

Geometric Algorithms for Sensor Networks

Jie Gao*

Leonidas Guibas[†]

June 7, 2011

1 Introduction

Networked embedded sensors provide a unique opportunity for real time, large scale, high resolution environmental monitoring. Such systems are becoming ubiquitous across many activities important to our economy and life, from manufacturing and industrial sensing, to traffic and powergrid management, to wildlife, agriculture and environmental monitoring, to hospital operations and patient observation, all the way to target tracking, battlefield awareness and other military applications.

The close relationship of sensor nodes with their embedded physical space imposes a unique geometric character to such systems. The physical locations of the sensor nodes greatly impact the system design in all aspects from low-level networking and organization to high-level information processing and applications. From a networking point of view, node placement clearly influences network connectivity and sensing coverage, which subsequently affects basic network organization such as clustering and localization, as well as naming and routing in the network. From an application point of view, sensor readings exhibit spatial correlations that can be exploited for data compression, approximation and validation. This article surveys algorithms which exploit geometric structures for sensor network operation and design. It is worth noting, however, that classical geometric techniques assume precise geometric information (e.g., node locations) which is often lacking in sensor deployments or operate in a heavily centralized fashion. Their adaptation to the sensor in sensor network context raises many novel challenges. In particular, the article covers the following set of research problems.

- *Network Localization.* Since wireless communications in sensor networks are short range, two nodes can be connected by a link only when they are in close proximity. The problem of network localization is to discover sensor node positions from the graph connectivity, possibly using additional distance or angle measurements.
- *Geometric Routing.* In resource-constrained and highly volatile networks, traditional routing table approaches cannot be applied. Geographical routing and its many variations, using node locations, provide alternative scalable solutions. A number of geometric ideas have been developed for routing using only local information and greedy decisions, even when true node geolocations are unavailable. We also cover recent research work for load balanced routing and routing using landmarks.

*Department of Computer Science, Stony Brook University, Stony Brook, NY 11794. Email: jgao@cs.sunysb.edu.

[†]Department of Computer Science, Stanford University, Stanford, CA 94305. Email: guibas@cs.stanford.edu.

- *Information Discovery and Data-Centric Routing.* Emerging applications of ubiquitous sensor deployments in spaces where people live and work envision a smart environment where situations are continuously assessed and responses generated in a real time manner. The sensors serve two purposes: discovering/detecting the events of interest; and forming a supporting infrastructure for distributed resources/users to be informed about and act on the detected events. In this setting, sensors that detect interesting data are named as information producers. Users in the same space, searching for such sensor data are called information consumers. Unfortunately, neither information producers nor consumers are directly aware of each other, unless some information brokerage mechanism brings them together. We survey various information processing, discovery and brokerage schemes, for different application requirements (query latency, storage, structures in the data, user query patterns), that best serve users' requests in a resource-constrained network. We also cover distributed data aggregation schemes.
- *Network Topology Discovery.* When a sensor network grows large in scale it is unrealistic to assume that nodes are placed uniformly in a regularly shaped region. Terrain variations and obstacles often forbid sensor placement; random sensor deployment also likely leads to deployment "holes" or "voids". The problem of network topology discovery is detecting and representing non-trivial sensor network topology — and then using this knowledge to improve routing and information discovery/delivery in the network.

We note that there are a number of topics which also heavily involve geometric techniques, such as topology control [103, 72] or geometry-based communication models, e.g., Signal-to-Interference-and-Noise-Ratio (SINR) [83, 51]. These are not covered in this survey due to space constraints. Interested readers are referred to the cited books or survey papers.

2 Network Localization

Sensor location information is indispensable for both sensor data integrity and important network management issues such as coverage and data delivery. Traditional approaches to obtain location information mainly rely on global positioning systems (GPS) [57]. But GPS is not appropriate for large-scale sensor localization due to its high cost, large form factor and environmental constraints. GPS requires direct line of sight to satellites and thus does not work indoor, underground, underwater and often has poor accuracy in large metropolitan areas (due to the "urban canyon" problem). There has been a lot of study on localization algorithms that derive the locations of sensor nodes from local measurements such as distance and angle estimations between neighbors [109, 110, 90, 91, 92, 113, 116, 7, 88, 50, 22]. The distance between two communicating nodes can be estimated by Received Signal Strength Indicator (RSSI) or Time of Arrival (ToA) techniques. Angles between adjacent edges can be measured by antenna arrays, multiple ultrasound receivers [96], or directional antennas.

Generally speaking, localization algorithms can be classified as anchor-based and anchor-free methods. Anchor-based methods assume that a (sometimes large) number of anchor nodes know their positions already [109, 110, 90, 91, 113, 92, 22]. In anchor-free methods, no node knows its absolute coordinates, the output is the relative positioning of the sensors, subject to a global translation and rotation.

Basic trilateration/triangulation. In the basic planar trilateration scheme, a node can determine its location

with distance measurements to three anchor nodes that are not co-linear [109, 110, 90]. For localization of a sensor network, we can repeat trilateration iteratively, called *iterative trilateration*. Nodes who obtain their locations become new anchor nodes and help to localize others. Similar methods can also be done by using angle measurements [91, 92]. For these incremental methods, two issues need to be addressed. One is to deal with cascading error accumulation in large networks. One can adopt optimization techniques such as mass-spring relaxation to smooth out error distribution, or use robust statistics to handle outliers in input measurements [78]. The other issue is to handle an insufficient number of initial anchor nodes. If the number of anchors is too small or the anchors are not well distributed, some nodes may not be able to find three neighboring anchor nodes to locate themselves. In this case, one can use distance estimations to anchor nodes via multi-hop paths [90] or adopt collaborative multi-lateration by solving a larger optimization problem [109, 110].

When the sensor nodes are not uniformly distributed and the network contains big holes, using hop counts as estimates of the Euclidean distances may introduce large errors as bending around hole boundaries may artificially inflate the Euclidean distance. To address this problem, paths that touch hole boundaries can be detected and then either eliminated from network localization [122] or amended to get a better distance approximation [76].

Graph rigidity and network localizability. Existing anchor-free algorithms take either the connectivity graph [113], or the distances between neighboring sensor nodes [116, 7, 88, 50] as input. One major challenge using this approach is localization ambiguity — when there are multiple, different localization solutions which satisfies all the distance constraints but are far from each other. Indeed, with range information local optimization such as mass-spring relaxation techniques may get stuck at local minima [24, 88].

The problem of whether a graph with given edge length constraints admits a unique embedding in the plane is studied in rigidity theory [53]. A graph is *rigid* in the plane if one cannot continuously deform the shape of the graph without altering the lengths of the edges. A graph is *globally rigid* if it admits a unique embedding in the plane, subject to global rotations and translations. The theory of graph rigidity in 2D has been relatively well understood. There is a combinatorial condition, the Laman condition, to characterize graphs that are generically rigid. There is also an efficient algorithm, the pebble game [59], to test whether a graph is generically rigid in time $O(nm)$ where n is the number of nodes and m is the number of edges. Similarly, both a combinatorial characterization of globally rigid graphs and polynomial algorithms for testing such graphs are known [56, 6]. It is however not trivial to apply these rigidity results in the development of efficient localization algorithms. Given a graph with edge lengths specified, finding a valid graph realization in \mathbb{R}^d for a fixed dimension d is an NP-complete problem [3, 4, 111].

The pioneer work of using rigidity theory in network localization [26, 48, 49, 88, 1, 116, 8] focuses on identifying special graphs that do admit efficient localization algorithms. The first idea is to use trilateration graphs [26, 48, 49, 88] — intuitively, the graphs derived by iterative trilateration methods. A trilateration graph is defined recursively. It is either a triangle or a trilateration graph with a trilateration extension, defined as adding an additional vertex with three edges to existing vertices. If the network contains a trilateration graph, one can exhaustively search for the ‘seed’ triangle in the graph and greedily find the trilateration extensions. Then an incremental algorithm, i.e., iterative trilateration, can be adopted to find the realization of the network. The second idea is to examine *uniquely d -localizable graphs*. A graph with known edge lengths is called uniquely d -localizable if there is a unique realization of the graph in \mathbb{R}^d and there is no non-trivial realization in \mathbb{R}^k with $k > d$. In other words, any realization in \mathbb{R}^k with $k > d$ stays in an affine subspace of dimension d in \mathbb{R}^k . For example, a generic simplex of $d + 1$ vertices is

uniquely d -localizable. For uniquely d -localizable graphs, So and Ye [116, 8] have shown that semi-definite programming can be used to find the realization. However, it is not known whether there is a combinatorial characterization of graphs that are d -localizable.

Tractability of network localization. From a theoretical point of view, one can formulate the localization problem as embedding a unit disk graph in the plane, if we assume that the communication graph follows the unit disk graph model. In a unit disk graph, two nodes are connected by an edge if and only if they are within distance 1. With purely the connectivity information, determining whether a combinatorial graph is a unit-disk graph is NP-complete, and thus finding such an embedding in the plane is also hard [14]. In fact, even a relaxed version of the problem is still hard. Kuhn *et al.* proved that finding an embedding such that non-neighboring pairs are at least 1 away and neighboring pairs are within $\sqrt{3}/2$ is NP-hard [69]. Distance measurements of neighboring nodes or angle measurements of neighboring edges alone do not help either [3, 4, 15]. Not much is known on approximation algorithms for unit disk graph embedding. So far the only positive result is an algorithm with an upper bound $O(\log^{2.5} n \sqrt{\log \log n})$ on the ratio of the longest distance between neighboring pairs to the shortest distance between non-neighboring pairs [69].

Global optimization methods. A number of localization methods use global optimization tools. These algorithms in general have the best accuracy but the downside is that they require the global network topology and thus are centralized methods, besides being computationally expensive.

In multi-dimensional scaling (MDS) [10, 113], the input is assumed to be a distance matrix on *all* pairs of nodes. By shifting to the center, the distance matrix is re-organized to derive the Gram matrix $G = XX^T$, where X is the coordinate vector on all sensor nodes. Since the Gram matrix G is symmetric and positive semi-definite, we can perform eigen-decomposition and derive $G = VAV^T$ where A is the diagonal matrix of all eigenvalues and V is the matrix of the corresponding eigenvectors. Now we can obtain the node coordinates X as $VA^{1/2}$. If we require a 2D or 3D embedding, we can take only the largest two or three eigenvalues. In practice when it is difficult to get all pairs' distance measurements, it is common practice to use the minimum hop count to approximate the Euclidean distance [113]. In a network with uniform sensor density and a regular shape, MDS typically works very well. But its performance may deteriorate significantly when the network shape becomes complex [73].

Biswas and Ye [7] used semi-definite programming (SDP) for network localization. The input is a graph with possible anchor nodes and distance measurements on the edges. The basic idea is to introduce the Gram matrix variables Y . Ideally $Y = XX^T$, where X is the coordinate vector on all sensor nodes. This constraint is relaxed such that $Y - XX^T$ should be a positive semi-definite matrix. Together with other distance constraints it can be shown that the system is a semi-definite program and can be solved using standard interior point methods. Similarly, with angle measurements one can formulate a linear program that solves for the node locations [15].

Cut, embed, and paste. A family of practical algorithms uses a bottom up approach by first localizing small components and assembling them together. Moore *et al.* [88] proposed to use robust quadrilaterals as the atomic components. A quadrilateral is a complete graph on four nodes and the smallest globally rigid graph. When the distance measurements are noisy, the skinny quadrilaterals, when three of the four nodes are almost co-linear, can possibly have flip ambiguities. Thus only a set of robust quadrilaterals, i.e., those that are fat, are used. The global layout is obtained by gluing locally identified robust quadrilaterals. Similarly, with ideas from rigidity theory to improve iterative multi-lateration on sparse networks, Goldenberg *et al.* proposed to record, propagate and verify multiple possible locations of sensors to discover the truth network

layout [49]. The idea of decomposing the network into small pieces, embedding them and then gluing them together can also be applied in a multi-resolution manner [62].

In [73], Lederer *et al.* considered network localization using only connectivity information when the network shape is complex. The algorithm first selects landmarks on network boundaries with sufficient density, then constructs the landmark Voronoi diagram (in which all nodes closest to the same landmark are grouped in the same Voronoi cell¹) and its dual combinatorial Delaunay complex on these landmarks (in which an edge is constructed between two landmarks if they share common nodes). The key insight is that the combinatorial Delaunay complex is provably *globally rigid* and has a *unique* realization in the plane, provided that the landmarks are sampled according to the local geometric complexity of the sensor field. Thus an embedding of the landmarks by simply gluing the Delaunay triangles properly recovers the faithful network layout. With the landmarks nicely localized, the rest of the nodes can easily localize themselves by using trilateration to nearby landmark nodes. This leads to a practical and accurate localization algorithm for large networks using only network connectivity.

3 Geometric Routing

3.1 Geographical routing

In geographical routing, the physical locations of the sensor nodes are used to guide the path a packet takes in the network. Geographical routing was originally proposed by Bose *et al.* [12], and independently by Karp and Kung [61]. Geographical routing has two modes: the greedy mode and the recovery mode.

Greedy mode. In this mode, the node currently holding the packet “advances” it towards the destination, based only on the location of itself, its immediate neighbors and the destination. The advance rule may be defined in many ways. Examples are, Closest to Destination [121, 32, 61, 12], Most Forward within Radius (MFR) [121], Nearest Forward Progress (NFP), Nearest Closer (NC) [120], Geographic Distance Routing (GEDIR) [119] and Compass routing [66]. The most popular way of defining the advance is by decreasing the Euclidean distance to the destination. Greedy routing often suffices to deliver a packet in a dense network, but may fail in a sparse network, where the packet may reach a *local minimum* whose neighbors are all further away from the destination than itself.

Recovery mode. The recovery mode defines how to forward the packet at a local minimum. Some examples of the recovery methods are simple flooding [119], terminode routing [9], breadth-first search or depth-first search [60], the face algorithm [12] and perimeter routing [61].

Here we use a specific routing protocol, the Greedy Perimeter Stateless Routing (GPSR) Protocol [61], to explain the two modes in detail. A nice summary of geographical routing schemes with a focus on some subtle issues is presented in [35]. In GPSR, the greedy mode uses minimum distance to destination as the rule of making progresses. To get out of the local minimum, the protocol maintains a planar and connected subgraph, e.g. the *Gabriel Graph* (GG) or the *Relative Neighborhood Graph* (RNG), and applies routing along the faces of the subgraph that intersect the imaginary line segment between the source and the destination. Greedy routing combined with the perimeter routing guarantees the delivery of a packet to the destination if there is indeed a way of doing so. In addition, the construction of the planar graphs,

¹One node may have multiple closest landmarks and thus may stay in multiple Voronoi cells.

as well as the routing decisions, are localized — requiring only information from neighboring nodes. This makes geographical routing very attractive for large-scale networks as the state information maintained in each node and in the packet is nearly minimum.

Theoretical bounds on path stretch. The path stretch measures how much longer the routing path is compared with the shortest path. A lower bound proven in [71] states that any deterministic or randomized algorithm using only local information cannot find a path of length shorter than $\Omega(k^2)$ in the worst case, where k is the length of the shortest path. On the other hand, it is shown in [45] that greedy routing, in the case of successful delivery, uses $O(k^2)$ hops. The GOAFR+ family protocols [71, 70], or, alternatively, using restricted flooding with doubling radii, both guarantee a worst case path length of $O(k^2)$.

Geographical routing in 3D. Durocher *et al.* [23] showed a negative result that it is impossible to find a constant state routing algorithm in 3D unit ball networks. Thus one cannot hope to directly carry the ideas to routing in 3D networks.

Geographical routing in practice. The recovery scheme with face routing or perimeter routing works nicely in theory but encounters a few problems in practice. The correct construction of a planar subgraph assumes accurate location information, which is hard to obtain, and that the communication graph is modeled by a unit disk graph, which does not hold in practice. Experimental evaluations of the communication model on sensor nodes reveal various spatial and temporal radio irregularities [41, 132]. In practice, the planar subgraph extracted may become disconnected or still contain crossing edges, as shown in [112, 64]. This subsequently causes the delivery rate to drop to about 68% on a real testbed [64]. Later, a number of repair mechanisms propose to remove crossing links by probing [64, 63, 52]. These mechanisms improve the delivery rate at the cost of higher processing overhead.

Another problem of face routing, especially in a sensor field with big holes, is that the routes are likely to be ‘hugging’ the hole boundaries. Thus the nodes on the hole boundaries deliver higher traffic than the average. The load imbalance may cause the boundary nodes run out of battery earlier than the others, enlarging the size of the holes and disconnecting the network.

These two issues motivated researchers to go beyond geographical routing and design virtual coordinates that still retain the good properties of geographical routing but are more robust to radio models, topological changes and node failures.

3.2 Greedy routing with virtual coordinates

NoGeo: using rubberband representations. NoGeo [99] is one of the earliest proposals for efficient geometric routing with virtual coordinates. The virtual coordinates are defined by the rubberband representation of a graph [81]. In short, suppose some nodes (preferably on the network perimeter) are given fixed locations; the rubberband representation is obtained by each node (that is not fixed) running an iterative algorithm of putting itself at the center of mass of the neighbors’ current locations (also known as Laplacian smoothing). This algorithm is known to reduce the total energy of the system, if all the edges are considered as rubberbands whose rest length is zero, and converge to the *rubberband representation*. NoGeo has no guarantee on delivery by greedy routing on the virtual coordinates but works well in simulations. To understand that, one can see that any node (other than those with fixed locations) cannot be a reflex vertex, i.e., it cannot have all neighbors on the same half plane. Since a reflex vertex is bad for greedy routing, this observation intuitively explains the good performance of routing on rubberband representations in practice.

Greedy embeddings. Motivated by NoGeo, researchers asked whether one can find an embedding of a given graph in the plane such that greedy routing always works [94]. Such an embedding is called a *greedy embedding* or a *greedy drawing*. It is known that not every graph admits a greedy embedding, such as a star with 7 leaves [94]. Some graphs are known to have a greedy embedding, for example, graphs that contain a Hamiltonian path, complete graphs, 4-connected planar graphs (as they are Hamiltonian [123]), power diagrams [17] and Delaunay triangulations [11, 79]. But it still remains open to fully characterize the class of graphs that admit a greedy embedding.

Papadimitriou and Ratajczak [94] showed that any planar 3-connected graph (A graph is 3-connected if it remains connected after the removal of any 2 nodes) has an embedding in 3D such that with a special distance function greedy routing always works. This is due to a famous graph theory result that a 3-connected planar graph is actually the edge graph of a 3D convex polytope [133]. They also conjectured that a planar 3-connected graph has a greedy embedding in the plane. Dhandapani [21] discovered that for any planar triangulation there *exists* a greedy embedding in the plane using Schnyder’s realizer concept. The conjecture was proved to be true in 2008 by Leighton and Moitra [74], and independently by Angelini *et al.* [2].

A tricky and important issue, as pointed out by Eppstein and Goodrich [25], is that the greedy embedding as proposed earlier (e.g., in [74]) may require $\Omega(n \log n)$ -bit size coordinates. This makes such greedy drawing methods impractical as the space requirement is the same as standard routing table approach. Ideally, we would like *succinct greedy drawing* in the plane, in which the virtual coordinates are represented by $O(\log n)$ bits. Angelini *et al.* [5] proved that this is not possible. They show that there are infinitely many trees whose greedy embeddings need exponential size grids. Recently, Goodrich and Strash used the embedding presented in [74] and developed succinct $O(\log n)$ -bit coordinates (not Euclidean though) that support greedy routing using a different metric function. In a similar way, He and Zhang showed that the classical Schnyder drawing admits greedy routing with succinct coordinates under a special metric function [55].

Sarkar *et al.* [105] considered a triangulated domain with holes and produced an embedding such that all the holes are mapped to circular disks, by using conformal mappings and discrete Ricci flows. Thus greedy routing never gets stuck at hole boundaries. One other good property is that the mapping can be computed with a distributed gossip-style algorithm.

Most of the embeddings are in low dimensional (2D or 3D) Euclidean spaces. Flury *et al.* [33] considered embedding a graph in an $O(\log n)$ -dimensional space and the routing stretch is bounded by $1 + \varepsilon$ for any ε .

Embedding in hyperbolic spaces. Robert Kleinberg [65] showed that greedy routing in hyperbolic space is easy, since any *tree* has a greedy embedding in the hyperbolic plane. Unfortunately, the hyperbolic embedding in [65] does not provide succinct coordinates either. In [86] Maymoukov provides a greedy drawing method for three-dimensional hyperbolic space using vertices that can be represented with $O(\log^2 n)$ bits. Eppstein and Goodrich [25] proposed a method to embed a tree in a ‘balanced’ manner and therefore showed that any n -vertex connected graph can be embedded in two-dimensional hyperbolic space with succinct coordinates that support greedy routing.

3.3 Load balanced routing

In large scale sensor networks it is critical to balance work load on different sensors, to prevent some nodes from running out of battery immaturely. For load balanced routing, one formulation of the problem is to

select routes such that the maximum load is minimized. Unfortunately solving this problem is NP-hard (modeled as unsplittable flow problem) even in very simple networks (such as a grid). Approximation algorithms [97, 98] have been developed for the problem. But they all require global coordination, and are thus inappropriate for low resource sensor networks. In this section we focus on the load balancing issue of *greedy routing* solutions.

Routing in a narrow band. In [46], Gao and Zhang proposed a greedy routing scheme that achieves a routing stretch factor of four (compared with shortest path routing) and a load balancing ratio of three (compared with the optimal load balanced routing), for wireless nodes distributed in a narrow strip of width $\sqrt{3}/2$, when the communication graph follows the unit disk graph model. Even for this restricted scenario, it is still NP-hard to compute the most balanced routes.

Routing in a disk. In a simple network with sensors uniformly spread in a disk, shortest path routing (or geographical greedy routing) will incur higher traffic load on the nodes near the center. Popa *et al.* [95] proposed the curveball method to use stereographic projection to map the network to a hemisphere. Routing is done by greedy routing using the spherical distance instead of the Euclidean distance. Intuitively the paths now move away from the center. The same idea is used by Li *et al.* [75] to resolve routing congestion at network center.

In a recent work [130], the load balancing property is studied by using the projection of a 3-connected planar graph to a convex 3D polyhedron as described in the Koebe-Andreiev-Thurston's embedding [133]. When the network is dense and with a proper embedding, the polyhedron is approximately a sphere. Thus greedy routing on the polyhedron also avoids the problem of causing a dense center. Different from the curveball method, greedy routing on the convex polytope guarantees delivery [94].

Routing in a square. Mei and Stefa [87] studied load balancing for a square shape sensor network. The network is quadrupled by successive reflections in the X and Y axes and the network is now four times as big and has the topology of a torus. Each node is mapped to an image in each of the 4 copies. For routing, each of the four copies of the destination is chosen with equal probability. And routing is executed using greedy geographical routing on the torus, which equals to reflecting back from the boundary in the original network. This method has good load balancing properties and the center is not overloaded, but at the cost of increasing the average traffic load, as the routing path may now be much longer.

Routing in the covering space. For a general network with possibly big holes, none of the previous methods work very well. In particular, most greedy routing methods, using geographical locations or other virtual coordinates, tend to follow the hole boundaries and heavily overload these nodes. Sarkar *et al.* [106] proposed to use the notion of covering space to deal with the problem. This builds on top of the virtual embedding [105] in which all boundaries are circular. Now we have a network with k circular holes. For each interior hole, one can take a Möbius transformation that essentially 'reflects' the network inward with respect to the hole. This partially fills up the hole, except that there are k smaller circular holes. Now we can continue such transformations with respect to the remaining holes so that all the holes are eventually filled up, with infinitely many transformed copies of the original sensor field. In practice we stop after a number of rounds when the holes left are sufficiently small. For routing in the covering space a message can 'enter' the hole to route in another copy of the network, effectively reflecting on the hole boundary. The nodes on the boundary are now treated in the same way as the other nodes with respect to routing and the problem of them being overloaded is substantially alleviated.

3.4 Landmark-based routing

A family of routing algorithms use landmark schemes [28, 34, 16, 89]. The idea of using landmarks for generating naming and routing schemes in large networks was initially proposed two decades ago, such as the landmark hierarchy by Tsuchiya [126]. We will discuss hierarchical routing in a later section. The landmark-based routing schemes described below are all two-level hierarchy.

In sensor network setting, a small subset of the nodes are selected as *landmarks*. The landmarks flood the network and every node records its hop count distances to these landmarks. The landmark distances are then used to generate virtual coordinates for the nodes. Routing is typically guided in a greedy way by a potential function on the landmark-based distances. Landmark-based schemes are favored by their simplicity and independence on the dimensionality of the network.

Gradient Landmark-based Routing (GLIDER). GLIDER [28] is concerned about routing in a sensor network with holes or an irregular shape. A number of landmarks are selected and the network is partitioned into Voronoi tiles such that nodes inside the same tile are closest to the same landmark. Two Voronoi tiles are adjacent if there are two neighboring nodes in different Voronoi tiles. The tile adjacency graph, denoted by the *combinatorial Delaunay graph*, is abstracted and known to all the nodes, for global routing across tiles. The landmarks are selected to capture the global topology of the sensor field. Thus the number of landmarks is likely to be dependent on the topological complexity, such as the number of holes, typically much fewer than the number of sensors [93].

Each node p is given virtual coordinates, defined by the distance to its closest landmark u and the distance to u 's neighboring landmarks in the Voronoi diagram. To route to the destination node, the source node will first consult with the combinatorial Delaunay graph to find a sequence of tiles to visit. The actual route will be composed of *inter-tile routing* (routing to an adjacent tile) and *intra-tile routing* (routing within a tile), both using greedy, local routing schemes based on landmark distances.

The same combinatorial Delaunay structure is used by Funke and Milosavljević to account for inaccurate location information and small network holes and irregularities [39, 40].

Beacon Vector Routing (BVR). Beacon Vector Routing (BVR) [34] uses a potential function that depends on the distances to the k landmarks closest to the destination, where k is a system parameter. Out of the k landmarks, the ones that are closer to the destination than to the source impose a “pulling” force while the rest impose a “pushing” force. The potential function is a combination of the two. BVR has received significant attention and has been selected in a number of cases as a comparison benchmark [89, 85].

Gradient landmark descent routing. None of the previous methods guarantee routing stretch. Nguyen *et al.* [89] proposed a method with bounded stretch in the continuous domain. The landmarks are selected to be an r -sampling: any node is within hop count r from a landmark and no two landmarks are within distance r of each other. r -sampling has better uniformity than random sampling and can also improve the performance of other methods such as BVR. Routing is repeated by greedy forwarding towards the landmark that *maximizes* the ratio of the distances to source and destination.

Small State Small Stretch Routing (S4). Mao *et al.* [85] adopted the idea of compact routing schemes (by Thorup and Zwick [124, 125]). When the destination is far away, it is a reasonable approach to route towards the landmark closest to the destination. Only when a message gets close to the destination, when all landmarks are equally far away, do we need a separate scheme to deliver the message to the destination. In

S4, a local routing table is constructed for the last phase. The routing scheme guarantees worst case stretch of 3 and routing table size of roughly $\tilde{O}(\sqrt{n})$ for a uniform regular network of n nodes.

To conclude, we remark that a few other protocols, whose ideas are similar to one of these, were not covered in details here [16, 128]. These landmarks in these schemes are mainly used as points of reference, and do not serve the functionalities of gateway nodes (which attract traffic). The schemes described here differ in how routing to a distant or a nearby destination is designed and integrated.

3.5 Information Discovery and Data-Centric Routing

In data-centric routing, a node poses a query for data of certain types, and a routing algorithm must route the query to retrieve the relevant data in the sensor network. Data-centric routing is fundamentally different from traditional routing paradigms as the problem also involves the discovery of the destination, the nodes who have the required data. A dual version is also possible, in which data is routed looking for relevant queries. Both versions can be formulated as the problem of *information brokerage* which, matches and carries information between *information producers* (also known as sources), nodes that perform data acquisition and event detection, with *information consumers* (also known as sinks), nodes that search for or need this information.

This problem in sensor networks was initiated in the directed diffusion work of Intanagonwiwat, Govindan and Estrin [58]. The sensor nodes locally process their data and organize them by attribute-value pairs. A query (called an interest) is disseminated to the entire network and matched data will be routed back, possibly along multiple paths. Reinforcement learning is used to prune the bad paths. A similar routing paradigm is also adopted in TinyDB [84] to support SQL-style aggregation queries on information distributed in the network. These two approaches aim at infrequent queries for streaming data type so that the cost of flooding can be justified and amortized over the following long-term data delivery. In the following we describe a number of schemes that exploit the geometric properties of the network for more efficient information discovery and brokerage.

Geographical hash tables. A parallel approach to data-centric routing is to adopt data-centric storage, which aims at large-scale networks with many simultaneously detected events that are not necessarily desirable for all users [114, 101]. The idea is similar to Distributed Hash Tables (DHTs) on the Internet [102, 131, 117, 100]. A producer leaves data on rendezvous nodes for consumers to retrieve. Thus data across space and time can be aggregated at rendezvous nodes. In geographical hash tables (GHTs) [101], data is hashed by its data type to geographical locations. The nodes close to the hashed location serve as rendezvous nodes. Data and query delivery to the rendezvous nodes is implemented by geographical routing such as GPSR [61].

GHTs have greatly reduced the communication cost and energy consumption by avoiding network-wide flooding for information discovery. Their simplicity is also very appealing. However, the basic GHT scheme does not have *distance sensitivity*. In particular, if the producer is actually close to the data producer (although neither has the knowledge of each other), we would like the data consumer to discover the data producer quickly. This is an attractive property in many applications, as information will be most useful, thus queried more frequently, in the spatiotemporal locale where it was collected.

In [37] a discrete center hierarchy [43] is built for distance sensitive retrieval. The hierarchy is composed of clusters with a parent cluster including all children clusters. A data item is hashed and stored at certain

other nodes, called the *data servers*, chosen from nearby clusters and fewer servers in far away but large clusters. The retrieval scheme starts to search in local neighborhood until it discovers a nearby data server. The total cost is proportional to the distance to the producer.

Double rulings. Since data delivery from data source to a rendezvous node is implemented by multi-hop routing, it is natural to leave information hints along the trail that the data travels on, at no extra communication cost. In a basic *double-ruling* scheme works as follows: data or data pointers are stored at nodes that follow a *replication curve* while a data request travels along a *retrieval curve*. The crucial property is that any retrieval curve intersects the replication curve for the desired data. Thus successful retrieval can be guaranteed.

For an easy case, assume the network is a two-dimensional grid. The information storage curves follow the horizontal lines. The information retrieval curves follow the vertical lines. We name this scheme the *rectilinear double rulings* [82, 129, 118]. The spherical double rulings scheme [107] has both GHTs and the rectilinear double rulings as special cases. It considers the sensor field as the stereographic projection of a sphere and the producer curve is a circle through the producer and the hashed location, which in a stereographic projection maps to the great circle of a sphere. Both double-ruling schemes are distance sensitive — if the producer and consumer are of distance d , the producer can find the data with cost $O(d)$. The spherical double ruling also has many good properties as it allows natural data aggregates, supports range queries, and is robust to node failures.

Double rulings schemes can also be defined in the virtual coordinates space [29] or in a probabilistic manner (as in rumor routing [13]).

3.6 Data collection and aggregation

Many applications of sensor networks require that data be collected from all nodes of the network or a subset thereof. Most common approaches for data collection in a network are tree-based. In this subsection we introduce a few geometric ideas that take a flat, non-tree approach and use sweeps or sensor initiated probes for data collection and aggregations.

Sweeps over sensor networks. The network sweep, proposed in [115], is implemented by a narrow band of active nodes that ‘moves’ over the network by issuing invitations to new nodes to join the band, and dropping nodes that have already been processed and serve no other essential purpose.

By using simple local diffusion algorithms, one can pre-compute a certain *potential* at each node, essentially obtained by solving a simple PDE, a Laplace’s equation with Dirichlet boundary conditions, over the network. The algorithm, in its simplest form, is just the familiar Gauss-Seidel iteration (each node sets its value to the average values of its neighbors). The construction guarantees that no local extrema where the sweep could get stuck are present in the field. In general, such potentials and their gradients are rather resilient to link connectivity changes in the network because their values integrate information from large areas of the network. They also smooth out local variations in node density. They are used by the sweep protocol to provide global guidance so that the wavefront propagates in an orderly way without self-collision while allowing the sweep control to remain completely local.

Such diffusion-based gradients have also been used for information routing and discovery [80]. In particular, the linear space algebraic structure of these gradients allows certain elegant optimizations on load balancing, boolean query routing, and more sophisticated aggregations.

Sparse aggregations. Most of the extant work assumes that the goal of the aggregation is to accumulate data from *all* the network nodes, while in certain situations only a relatively small subset of the nodes have useful data to report, or want to participate in an aggregation. Clearly, for rare events, frequent polling is inefficient as the query needs to reach all nodes although most of them will not participate. A purely event-driven approach, where each node with a detection independently transmits a record of the event to a base station, may also not be optimal — consider a group of nearby source nodes sending a message each to the base station.

In [44], the problem is named *sparse aggregation* and the nodes with data are named *hot nodes*, that are unaware of each other. Key to the solution is a probe protocol that can be initiated independently by a node and involves emitting a few packets out of the node that follow certain paths in the network while leaving a trail behind, until the probes are terminated. These paths are chosen as the double rulings paths (say horizontal and vertical rays) such that they support distance sensitive discovery. At the probes encounter each other the data from the nodes are aggregated. The two nodes can be considered as competing in a tournament and the winner of the tournament is decided based on the data carried or left behind by the probe packets. The losing node then has his probes recalled/terminated, while the probes of the winner continue to propagate. Eventually the champion’s probe, together with the final aggregated data, propagates to the network boundary, where a high-speed network is assumed to be available. The probes also implicitly build an aggregation tree on the hot nodes, in a bottom up fashion. In the end it can be shown that the weight of the aggregation tree, as well as the total message cost is $O(T \log n)$, where T is the weight of the minimum spanning tree of the hot nodes and n is the network size.

3.7 Range queries

For a typical range query, we are given a query region plus possibly a range of the sensor data, and then ask for all the sensors in the query region whether any sensor data is within the data range. This is a problem that has been studied a lot in computational geometry. Various distributed schemes have been proposed. In the case of a scalar field, one solution is to partition the information about large geographic regions into subsets according to smaller ranges of the field value, and store these subsets in different nodes. This is the approach taken in the DIFS system [54]. In the DIM system [77] a locality preserving hash function is used to map portions of a multidimensional attribute space to sensors so that all data needed to answer a range searching query can be located conveniently. In the fractional cascading approach [42], information is stored so that more detailed information is available about data obtained in the spatio-temporal locality of the sensor where the query is injected — but without sacrificing the ability to query distant regions or times as well.

All of these schemes are designed to support range queries for static sensor data and essentially use a quadtree-type hierarchical space decomposition. For mobile data, constant updates to a fixed space partitioning make these schemes too costly — small movement of a target may lead to updates up to a high level of the quadtree and possibly updates on *all* sensors, if the mobile target happens to cross a high-level boundary.

For range queries of mobile targets, a naive solution is to do a geocast to the query region, flood the nodes in the query range and collect the target counts. Only the sensors that detect the target keep information about the target. This has an update cost proportional to the target movement distance, but a query cost proportional to the area of the range. Sarkar [104] recently proposed to store target movements implicitly

using the differential one-form such that the range queries have cost proportional to the perimeter of range R . Consider a planar graph embedded in the plane, and one target lies within a face f_0 and has a weight of w . The differential one-form is represented by a function ξ on directed edges. That is, the value for $\xi(ab)$ is the negation of the value for $\xi(ba)$. The differential one-form maintains the property that for the face f_0 , the summation of all the values of the edges on its boundary, in clockwise order, is w . It is easy to see that for a query range that contains multiple faces, and if one sums up the edge weights in clockwise order for all the faces in the range, the edges that are completely within the range will be counted twice in opposite directions – so their weights cancel out. Thus a range query can simply ignore the edges that are in the interior of R and take the clockwise integration along the range boundary only. The differential one-form provides great flexibility that allows low maintenance cost under both network dynamics and target movements. Network dynamics such as link addition and removal, or node insertion and removal, can be handled in constant time as well.

4 Network Topology Discovery

Large scale sensor networks are unlikely to be uniform and regular due to obstacles and terrain variations. A sensor network is regarded as a discrete sampling of an underlying geometric environment or space. Thus, the holes of the sensor field can be intrinsic features indicating important structures of the underlying environment such as inaccessible terrain, or obstacles (buildings, lakes, etc), as may be indicated on a building floor plan, a map of a transportation network, etc. Holes may also map to events that are being monitored by the sensor network. If we consider the sensors with readings above a threshold to be “invalid”, then the hole boundaries are basically isocontours of the landscape of the attribute of interest. Holes are also important indicators of the general health of a sensor network, indicating e.g. traffic hot spots or sensors with low battery levels.

To discover the nodes on the ‘hole boundaries’ of the sensor field, various approaches with different assumptions have been investigated, including geometric methods that use the node locations [27], statistical methods [31, 30], graph methods [67, 36, 38] and topological methods [47, 20, 19, 127]. In this section we describe some representative ideas to extract and represent the network hole structure.

Geometric and statistical methods. Geometric methods for boundary detection use geographical location information. The first paper on this topic, by Fang *et al.* [27], assumes that the nodes know their geographical locations and that the communication graph follows the unit-disk graph assumption. The definition of holes in [27] is intimately associated with geographical forwarding such that a packet can only get stuck at a node on hole boundaries. Fang *et al.* also proposed a simple algorithm that greedily sweeps along hole boundaries and eventually discovers boundary cycles.

Statistical methods for boundary detection make assumptions about the probability distribution of the sensor deployment. One can detect boundary nodes as those with relatively low degree [31] or with relatively low “restricted stress centrality” (the number of shortest paths going through) [30], in a network with uniform random sensor distributions. The major limitation of these two algorithms is the requirement on sensor distribution and typically very high node density.

Graph methods. Funke [36, 38] developed a simple algorithm using only network connectivity information. The basic idea is to construct isocontours based on hop count from a root node and identify where the contours are broken. Kröller *et al.* [68] presented an algorithm searching for combinatorial structures called

flowers and augmented cycles, in a quasi-unit disk graph. A flower would ensure the node in the center of the flower is not on the boundary. Saukh *et al.* [108] considered the embedding together with the graph structure and defined a node to be on the boundary if there is a quasi-unit disk graph embedding in which the node stays on the *geometric* boundary. Then they considered subgraph *patterns* that ensure a certain subset of the nodes to be always in the interior of any quasi-UDG embedding. Only the nodes that are certified to be in the interior will not be on the boundary. The patterns can be considered as generalizations of the flower structure in [68].

Topological methods. Ghrist *et al.* in [47, 20, 19] took an algebraic topology approach to study sensor coverage problem: check whether the domain is entirely covered by the union of the sensing ranges of all the sensors, assuming each node has a disk sensing range. They abstracted the coverage problem by forming the Čech complex $C(r)$, in which the k -simplices correspond to nonempty intersections of the sensing ranges of $k+1$ distinct sensors, with r as the radius of the sensing range. Now, by the nerve theorem, when the sensing ranges and all their non-empty finite intersections are contractible, the union of the sensing ranges has the homotopy type of the covered domain. Thus the detection of sensing holes is directly solved by checking whether the domain boundary cycle is null-homologous in the Čech complex. However, the computation of a Čech complex requires the location of the sensor nodes and expensive computations for the common intersection of disks. Instead, they showed that one can use two Vietoris-Rips complexes to sandwich the Čech complex. A Rips complex $\mathcal{R}(r)$ contains a k -simplex if the $k+1$ vertices are pairwise within distance r . In particular, the following inclusion relationship holds: $\mathcal{R}(r') \subset C(r) \subset \mathcal{R}(r)$, for $r/r' \geq 2/\sqrt{3}$. Therefore the homology of the Rips complex gives sufficient but not necessary conditions for the coverage holes in the Čech complex — if there is a nontrivial cycle in $\mathcal{R}(r)$, then there must be a cycle in $C(r)$; if there are no nontrivial cycles in $\mathcal{R}(r')$ then $\mathcal{R}(r)$ provides full coverage.

In a different approach [127], the structure of the homotopy classes of the shortest paths from a single node is used to detect holes. Holes in a sensor field create irregularities in hop count distances. This phenomenon gives rise in the continuous setting to the *cut locus* [18], defined as the collection of points with distinct shortest paths of different homotopy type — paths that cannot be “continuously deformed” into one other but get around the hole(s) in different ways and terminate by touching each other, trapping the hole(s) between them. The nodes in a cut can be easily identified, since they have the property that their common ancestor in the shortest path tree is fairly far away, at the other side of the hole. One nice property of the cut detection is that the operations can be done independently and locally at each pair of adjacent nodes, allowing natural parallelism.

References

- [1] B. D. O. Anderson, P. N. Belhumeur, T. Eren, D. K. Goldenberg, A. S. Morse, W. Whiteley, and Y. R. Yang. Graphical properties of easily localizable sensor networks. *Wireless Networks*, 15(2):177–191, 2007.
- [2] P. Angelini, F. Frati, and L. Grilli. An algorithm to construct greedy drawings of triangulations. In *Proc. of the 16th International Symposium on Graph Drawing*, pages 26–37, 2008.
- [3] J. Aspnes, D. Goldenberg, and Y. R. Yang. On the computational complexity of sensor network localization. In *The First International Workshop on Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSORS)*, pages 32–44, 2004.

- [4] M. Badoiu, E. D. Demaine, M. T. Hajiaghayi, and P. Indyk. Low-dimensional embedding with extra information. In *SCG '04: Proceedings of the twentieth annual symposium on Computational geometry*, pages 320–329, 2004.
- [5] G. D. Battista, F. Frati, and P. Angelini. Succinct greedy drawings do not always exist. In *Proc. of the 17th International Symposium on Graph Drawing*, volume 5849 2009, pages 171–182. Springer, 2009.
- [6] A. R. Berg and T. Jordán. A proof of connelly’s conjecture on 3-connected circuits of the rigidity matroid. *J. Comb. Theory B*, 88(1):77–97, May 2003.
- [7] P. Biswas and Y. Ye. Semidefinite programming for ad hoc wireless sensor network localization. In *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks*, pages 46–54, 2004.
- [8] P. Biswas and Y. Ye. Semidefinite programming for ad hoc wireless sensor network localization. In *Proceedings of the third international symposium on Information processing in sensor networks*, pages 46–54, 2004.
- [9] L. Blazević, L. Buttyán, S. Capkun, S. Giordano, H. Hubaux, and J. L. Boudec. Self-organization in mobile ad hoc networks: the approach of terminodes. *IEEE Communications Magazine*, pages 166–175, 2001.
- [10] I. Borg and P. Groenen. *Modern Multidimensional Scaling: Theory and Applications*. Springer-Verlag, 1997.
- [11] P. Bose and P. Morin. Online routing in triangulations. In *Proceedings of the 10th International Symposium on Algorithms and Computation, ISAAC '99*, pages 113–122, London, UK, 1999. Springer-Verlag.
- [12] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(6):609–616, 2001.
- [13] D. Braginsky and D. Estrin. Rumor routing algorithm for sensor networks. In *Proc. of the 1st ACM Int’l Workshop on Wireless Sensor Networks and Applications (WSNA)*, pages 22–31, September 2002.
- [14] H. Breu and D. G. Kirkpatrick. Unit disk graph recognition is NP-hard. *Computational Geometry. Theory and Applications*, 9(1-2):3–24, 1998.
- [15] 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’05)*, pages 181–192, May 2005.
- [16] A. Caruso, A. Urpi, S. Chessa, and S. De. GPS free coordinate assignment and routing in wireless sensor networks. In *Proc. of the 24th Conference of the IEEE Communication Society (INFOCOM)*, volume 1, pages 150–160, March 2005.
- [17] M. B. Chen, C. Gotsman, and C. Wormser. Distributed computation of virtual coordinates. In *SCG '07: Proceedings of the twenty-third annual symposium on Computational geometry*, pages 210–219, 2007.
- [18] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, 1997.
- [19] V. de Silva and R. Ghrist. Coverage in sensor networks via persistent homology. *Algebraic and Geometric Topology*, 7:339–358, 2007.
- [20] V. de Silva and R. Ghrist. Homological sensor networks. *Notices American Mathematical Society*, 54(1):10–17, 2007.
- [21] R. Dhandapani. Greedy drawings of triangulations. In *SODA '08: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 102–111, 2008.
- [22] L. Doherty, L. E. Ghaoui, and S. J. Pister. Convex position estimation in wireless sensor networks. In *IEEE Infocom*, volume 3, pages 1655–1663, April 2001.

- [23] S. Durocher, D. Kirkpatrick, and L. Narayanan. On routing with guaranteed delivery in three-dimensional ad hoc wireless networks. In *Proceedings of the 9th international conference on Distributed computing and networking*, ICDCN'08, pages 546–557, Berlin, Heidelberg, 2008. Springer-Verlag.
- [24] A. Efrat, C. Erten, D. Forrester, A. Iyer, and S. G. Kobourov. Force-directed approaches to sensor localization. In *Proc. of the 8th Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 108–118, 2006.
- [25] D. Eppstein and M. T. Goodrich. Succinct greedy graph drawing in the hyperbolic plane. In *Proc. of the 16th International Symposium on Graph Drawing*, pages 14–25, 2008.
- [26] T. Eren, D. Goldenberg, W. Whitley, Y. Yang, S. Morse, B. Anderson, and P. Belhumeur. Rigidity, computation, and randomization of network localization. In *Proceedings of the 23rd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'04)*, volume 4, pages 2673–2684, March 2004.
- [27] Q. Fang, J. Gao, and L. Guibas. Locating and bypassing routing holes in sensor networks. In *Mobile Networks and Applications*, volume 11, pages 187–200, 2006.
- [28] 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.
- [29] Q. Fang, J. Gao, and L. J. Guibas. Landmark-based information storage and retrieval in sensor networks. In *The 25th Conference of the IEEE Communication Society (INFOCOM'06)*, volume 1, pages 339–350, April 2006.
- [30] S. P. Fekete, M. Kaufmann, A. Kröller, and N. Lehmann. A new approach for boundary recognition in geometric sensor networks. In *Proceedings 17th Canadian Conference on Computational Geometry*, pages 82–85, 2005.
- [31] S. P. Fekete, A. Kröller, D. Pfisterer, S. Fischer, and C. Buschmann. Neighborhood-based topology recognition in sensor networks. In *ALGOSENSORS*, volume 3121 of *Lecture Notes in Computer Science*, pages 123–136. Springer, 2004.
- [32] G. G. Finn. Routing and addressing problems in large metropolitan-scale internetworks. Technical Report ISU/RR-87-180, ISI, March 1987.
- [33] R. Flury, S. Pemmaraju, and R. Wattenhofer. Greedy routing with bounded stretch. In *Proc. of the 28th Annual IEEE Conference on Computer Communications (INFOCOM)*, April 2009.
- [34] R. Fonesca, S. Ratnasamy, J. Zhao, C. T. Ee, D. Culler, S. Shenker, and I. Stoica. Beacon vector routing: Scalable point-to-point routing in wireless sensornets. In *Proc. of the 2nd Symposium on Networked Systems Design and Implementation (NSDI)*, pages 329–342, May 2005.
- [35] H. Frey and I. Stojmenovic. On delivery guarantees of face and combined greedy-face routing in ad hoc and sensor networks. In *Proceedings of the 12th annual international conference on Mobile computing and networking*, MobiCom '06, pages 390–401, New York, NY, USA, 2006. ACM.
- [36] S. Funke. Topological hole detection in wireless sensor networks and its applications. In *DIALM-POMC '05: Proceedings of the 2005 Joint Workshop on Foundations of Mobile Computing*, pages 44–53, 2005.
- [37] S. Funke, L. J. Guibas, A. Nguyen, and Y. Wang. Distance-sensitive routing and information brokerage in sensor networks. In *IEEE International Conference on Distributed Computing in Sensor System (DCOSS'06)*, pages 234–251, June 2006.
- [38] S. Funke and C. Klein. Hole detection or: “how much geometry hides in connectivity?”. In *SCG '06: Proceedings of the twenty-second annual symposium on Computational geometry*, pages 377–385, 2006.
- [39] S. Funke and N. Milosavljević. Guaranteed-delivery geographic routing under uncertain node locations. In *Proceedings of the 26th Conference of the IEEE Communications Society (INFOCOM'07)*, pages 1244–1252, May 2007.

- [40] S. Funke and N. Milosavljević. Network sketching or: “how much geometry hides in connectivity? - part II”. In *SODA '07: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 958–967, 2007.
- [41] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. Technical Report UCLA/CSD-TR 02-0013, UCLA, 2002.
- [42] J. Gao, L. Guibas, J. Hershberger, and L. Zhang. Fractionally cascaded information in a sensor network. In *Proc. of the 3rd International Symposium on Information Processing in Sensor Networks (IPSN'04)*, pages 311–319, April 2004.
- [43] J. Gao, L. Guibas, and A. Nguyen. Deformable spanners and their applications. *Computational Geometry: Theory and Applications*, 35(1-2):2–19, 2006.
- [44] J. Gao, L. J. Guibas, J. Hershberger, and N. Milosavljevic. Sparse data aggregation in sensor networks. In *Proc. of the International Conference on Information Processing in Sensor Networks (IPSN'07)*, pages 430–439, April 2007.
- [45] J. Gao, L. J. Guibas, J. Hershberger, L. Zhang, and A. Zhu. Geometric spanners for routing in mobile networks. *IEEE Journal on Selected Areas in Communications Special issue on Wireless Ad Hoc Networks*, 23(1):174–185, 2005.
- [46] J. Gao and L. Zhang. Load balanced short path routing in wireless networks. *IEEE Transactions on Parallel and Distributed Systems, Special Issue on Localized Communications*, 17(4):377–388, April 2006.
- [47] R. Ghrist and A. Muhammad. Coverage and hole-detection in sensor networks via homology. In *Proc. the 4th International Symposium on Information Processing in Sensor Networks (IPSN'05)*, pages 254–260, 2005.
- [48] D. Goldenberg, A. Krishnamurthy, W. Maness, Y. R. Yang, A. Young, A. S. Morse, A. Savvides, and B. Anderson. Network localization in partially localizable networks. In *Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'05)*, volume 1, pages 313–326, March 2005.
- [49] D. K. Goldenberg, P. Bihler, Y. R. Yang, M. Cao, J. Fang, A. S. Morse, and B. D. O. Anderson. Localization in sparse networks using sweeps. In *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*, pages 110–121, New York, NY, USA, 2006. ACM Press.
- [50] C. Gotsman and Y. Koren. Distributed graph layout for sensor networks. In *Proceedings of the International Symposium on Graph Drawing*, pages 273–284, September 2004.
- [51] O. Goussevskaia, Y.-A. Pignolet, and R. Wattenhofer. Efficiency of wireless networks: Approximation algorithms for the physical interference model. *Found. Trends Netw.*, 4:313–420, March 2010.
- [52] Y.-J. K. R. Govindan, B. Karp, and S. Shenker. Lazy cross-link removal for geographic routing. In *SenSys '06: Proceedings of the 4th international conference on Embedded networked sensor systems*, pages 112–124, 2006.
- [53] J. E. Graver, B. Servatius, and H. Servatius. *Combinatorial Rigidity*. Graduate Studies in Math., AMS, 1993.
- [54] 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 International Workshop on Sensor Network Protocols and Applications*, pages 163–173, Anchorage, Alaska, May 2003.
- [55] X. He and H. Zhang. On succinct convex greedy drawing of 3-connected plane graphs. In *Proceedings of the ACM-SIAM symposium on Discrete algorithms*, pages 1477–1486, January 2011.
- [56] B. Hendrickson. Conditions for unique graph realizations. *SIAM J. Comput.*, 21(1):65–84, 1992.
- [57] B. Hofmann-Wellenhof, H. Lichtenegger, and J. Collins. *Global Positioning Systems: Theory and Practice*. Springer, 5 edition, 2001.

- [58] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed diffusion: a scalable and robust communication paradigm for sensor networks. In *ACM Conf. on Mobile Computing and Networking (MobiCom)*, pages 56–67, 2000.
- [59] D. J. Jacobs and B. Hendrickson. An algorithm for two-dimensional rigidity percolation: the pebble game. *J. Comput. Phys.*, 137(2):346–365, 1997.
- [60] R. Jain, A. Puri, and R. Sengupta. Geographical routing using partial information for wireless ad hoc networks. *IEEE Personal Communications*, 8(1):48–57, Feb. 2001.
- [61] B. Karp and H. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *Proc. of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 243–254, 2000.
- [62] B. Katz and D. Wagner. Multiscale anchor-free distributed positioning in sensor networks. In S. K. Makki, X.-Y. Li, N. Pissinou, S. Makki, M. Karimi, and K. Makki, editors, *Sensor and Ad-Hoc Networks*, volume 7 of *Lecture Notes in Electrical Engineering*, pages 1–18. Springer US, 2008.
- [63] Y.-J. Kim, R. Govindan, B. Karp, and S. Shenker. Geographic routing made practical. In *Proceedings of the Second USENIX/ACM Symposium on Networked System Design and Implementation (NSDI 2005)*, May 2005.
- [64] Y.-J. Kim, R. Govindan, B. Karp, and S. Shenker. On the pitfalls of geographic face routing. In *DIALM-POMC '05: Proceedings of the 2005 joint workshop on Foundations of mobile computing*, pages 34–43, 2005.
- [65] R. Kleinberg. Geographic routing using hyperbolic space. In *Proceedings of the 26th Conference of the IEEE Communications Society (INFOCOM'07)*, pages 1902–1909, 2007.
- [66] E. Kranakis, H. Singh, and J. Urrutia. Compass routing on geometric networks. In *Proc. 11th Canadian Conference on Computational Geometry*, pages 51–54, 1999.
- [67] A. Kröller, S. P. Fekete, D. Pfisterer, and S. Fischer. Deterministic boundary recognition and topology extraction for large sensor networks. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1000–1009, 2006.
- [68] A. Kröller, S. P. Fekete, D. Pfisterer, and S. Fischer. Deterministic boundary recognition and topology extraction for large sensor networks. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'06)*, pages 1000–1009, 2006.
- [69] F. Kuhn, T. Moscibroda, and R. Wattenhofer. Unit disk graph approximation. In *Proceedings of the 2004 Joint Workshop on Foundations of Mobile Computing*, pages 17–23, 2004.
- [70] F. Kuhn, R. Wattenhofer, Y. Zhang, and A. Zollinger. Geometric ad-hoc routing: Of theory and practice. In *Proc. 22nd ACM Int. Symposium on the Principles of Distributed Computing (PODC)*, pages 63–72, 2003.
- [71] F. Kuhn, R. Wattenhofer, and A. Zollinger. Asymptotically optimal geometric mobile ad-hoc routing. In *Proc. of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, pages 24–33, 2002.
- [72] M. A. Labrador and P. M. Wightman. *Topology Control in Wireless Sensor Networks: With a Comparison Simulation Tool for Teaching and Research*. Springer, 2009.
- [73] S. Lederer, Y. Wang, and J. Gao. Connectivity-based localization of large scale sensor networks with complex shape. *ACM Transactions on Sensor Networks*, 5(4), November 2009.
- [74] T. Leighton and A. Moitra. Some results on greedy embeddings in metric spaces. In *Proc. of the 49th IEEE Annual Symposium on Foundations of Computer Science*, pages 337–346, October 2008.
- [75] F. Li and Y. Wang. Circular sailing routing for wireless networks. In *INFOCOM*, pages 1346–1354, 2008.
- [76] M. Li and Y. Liu. Rendered path: range-free localization in anisotropic sensor networks with holes. In *Proceedings of the 13th annual ACM international conference on Mobile computing and networking, MobiCom '07*, pages 51–62, New York, NY, USA, 2007. ACM.

- [77] X. Li, Y. J. Kim, R. Govindan, and W. Hong. Multi-dimensional range queries in sensor networks. In *Proceedings of the first international conference on Embedded networked sensor systems*, pages 63–75, 2003.
- [78] Z. Li, W. Trappe, Y. Zhang, and B. Nath. Robust statistical methods for securing wireless localization in sensor networks. In *IPSN '05: Proceedings of the 4th international symposium on Information processing in sensor networks*, pages 91–98, Piscataway, NJ, USA, 2005. IEEE Press.
- [79] K. M. Lillis and S. V. Pemmaraju. On the efficiency of a local iterative algorithm to compute delaunay realizations. In *Proceedings of the 7th international conference on Experimental algorithms*, WEA'08, pages 69–86, Berlin, Heidelberg, 2008. Springer-Verlag.
- [80] H. Lin, M. Lu, N. Milosavljević, J. Gao, and L. Guibas. Composable information gradients in wireless sensor networks. In *Proc. of the International Conference on Information Processing in Sensor Networks (IPSN'08)*, pages 121–132, April 2008.
- [81] N. Linial, L. Lovász, and A. Wigderson. Rubber bands, convex embeddings and graph connectivity. *Combinatorica*, 8(1):91–102, 1988.
- [82] X. Liu, Q. Huang, and Y. Zhang. Combs, needles, haystacks: balancing push and pull for discovery in large-scale sensor networks. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 122–133, 2004.
- [83] Z. Lotker and D. Peleg. Structure and algorithms in the sinr wireless model. *SIGACT News*, 41:74–84, June 2010.
- [84] S. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. TAG: a tiny aggregation service for ad-hoc sensor networks. *SIGOPS Oper. Syst. Rev.*, 36(SI):131–146, 2002.
- [85] Y. Mao, F. Wang, L. Qiu, S. S. Lam, and J. M. Smith. S4: Small state and small stretch compact routing protocol for large static wireless networks. *IEEE/ACM Trans. Netw.*, 18:761–774, June.
- [86] P. Maymounkov. Greedy embeddings, trees, and euclidean vs. lobachevsky geometry. manuscript, 2006.
- [87] A. Mei and J. Stefa. Routing in outer space: fair traffic load in multi-hop wireless networks. In *MobiHoc '08: Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*, pages 23–32, New York, NY, USA, 2008. ACM.
- [88] D. Moore, J. Leonard, D. Rus, and S. Teller. Robust distributed network localization with noisy range measurements. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 50–61, New York, NY, USA, 2004. ACM Press.
- [89] A. Nguyen, N. Milosavljevic, Q. Fang, J. Gao, and L. J. Guibas. Landmark selection and greedy landmark-descent routing for sensor networks. In *Proceedings of the 26th Conference of the IEEE Communications Society (INFOCOM'07)*, pages 661–669, May 2007.
- [90] D. Niculescu and B. Nath. Ad hoc positioning system (APS). In *IEEE GLOBECOM*, pages 2926–2931, 2001.
- [91] D. Niculescu and B. Nath. Ad hoc positioning system (APS) using AOA. In *IEEE INFOCOM*, volume 22, pages 1734–1743, March 2003.
- [92] D. Niculescu and B. Nath. Error characteristics of ad hoc positioning systems (APS). In *MobiHoc '04: Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 20–30, 2004.
- [93] S. Y. Oudot, L. J. Guibas, J. Gao, and Y. Wang. Geodesic delaunay triangulations in bounded planar domains. *ACM Transactions on Algorithms (TALG)*, 6(4):61:1–61:47, 2010.
- [94] C. H. Papadimitriou and D. Ratajczak. On a conjecture related to geometric routing. *Theor. Comput. Sci.*, 344(1):3–14, 2005.

- [95] L. Popa, A. Rostamizadeh, R. Karp, C. Papadimitriou, and I. Stoica. Balancing traffic load in wireless networks with curveball routing. In *MobiHoc '07: Proceedings of the 8th ACM international symposium on Mobile ad hoc networking and computing*, pages 170–179, New York, NY, USA, 2007. ACM.
- [96] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *MobiCom '00: Proceedings of the 6th ACM Annual International Conference on Mobile Computing and Networking*, pages 32–43, 2000.
- [97] P. Raghavan. Probabilistic construction of deterministic algorithms: approximating packing integer programs. *J. Comp. and System Sciences*, pages 130–143, 1988.
- [98] P. Raghavan and C. D. Thompson. Provably good routing in graphs: regular arrays. In *Proceedings of the 17th annual ACM Symposium on Theory of Computing*, pages 79–87, 1985.
- [99] A. Rao, C. Papadimitriou, S. Shenker, and I. Stoica. Geographic routing without location information. In *Proceedings of the 9th annual international conference on Mobile computing and networking*, pages 96–108, 2003.
- [100] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content-addressable network. In *SIGCOMM '01: Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 161–172, 2001.
- [101] S. Ratnasamy, B. Karp, L. Yin, F. Yu, D. Estrin, R. Govindan, and S. Shenker. GHT: A geographic hash table for data-centric storage in sensor networks. In *Proc. 1st ACM Workshop on Wireless Sensor Networks and Applications*, pages 78–87, 2002.
- [102] A. I. T. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In *Middleware '01: Proceedings of the IFIP/ACM International Conference on Distributed Systems Platforms Heidelberg*, pages 329–350, London, UK, 2001. Springer-Verlag.
- [103] P. Santi. *Topology Control in Wireless Ad Hoc and Sensor Networks*. John Wiley & Sons, Sept. 2005.
- [104] R. Sarkar and J. Gao. Differential forms for target tracking and aggregate queries in distributed networks. In *Proc. of the 16th Annual International Conference on Mobile Computing and Networking (MobiCom'10)*, pages 377–388, September 2010.
- [105] R. Sarkar, X. Yin, J. Gao, F. Luo, and X. D. Gu. Greedy routing with guaranteed delivery using ricci flows. In *Proc. of the 8th International Symposium on Information Processing in Sensor Networks (IPSN'09)*, pages 97–108, April 2009.
- [106] R. Sarkar, W. Zeng, J. Gao, and X. D. Gu. Covering space for in-network sensor data storage. In *Proc. of the 9th International Symposium on Information Processing in Sensor Networks (IPSN'10)*, pages 232–243, April 2010.
- [107] R. Sarkar, X. Zhu, and J. Gao. Double rulings for information brokerage in sensor networks. In *Proc. of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 286–297, September 2006.
- [108] O. Saukh, R. Sauter, M. Gauger, and P. J. Marrón. On boundary recognition without location information in wireless sensor networks. *ACM Trans. Sen. Netw.*, 6:20:1–20:35, June 2010.
- [109] A. Savvides, C.-C. Han, and M. B. Strivastava. Dynamic fine-grained localization in *ad-hoc* networks of sensors. In *Proc. 7th Annual International Conference on Mobile Computing and Networking (MobiCom 2001)*, pages 166–179, Rome, Italy, July 2001. ACM Press.
- [110] A. Savvides, H. Park, and M. B. Strivastava. The n -hop multilateration primitive for node localization problems. *Mobile Networks and Applications*, 8(4):443–451, 2003.
- [111] J. B. Saxe. Embeddability of weighted graphs in k -space is strongly NP-hard. In *Proceedings of the 17th Allerton Conference in Communications, Control and Computing*, pages 480–489, 1979.

- [112] K. Seada, A. Helmy, and R. Govindan. On the effect of localization errors on geographic face routing in sensor networks. In *IPSN '04: Proceedings of the third international symposium on Information processing in sensor networks*, pages 71–80, 2004.
- [113] Y. Shang, W. Ruml, Y. Zhang, and M. P. J. Fromherz. Localization from mere connectivity. In *MobiHoc '03: Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 201–212, 2003.
- [114] S. Shenker, S. Ratnasamy, B. Karp, R. Govindan, and D. Estrin. Data-centric storage in sensor networks. *SIGCOMM Comput. Commun. Rev.*, 33(1):137–142, 2003.
- [115] P. Skraba, Q. Fang, A. Nguyen, and L. Guibas. Sweeps over wireless sensor networks. In *IPSN '06: Proceedings of the fifth international conference on Information processing in sensor networks*, pages 143–151, 2006.
- [116] A. M.-C. So and Y. Ye. Theory of semidefinite programming for sensor network localization. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 405–414, 2005.
- [117] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *SIGCOMM '01: Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 149–160, 2001.
- [118] I. Stojmenovic. A routing strategy and quorum based location update scheme for ad hoc wireless networks. Technical Report TR-99-09, SITE, University of Ottawa, September, 1999.
- [119] I. Stojmenovic and X. Lin. Loop-free hybrid single-path/flooding routing algorithms with guaranteed delivery for wireless networks. *IEEE Trans. Parallel Dist. Sys.*, 12(10):1023–1032, 2001.
- [120] I. Stojmenovic and X. Lin. Power-aware localized routing in wireless networks. *IEEE Trans. Parallel Dist. Sys.*, 12(11):1122–1133, 2001.
- [121] H. Takagi and L. Kleinrock. Optimal transmission ranges for randomly distributed packet radio terminals. *IEEE Trans. Commun.*, 32(3):246–257, 1984.
- [122] G. Tan, H. Jiang, S. Z. Zhang, and A.-M. Kermarrec. Connectivity-based and anchor-free localization in large-scale 2d/3d sensor networks. In *Proceedings of the eleventh ACM international symposium on Mobile ad hoc networking and computing*, MobiHoc '10, pages 191–200, New York, NY, USA, 2010. ACM.
- [123] R. Thomas and X. Yu. 4-connected projective-planar graphs are hamiltonian. *J. Comb. Theory Ser. B*, 62(1):114–132, 1994.
- [124] M. Thorup and U. Zwick. Approximate distance oracles. In *Proc. ACM Symposium on Theory of Computing*, pages 183–192, 2001.
- [125] M. Thorup and U. Zwick. Compact routing schemes. In *SPAA '01: Proceedings of the thirteenth annual ACM symposium on Parallel algorithms and architectures*, pages 1–10, 2001.
- [126] P. F. Tsuchiya. The landmark hierarchy: a new hierarchy for routing in very large networks. In *SIGCOMM '88: Symposium proceedings on Communications architectures and protocols*, pages 35–42, 1988.
- [127] Y. Wang, J. Gao, and J. S. B. Mitchell. Boundary recognition in sensor networks by topological methods. In *Proc. of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 122–133, September 2006.
- [128] M. Wattenhofer, R. Wattenhofer, and P. Widmayer. Geometric Routing without Geometry. In *12th Colloquium on Structural Information and Communication Complexity (SIROCCO)*, Le Mont Saint-Michel, France, May 2005.
- [129] F. Ye, H. Luo, J. Cheng, S. Lu, and L. Zhang. A two-tier data dissemination model for large-scale wireless sensor networks. In *MobiCom '02: Proceedings of the 8th annual international conference on Mobile computing and networking*, pages 148–159, 2002.

- [130] X. Yu, X. Ban, R. Sarkar, W. Zeng, X. D. Gu, and J. Gao. Spherical representation and polyhedron routing for load balancing in wireless sensor networks. In *Proc. of 30th Annual IEEE Conference on Computer Communications (INFOCOM'11)*, April 2011.
- [131] B. Y. Zhao, J. D. Kubiatowicz, and A. D. Joseph. Tapestry: An infrastructure for fault-tolerant wide-area location and. Technical report, Berkeley, CA, USA, 2001.
- [132] G. Zhou, T. He, S. Krishnamurthy, and J. A. Stankovic. Impact of radio irregularity on wireless sensor networks. In *MobiSys '04: Proceedings of the 2nd international conference on Mobile systems, applications, and services*, pages 125–138, 2004.
- [133] G. Ziegler. *Lectures on Polytopes*. Springer-Verlag, 1995.