of geographic addressing but propose novel routing algorithms that can tolerate incongruencies between physical distance and connectivity. This approach, together with continued developments in GPS technology, could enable geographic addressing to serve our needs in a robust and low-cost manner.

6.2 Virtual coordinates and routing

Location-based addressing and naming typically undertake an implicit assumption that the sensors are deployed uniformly densely in a regular region. Thus sensors' geographical location correlates well with the real network connectivity. This is no longer true when the sensor field is too sparse, has holes or complex shape. Two nodes that are geographically close may actually be far away in the connectivity graph. Thus, when these topological features (e.g., holes) become prominent, the naming and its coupled routing protocol should represent the real network connectivity, rather than the geographical locations, and adjust to these topological features accordingly. In the literature, uniformly randomly deployed sensors is arguably the most commonly adopted assumption on sensor distribution — but is rarely the case in practice.

This motivates the approach of eschewing the use of geography altogether, and construct virtual coordinate systems from only the measured connectivity between the nodes (e.g., in schemes such as GEM[37] and NoGeo[41]). Virtual coordinates offer a cheap addressing solution for the purpose of routing and as a network-wide frame of reference. Additional work along this line propose to investigate the topology of a sensor field and adapt node naming and network routing accordingly. Here by topology we refer to the algebraic topology such as holes or high-order features.

A second motivation to investigate the topology of a sensor field is that the topological features mostly reflect the underlying structure of the environment (e.g. obstacles, buildings, etc), therefore they are typically only a few in number and likely to remain stable. In most cases isolated or sporadic link failures will not change these large topological features. Thus we can afford to explicitly compute an abstraction of these large features (e.g., holes) and store this compact structure at each node to facilitate routing or information processing in such a sensor field. For example, we can carry out proactive protocols at this abstract level, which is stable and compact, such that this high-level combinatorial guidance can be realized with localized and reactive protocols at sensor nodes with high link volatility. Landmark-based gradient descent routing [12] and medial-axis based road map [6] are two such examples. Both algorithms apply local greedy routing rule accompanied by a global routing guidance on how to route around holes.

6.3 Challenges

The current state of addressing/routing is sufficient for only fairly primitive uses of sensor networks. If we are to achieve the vision of highly scalable sophisticated sensor networks, then we are going to need correspondingly more scalable routing support and geometric-based approaches appear to be the most promising in this regard. The challenge will be finding geometric-based solutions that are simple and yet robust to (i) location inaccuracies; (ii) violations of traditional assumptions made in these algorithms, such as unit disk graph; (iii) network deployment dimension (2D or 3D); (iv) link volatility and node failure.

Finally, much of the above discussion stems from a fairly narrow view of a sensor network as a standalone deployment of homogeneous nodes tailored to a single well-defined task. Broadening our view to encompass networks in which heterogeneous sensors are more generally deployed within a larger interconnected world – mobile phones and cameras in an urban environment, sensors in a home, vehicular networks and so forth – adds additional considerations. For example, a closer coupling with the wired Internet and its users and applications would require addressing schemes that focus on interoperability with, and management within, the public Internet. The requirements of, and options for, addressing in such highly-networked environments adds an interesting dimension to the role of addressing in sensor networks that is as yet not well understood.

7 Sensor Networks and Mobility

In the wireless communications literature, "mobility" conventionally refers to unintentional or uncontrolled motion (*e.g.* the movements of cell phone users). We are mainly interested in systems at the opposite end of the spectrum - those in which mobility is controllable to a reasonable extent (*e.g.* sensor nodes and radios are the payload on mobile robots). However, in some cases, it is also useful to exploit unintentional or serendipitous motion.

We assume that in the future we will see more systems that are combinations and collaboration between static sensors and autonomous mobile robots. These systems can have applications in environmental measurements and monitoring, military and surveillance applications, and many others.

We believe that to benefit from the whole spectrum of potential that combines robotics and sensors networks, several research challenges need to be addressed. These challenges can be categorized into different families of problems, which in turn depend on the level of collaboration between robots and static sensors, and robots and other robots. We discuss this in detail below. All these problems require novel solutions, as well as combinations of new and existing techniques in sensor networks and robotics.

We present the challenges hierarchically, where addressing the challenges of one family requires (full or partial) solutions to the families above it.

First family: Here we assume that a single robot is used, and it carries sensor(s) onboard, so it functions as a mobile sensor(s). The main challenges that we predict are a combination of path and motion planning while interpolating (or otherwise interpreting) the data gathered. These problems are of particular interest when the path of the robot needs to be updated online, according to the sampled data from the sensors. For example, the sampled data might signal that more samples are required in some areas rather than others. Also, special attention is needed when the environment is not fully known to the robot, *i.e.* it is studying and mapping this environment via SLAM (Simultaneous Localization and Mapping) methods. In addition, it is likely that the robot will need to communicate with a static base station. In this case we have effectively a Simultaneous Localization, Mapping, and Exploration problem with communication constraints.

Second family: Here we assume that robots, in addition to carrying sensors onboard, communicate with static sensors. This communication and interaction could take several forms. First, robots may act as relay stations for the static nodes. Second, robots may act as data mules for the static nodes (*i.e.* data are transported by robots ferrying datagrams back and forth). Third, the static nodes may act as 'signposts' for the robots, and provide navigation and path planning 'advice' or 'directives'. Fourth, robots may carry energy to nodes to replenish them, and finally, robots may pick up, drop off or reposition static nodes in order to physically reconfigure the network. All these imply dealing with questions of optimal path planning for the robots in a way that enables them to pass close enough to the sensors. It also implies designing routing protocols that would enable sensors to send data to other nodes which are close enough to the robots' paths. This setting is squarely at the intersection of conventional (*i.e.*static) sensor networks and multi-robot systems.

Third family: Here we assume that we are using a swarm of inexpensive robots. All, or the majority of, the robots are inexpensive and have limited localization and computational capabilities. However, some of them might have additional capabilities and be used as shepherds to the others. The motion, reasoning, data aggregation, communication with static sensors and other tasks performed by a swarm should be accomplished collaboratively. Essentially this is a multi-robot setting with communication constraints.

In the last few years, some preliminary attempts to address these challenges appeared in the literature. The algorithmic issues have been studied mainly using simulations, but quite a few working prototypes exist. Nicholas Roy and his colleagues use sensors to help robot localizations. James McLurkin is using a swarm of 112 robots with simple sensing and communications abilities to accomplish different exploration and sensing tasks. As an example of unintentional motion, Samuel Madden and his colleagues use sensors installed on cars to collect environmental data.

8 Wireless network models

8.1 Introduction

Wireless network models much like previous models of the Internet and the PSTN (Public switched telephone network) are invaluable tools for network researchers. Such models are used to develop algorithms for sensor networks and to give mathematical proofs of their correctness and performance. Furthermore, models of the wireless medium are used in simulation to test the performance of proposed protocols before they are implemented and tested on actual Wireless Sensor Network (WSN) deployments.

As with previous models, there is an inherent tension between deriving a model that simplifies reasoning about the behavior of distributed WSN algorithms and a model that captures all the pertinent properties of the network. A great algorithm in theory may be inefficient or even incorrect in practice if the analysis is based on idealistic assumptions. For example, an algorithm which ignores interference may fail in practice since communication happens over a shared medium.

Many models for sensor network have their origin in theoretical computer science and applied mathematics. Since the topology of a sensor network can be regarded as a *graph*, the distributed algorithms community uses models from *graph theory*, representing nodes by vertices and wireless links by edges. Another crucial ingredient of sensor network models is *geometry*. Geometry comes into play as the distribution of nodes in space, and the propagation range of wireless links, usually adhere to geometric constraints.

In what follows we present a series of models each capturing the behavior of wireless networks with increasing levels of accuracy. While we do not believe that there is a single wireless network model that is appropriate for all cases, we argue that algorithm designers and protocol developers must understand the limitations of the models they use and so select the ones that best matches their requirements.

8.2 Models

The models we present are organized as follows. In Section 8.2.1, various models for the network's *connectivity* are outlined. Connectivity models answer the question: Which nodes are "connected" to which other nodes and can therefore directly communicate with each other. Section 8.2.2 then enhances these connectivity models by adding *interference* aspects: Since sensor nodes communicate over a shared, wireless medium, a transmission may disturb a nearby concurrent transmission. Finally, in Section 8.2.3 we sketch models that describe power consumption due to message transfers.

8.2.1 Connectivity

A fundamental modeling question concerns the *connectivity* of sensor nodes: Given a set of nodes distributed in space, we need to specify which nodes can receive a transmission of a node. Throughout this section, if a node u is within a node v's transmission range, we say that u is *adjacent* to v, or equivalently, that u is a *neighbor* of v.

Typically, the connectivity of a sensor network is described by a graph $G = (V, E)$, where V (*vertices*) is the set of sensor nodes, and E (*edges*) describes the adjacency relation between nodes. That is, for two nodes $u, v \in V$, $(u, v) \in E$ if v is adjacent to u. In an undirected graph, it holds that if $(u, v) \in E$, then also $(v, u) \in E$, i.e., edges can be represented by sets $\{u, v\} \in E$ rather than tuples.

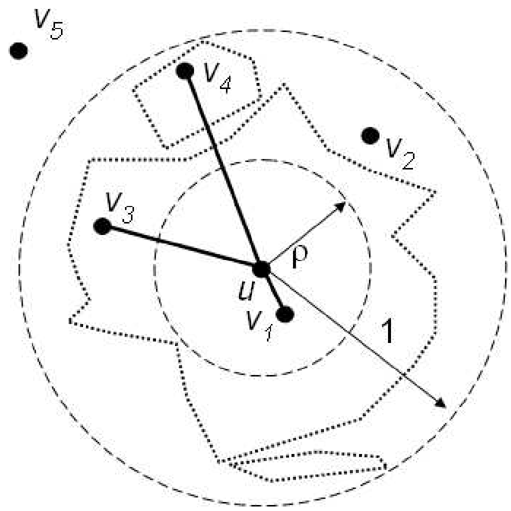

Figure 3: Quasi unit disk graph from the perspective of node u: Node u is always adjacent to node v_1 ($d(u, v_1) \leq \rho$) but never to v_5 ($d(u, v_5) > 1$). All other nodes may or may not be in u's transmission range. In this example, node u is adjacent to v_3 and v_4, but not to v_2.

The classic connectivity model is the so-called *unit disk graph* (UDG) [8]. In this model, nodes having omnidirectional radio antennas—*i.e.*, antennas with constant gain in all directions—are assumed to be deployed in a planar, unobstructed environment. Two nodes are adjacent if and only if they are within each other's transmission range (which is normalized to 1).

Model 8.1 (Unit Disk Graph (UDG)). *Let $V \subset \mathbb{R}^2$ be a set of nodes in the 2-dimensional plane. The Euclidean graph $G = (V, E)$ is called* unit disk graph *if any two nodes are adjacent if and only if their Euclidean distance is at most 1. That is, for arbitrary $u, v \in V$, it holds that $\{u, v\} \in E \Leftrightarrow |u, v| \leq 1$.*

The UDG model is quite idealistic: in reality, radios are not omnidirectional, and even small obstacles such as plants could change connectivity. The *quasi unit disk graph model* (QUDG) [1, 29] is a generalization of the UDG which takes imperfections into account as they may arise from non-omnidirectional antennas or small obstacles.

Model 8.2 (Quasi Unit Disk Graph (QUDG)). *The nodes are in arbitrary positions in $\mathbb{R}^2$. All pairs of nodes with Euclidean distance at most ρ for some given $\rho \in (0, 1]$ are adjacent. Pairs with a distance larger than 1 are never in each other's transmission range. Finally, pairs with a distance between ρ and 1 may or may not be neighboring. An example is shown in Figure 3.*

The QUDG as presented in Model 8.2 does not specify precisely what happens if the distance is between ρ and 1. There are several options. For example, one could imagine an *adversary* choosing for each node pair whether they are in each other's transmission range or not. Alternatively, there may be a certain *success probability* of being adjacent: the corresponding probability distribution could depend on the time and/or distance [48]. For example, the QUDG could be used to study *Rayleigh fading*, that is, the radio signal intensity could vary according to a Rayleigh distributed random variable. Also a *probabilistic on/off model* is reasonable, where in each round, a link's state changes from good to bad and vice versa with a given probability.

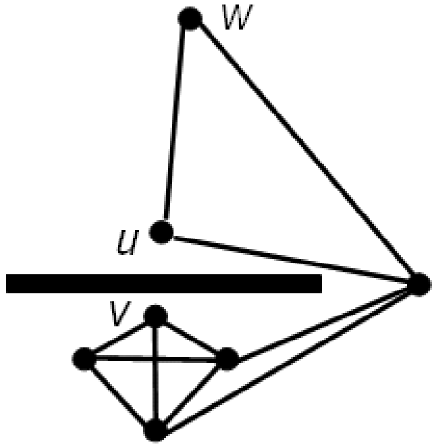

Figure 4: Nodes u and v are separated by a wall. Nodes on the same side of the wall are completely connected. However, due to the wall, although u can reach a distant node w, it cannot hear the close node v. Such situations can be modeled by the BIG, but not by the UDG or the QUDG.

While the QUDG can be attractive to model nodes deployed in fields with few obstacles, it does not make sense for inner-city or in-building networks where obstructions cannot be ignored: Since a node may be able to communicate with another node which is dozens of meters away, but not with a third node being just around the corner, ρ would be close to 0.

However, even in such heterogeneous environments, the connectivity graph is still far from being a general graph. Although nodes which are close but on different sides of a wall may not be able to communicate, a node is typically highly connected to the nodes which are in the same room, and thus many neighbors of a node are adjacent. In other words, even in regions with many obstacles, the total number of neighbors of a node which are not adjacent is likely to be small. This observation has motivated the **Bounded Independence Graph** model [28]. Figure 4 shows a sample scenario with a wall; in contrast to UDG and QUDG, the BIG model captures this situation well.

Recommendations. In closing, we present our recommendations regarding the models that should be used accurately capture the intricacies of the wireless medium.

First of all, since all WSNs are deployed in the physical environment which is inherently 3-dimensional, propagation models should consider nodes in $\mathbb{R}^3$. Moreover, in addition to Euclidean distance ($\mathbb{L}_1$ norm), distances between nodes could be modeled using the *Manhattan norm* ($\mathbb{L}_1$ norm) or the *maximum norm* ($\mathbb{L}_\infty$ norm), where $d(u, v) = \max\{|x_2 - x_1|, |y_2 - y_1|\}$. Besides omnidirectional antennas, there is a wide range of more sophisticated antenna models. For example, a node can have a directional radio antenna with more gain in certain directions. Given imperfections of the radio hardware transmission ranges of different nodes are not identical; one implication of this reality is the existence of *asymmetric* links (*i.e.* links in which u can hear v but not vice versa). Transmission models should attempt to take such variations across nodes into consideration. Given the large impact of obstacles to connectivity, *site specific* propagation models should be constructed, if the actual location in which a network will be deployed is known.

Finally, as mentioned in the discussion of the QUDG, links are not reliable: links may be up and down, *e.g.*, according to a stochastic process. The models must incorporate this reception uncertainty

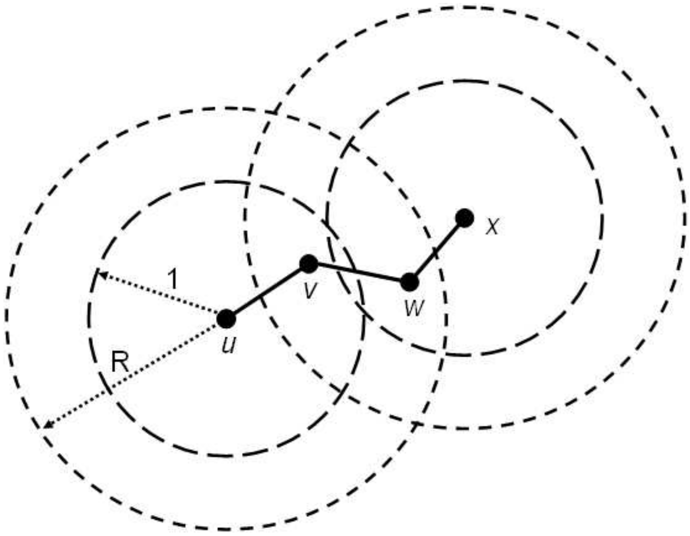

Figure 5: The UDI model has two radii: a transmission radius (length 1) and an interference radius (length $R \geq 1$). In this example, node v is not able to receive a transmission from node u if node x concurrently transmits data to node w—even though v is not adjacent to x.

preferably using data collected from experiments to derive the parameters of the stochastic process(es).

8.2.2 Interference

In wireless networks, the communication medium is shared, and transmissions are exposed to interference. Concretely, a node u may not be able to correctly receive a message of an adjacent node v because there is a concurrent transmission going on nearby.

The simplest interference model is then:

Model 8.3 (UDG with Distance Interference (UDI)). *Two nodes can communicate directly if and only if their Euclidean distance is at most 1,* and *if the receiver is not disturbed by a third node with Euclidean distance less or equal a constant* $R \geq 1$.

Figure 5 shows an example of UDI.

The *protocol model* (PM) is an extension that incorporates the effect of multiple competing transmissions [19].

Model 8.4 (Protocol Model (PM)). *Let* $u_1, u_2, ..., u_k$ *be the set of nodes transmitting simultaneously to receivers* $v_1, v_2, ..., v_k$ *respectively. The transmission of* u_i *is successfully received by* v_i *if for all* $j \neq i$*, it holds that* $d(u_j, v_i) > \lambda \cdot d(u_j, v_j)$*, where* $\lambda \geq 1$ *is a given constant. That is,* v_i *must not fall into a "guard zone" around any sender* u_j *which is a factor* $(1 + \lambda)$ *larger than* u_j*'s transmission range.*

The idea of quasi unit disk graphs (cf Model 8.2) could be adopted in these cases as well. For example, the UDI can be "quasified" as follows: if two nodes are closer than a given threshold R_1, concurrent transmissions will always interfere; if the distance is larger than a second threshold R_2, there will be no interference. Finally, if the distance is between R_1 and R_2, transmissions may or may not interfere.

All interference models above can be described in terms of the so-called *physical* or *SINR model* [19, 35, 42],. In this model, the successful reception of a message depends on the received signal strength, the ambient noise level, and the interference caused by simultaneously transmitting nodes.

Model 8.5 (Signal-to-Interference Plus Noise (SINR)). *Let P_r be the signal power received by a node v_r and let I_r denote the amount of interference generated by other nodes. Finally, let N be the ambient noise power level. Then, a node v_r receives a transmission if and only if $\frac{P_r}{N+I_r} \geq \beta$. The value of the received signal power P_r is a decreasing function of the distance $d(v_s, v_r)$ between transmitter v_s and receiver v_r. M*

8.2.3 Power Consumption

According to a popular model the energy consumed by a node is calculated by the sum over all its transmissions. Thereby, the energy needed to transmit one message is of the form $c \cdot d^{\alpha}$, where d is the distance between sender and receiver, α is the path-loss exponent (usually $\alpha > 2$), and c is a constant.

Although transmitting data is a costly operation, sensor nodes with short-range radios available today spend as much energy receiving or waiting for data. Therefore, techniques have been developed which allow nodes to change to a parsimonious *sleep mode* [40]. During the time periods a node is sleeping, it cannot receive any data. The idea is that if all nodes can somehow be synchronized to wake up at the same moment of time to exchange data (*e.g.*, every minute), much energy is saved. This motivates the following model.

Model 8.6 (Sleeping Time). *The energy consumed by a node is given by the accumulated time in which it is not in sleep mode.*

If there are no external disturbances, a node is assumed to live as long as it has some energy left.

9 Opportunities and Impact

Geometric approaches, through concepts and techniques, offer a number of opportunities in sensor networks to address problems at structural, functional and application levels. These opportunities are new in the sense that they have not been previously investigated in the context of internetworking. They are also expected to play an important role in the ongoing debate on whether the approaches and concepts that worked well in internetworking are also the right ones for sensor networks (albeit in a lightweight form).

Conversely, there are also tremendous opportunities for novel research in computational geometry. These opportunities go well beyond simply applying known geometric concepts to a new application domain. Indeed, the lightweight and distributed nature of the sensor platform, coupled with the constraints of energy, reliability, robustness, and security, necessitate a fresh approach to geometric algorithms. Moreover, as research on sensor networks discovers the true geometry of electromagnetic and other physical phenomena in diverse settings, we can expect new probabilistic, statistical, and model-based techniques to be devised for geometric solutions.

One possible approach to seed progress along these opportunities is through large-scale sensor network "community" testbeds. These testbeds would enable high-fidelity validation of geometric solutions as well as facilitate apples-to-apples comparison of geometry-inspired sensor network solutions with more traditional internet-style graph solutions. A large-scale testbed would also lend insight into how different algorithmic components in the sensor network (e.g., initialization, localization, structure discovery, routing, data storage, information brokerage, aggregation) interact with each other. Incidentally, early government and industrial research in sensor networks has already yielded a few large-scale testbeds, the GENI effort has recognized the importance of such testbeds, and efforts in the sensor network community to evolve and standardize these testbeds are presently underway. The participation by the computational geometry community can add valuable knowledge to this process and influence the development of a testbed that would be particularly appropriate for algorithms development.

A fundamental challenge in sensor networking has been that of scale, be it in terms of geographic coverage or network lifetime or network flexibility. Research at the interface of these two areas will help push the boundaries of scale, and perhaps even facilitate characterization of the limits of their communication capacity, information capacity, and observability.

References

[1] L. Barrière, P. Fraigniaud, and L. Narayanan. Robust Position-based Routing in Wireless Ad Hoc Networks with Unstable Transmission Ranges. In *Proc. 5th Intl. Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM)*, pages 19–27, 2001.

[2] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(2001), 609–616.

[3] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge UP, Mar. 2004.

[4] H. Breu and D. G. Kirkpatrick. “Unit Disk Graph Recognition is NP-hard,” *Comput. Geom. Theory Appl.*, 9(1-2):3–24, 1998.

[5] J. Bruck, J. Gao, and A. Jiang, “Localization and routing in sensor networks by local angle information,” in Proc. of the Sixth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc), 181-192, 2005.

[6] J. Bruck, J. Gao, and A. Jiang. MAP: Medial axis based geometric routing in sensor networks. In *Proc. ACM/IEEE Conference on Mobile Computing and Networking (MobiCom)*, 2005, 88–102,.

[7] Networked aquatic microbial observing system (namos). http://www-robotics.usc.edu/~namos/, 2005.

[8] B. N. Clark, C. J. Colbourn, and D. S. Johnson. Unit Disk Graphs. *Discrete Mathematics*, 86:165–177, 1990.

[9] N. A. Cressie. *Statistics for Spatial Data*. Wiley, 1991.

[10] 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 International Conference on Robotics and Automation (ICRA)*, 4285–4287, 2006.

[11] A. Efrat and S. Har-Peled, Locating guards in art galleries, *Information Processing Letters (IPL)*, to appear.

[12] Q. Fang, J. Gao, L. Guibas, V. de Silva, and L. Zhang. GLIDER: Gradient landmark-based distributed routing for sensor networks. In *Proc. of the 24th Conference of the IEEE Communication Society (INFOCOM)*, volume 1, pages 339–350, March 2005.

[13] R. Fonseca, S. Ratnasamy, J. Zhao, C.-T. Ee, D. Culler, S. Shenker, and I. Stoica. Beacon-vector: Scalable point-to-point routing in wireless sensor networks. In *Proceedings of the Second USENIX/ACM NSDI*, Boston, MA, May 2005.

[14] D. Ganesan, D. Estrin, and J. Heidemann. DIMENSIONS: Why do we need a new data handling architecture for sensor networks? In *Proceedings of the ACM HotNets*, pages 143–148, Princeton, NJ, USA, October 2002. ACM.

[15] R. Ghrist, D. Lipsky, S. Poduri, and G. Sukhatme, “Node isolation in coordinate-free networks,” *Proceedings of the Sixth International Workshop on Algorithmic Foundations of Robotics*, 2006.

[16] H. Gonzalez-Banos and J. Latombe. A randomized art-gallery algorithm for sensor placement. In *Proc. 17th ACM Symp. Comp. Geom.*, pages 232–240, 2001.

[17] B. Greenstein, D. Estrin, R. Govindan, S. Ratnasamy, and S. Shenker. DIFS: A Distributed Index for Features in Sensor Networks. In *Proceedings of First IEEE WSNA*, May 2003.

[18] C. Guestrin, A. Krause, and A. Singh. Near-optimal sensor placements in gaussian processes. In *ICML*, 2005.

[19] P. Gupta and P. Kumar. The Capacity of Wireless Networks. *Trans. Information Theory*, 46(2):388–404, 2000.

[20] D. S. Hochbaum and W. Maas. Approximation schemes for covering and packing problems in image processing and vlsi. *Journal of the ACM*, 32:130–136, 1985.

[21] C.-F. Huang and Y.-C. Tseng, “The coverage problem in a wireless sensor network,” in *ACM Intl. Workshop on Wireless Sensor Networks and Applications*, 115-121, 2003.

[22] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed Diffusion: a scalable and robust communication paradigm for sensor networks. In *Proc. of the 6th ACM/IEEE Conference on Mobile Computing and Networking MOBICOM*, 56–67, 2000.

[23] James san jacinto mountains reserve. http://www.jamesreserve.edu, 2005.

[24] B. Karp and H. T. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *Proc. 6th ACM/IEEE Conf. Mobile Computing and Networking MOBICOM*, 243–254, 2000.

[25] Y. J. Kim, R. Govindan, B. Karp, and S. Shenker. Geographic routing made practical. In *Proceedings of the Second USENIX/ACM NSDI*, Boston, MA, May 2005.

[26] H. Koskinen, “On the coverage of a random sensor network in a bounded domain,” in *Proceedings of 16th ITC Specialist Seminar*, pp. 11-18, 2004.

[27] A. Krause, C. Guestrin, A. Gupta, and J. Kleinberg. Near-optimal sensor placements: Maximizing information while minimizing communication cost. In *IPSN*, 2006.

[28] F. Kuhn, T. Nieberg, T. Moscibroda, and R. Wattenhofer. Local Approximation Schemes for Ad-hoc and Sensor Networks. In *Proc. Joint Workshop on Foundations of Mobile Computing (DIALM-POMC)*, pages 97–103, 2005.

[29] F. Kuhn, R. Wattenhofer, and A. Zollinger. Ad-hoc Networks Beyond Unit Disk Graphs. In *Proc. 1st ACM Joint Workshop on Foundations of Mobile Computing (DIALM-POMC)*, 2003.

[30] F. Kuhn, R. Wattenhofer, and A. Zollinger. Worst-Case Optimal and Average-case Efficient Geometric Ad-hoc Routing. In *Proc. 4th ACM Int. Symp. on Mobile Ad Hoc Networking & Computing (MobiHoc)*, pages 267–278, 2003.

[31] B. Leong, B. Liskov, and R. Morris. Geographic routing without planarization. In *Proceedings of the Third USENIX/ACM NSDI*, San Jose, CA, May 2006.

[32] X. Li, Y. J. Kim, R. Govindan, and W. Hong. Multi-dimensional range queries in sensor networks. In *Proceedings of the First SenSys*, pages 63–75. ACM Press, 2003.

[33] S. Madden, M. Franklin, J. Hellerstein, and W. Hong. TAG: a tiny aggregation service for ad hoc sensor networks. In *OSDI*, 2002.

[34] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. Srivastava, “Coverage problems in wireless ad-hoc sensor network,” in *IEEE INFOCOM*, 1380-1387, 2001.

[35] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *Proc. IEEE Infocom*, 2006.

[36] G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of the approximations for maximizing submodular set functions. *Mathematical Programming*, 14:265–294, 1978.

[37] J. Newsome and D. Song. Gem: Graph embedding for routing and data-centric storage in sensor networks without geographic information. In *Proceedings of the First SenSys*, 76–88, 2003.

[38] J. O’Rourke. *Art Gallery Theorems and Algorithms*. The International Series of Monographs on Computer Science. Oxford University Press, New York, NY, 1987.

[39] A. Ostfeld, J. G. Uber, and E. Salomons. Battle of the water sensor networks (bwsn): A design challenge for engineers and algorithms. http://www.eng.uc.edu/wdsa2006/BWSN_Rules_Final.pdf, 2006.

[40] V. Raghunathan, C. Schurgers, S. Park, and M. Srivastava. Energy Aware Wireless Microsensor Networks. *IEEE Signal Processing Magazine*, 19(2):40–50, 2002.

[41] A. Rao, S. Ratnasamy, C. Papadimitriou, S. Shenker, and I. Stoica. Geographic routing without location information. In *Proc. ACM/IEEE Conference on Mobile Computing and Networking MOBICOM*, 96–108, 2003.

[42] T. Rappaport. Wireless Communications: Principles and Practices. *Prentice Hall*, 1996.

[43] S. Ratnasamy, L. Yin, F. Yu, D. Estrin, R. Govindan, B. Karp, and S. Shenker. Ght: A Geographic Hash table for Data-centric Storage. Proc. 1st ACM workshop on Wireless sensor networks and applications, 2002, 78–87 Sept. 2002.

[44] S. Shenker, S. Ratnasamy, B. Karp, R. Govindan, and D. Estrin. Data-centric storage in sensornets. *SIGCOMM Comput. Commun. Rev.*, 33(1):137–142, 2003.

[45] V. de Silva and R. Ghrist, “Coordinate-free coverage in sensor networks with controlled boundaries,” submitted, 2006.

[46] V. de Silva, R. Ghrist, and A. Muhammad, “Blind swarms for coverage in 2-d,” in *Proc. Robotics, Systems and Science*, 2005.

[47] D. A. Sims and R. W. Pearcy. Sunfleck frequency and duration affects growth rate of the understorey plant, alocasia macrorrhiza. *JSTOR: Functional Ecology*, 7 (1993), 7683–7689.

[48] I. Stojmenović, A. Nayak, and J. Kuruvila. Design Guidelines for Routing Protcols in Ad Hoc and Sensor Networks with a Realistic Physical Layer. *IEEE Communications Magazine*, March 2005.

[49] F. Valladares, M. T. Allen, and R. W. Pearcy. Photosynthetic responses to dynamic light under field conditions in six tropical rainforest shrubs occuring along a light gradient. *Oecologia* 111 (1997), 505–514.

[50] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *Proceedings of the First SenSys*, 14–27, 2003.