

Spatial Distribution in Routing Table Design for Sensor Networks

Rik Sarkar Xianjin Zhu Jie Gao

Department of Computer Science, Stony Brook University. {rik, xjzhu, jgao}@cs.sunysb.edu

Abstract—We propose a generic routing table design principle for scalable routing on networks with bounded geometric growth. Given an inaccurate distance oracle that estimates the graph distance of any two nodes with constant factor upper and lower bounds, we augment it by storing the routing paths of pairs of nodes, selected in a spatial distribution, and show that the routing table enables $1 + \varepsilon$ stretch routing. In the wireless ad hoc and sensor network scenario, the geographic locations of the nodes serve as such an inaccurate distance oracle. Each node p selects $O(\log n \log \log n)$ other nodes from a distribution proportional to $1/r^2$ where r is the distance to p and the routing paths to these nodes are stored on the nodes along these paths in the network. The routing algorithm selects links conforming to a set of sufficient conditions and guarantees with high probability $1 + \varepsilon$ stretch routing with routing table size $O(\sqrt{n} \log n \log \log n)$ on average for each node. This scheme is favorable for its simplicity, generality and blindness to any global state. It is a good example that global routing properties emerge from purely distributed and uncoordinated routing table design.

I. INTRODUCTION

Scalable routing is one of the most challenging problems in distributed network design — considerations include compact storage with aggressive address aggregation, efficient ways to propagate and manage topology update, and most importantly, ways for the distributed and uncoordinated decisions to enable globally close to optimal routing properties.

Structures of large-scale wireless sensor networks are closely related to the underlying geometric domain in which they are embedded. Therefore various properties of the geographical embedding of the nodes are exploited for routing in a sensor network — mostly in an explicit manner, as the geographical locations used in geographical routing families [1]–[3], or as in many virtual coordinate system design [4], [5] that abstracts the global geometric/topological properties of the embedding.

A. Our problem and solution overview

The problem we study in this paper is as follows: Given an inaccurate distance oracle \mathcal{O} , what routing paths should be kept in the routing table, such that the average routing table size is small, the path stretch is close to optimal ($1 + \varepsilon$ for any given $\varepsilon > 0$), and both the preprocessing and the routing can be achieved by the nodes making decisions on their own, blind to any global state?

First we remark that if we are given an *accurate* distance oracle that returns the hop count distance of any two nodes in the network, then routing is trivial as greedy routing based on this oracle will guarantee delivery along the shortest path — a message can be delivered to the neighbor whose distance (in the graph) to the destination is the smallest. Of course, the construction, maintenance and compact representation of

an accurate distance oracle is not easy in a distributed setting. As shown in [6], accurate distance oracle would require about $\Omega(n)$ storage per node. If we relax the condition and only use approximate distances, we can hope to reduce the storage requirement. An approximate distance oracle is also easier to obtain. In many cases, some inaccurate distance estimation is readily available.

For an example, in the sensor network setting, one can use the Euclidean distance to approximate the hop count distance of two nodes in the network. Of course it is not an accurate distance oracle as the graph structure may not be aligned well with the Euclidean coordinates, and that a message can get stuck at a local minimum if the neighbor on the shortest path to the destination estimates its distance to be larger than the distance estimation between source and destination [1], [2]. Therefore we will have to augment the inaccurate distance oracle with additional routing information to help packets get out of the local minimum.

With any inaccurate distance oracle, we augment in the routing table some routing information between pairs of nodes that are not immediate neighbors, called *long links*. In particular, for some selected pairs u, v a path between u, v , $P(u, v)$, is recorded in the routing table of all nodes on this path. When a node p wants to send a message to a node q , it uses its immediate neighbors, together with the nodes with which p has long links. Based on a set of sufficient conditions, we define a *forwarding region* (see Fig. 1) from which p selects the next hop in the path. If the selected node x is a neighbor through a long link, then the routing information stored on the path $P(p, x)$ is used to deliver the message to x . Node x then repeats an identical procedure to advance the message. Now the question is, what long links should each node build and which node should be selected in the routing stage, without knowing the global state, such that the routing table size is small, the path stretch is low, and delivery rate is high?

Our main theoretical results are the following. We first describe the special case with Euclidean distance as an approximate distance oracle. When the sensors are deployed in an Euclidean plane such that the Euclidean distance is an approximate distance oracle to the graph distance metric $\sigma(\cdot, \cdot)$, i.e., $\delta_1 |pq| \leq \sigma(p, q) \leq \delta_2 |pq|$, with $\delta_1 \leq \delta_2$ as two constants, and $|pq|$ the Euclidean distance between p, q . This is a reasonable assumption that makes no unit disk graph requirement on the wireless radio communication model and uses two relaxation constants δ_1 and δ_2 incorporating both

local distance disturbances and possible holes that are fat¹. It also allows localization errors as accurate location discovery is difficult. The routing tables are built by each node selecting its long links randomly with a spatial distribution. The routing algorithm using the augmented long links is able to deliver the message along a path of stretch $1 + \varepsilon$.

In fact, the theoretical results in this paper address a general setting in which an inaccurate distance oracle is given for a graph (more generally, a metric space) with bounded growth rate. A graph has bounded growth rate ρ if the number of nodes within r hops from any node p in the network is bounded by $c_1 r^\rho$ and $c_2 r^\rho$ from below and above respectively, with two constants $c_1 \leq c_2$. This model has been used to capture any physical constraints that disallow too many nodes ‘packed’ within certain distance and the graph has a geometric growth pattern instead of an exponential growth pattern (e.g., a balanced binary tree). This kind of geometric growth has been observed in many different scenarios such as VLSI design, the delay metric on the Internet overlay networks, and in our setting, wireless sensor networks. When sensor nodes are roughly uniformly deployed in a geometric region with bounded density per unit area² and when the network is not too much fragmented by deployment holes, the graph growth rate is typically 2. It is this packing property that allows us to aggressively compress the routing table entries by a simple routing table neighbor selection rule dominated by a spatial distribution.

The long link p, q is selected by p with probability proportional to $1/r^\rho$, for $r = \sigma(p, q)$ being the network distance between p, q . Selecting $O(\log n \log \log n)$ such long links per node produces average routing table size of $O(n^{1/\rho} \log n \log \log n)$ per node, and we show that a $1 + \varepsilon$ stretch path is obtained with high probability. Thus this principle of using spatial distribution in routing table design is in fact quite general and can be applied to many other potential applications in which a decentralized search is desired with only some vague proximity information. For example, in an ad hoc network setting when location information is not available, one can use other distance estimation to the network distance (e.g., by landmark scheme [4], [7]). In the design of overlay networks on the Internet, one can estimate the distance between two peers by the round-trip delay estimation. For all these scenarios the results in this paper show a way to achieve distributed routing along approximate shortest paths with a modest sized routing table on each node.

We also report simulation evaluations of this approach in a sensor network setting, to complement the theoretical analysis. For a connectivity network in which geographical greedy routing only achieves a delivery rate of 50% or so, with about 7 long links per node, we are able to achieve a delivery rate of 99% or higher. The routing table construction can be implemented

¹A hole is fat if any two nodes on the boundary of a hole has its hop count distance to be at most a constant factor of the Euclidean distance.

²If the density in a region becomes too high, it is easily possible to cluster neighboring nodes and operate on clusterheads so that density of clusterheads is not high

in a completely distributed manner as in [8].

When location information is not available, we have a second implementation by using landmark-based routing to show the power of the spatial distribution in routing table design. In particular, we select $O(\log n \log \log n)$ landmarks that flood the entire network and each node records the distances to all landmarks (called the landmark distance vector). The approximate distance oracle is implemented by the centered virtual distance as proposed in [4] on the landmark distance vectors of any two nodes. We select on the paths to the landmarks long link neighbors to help improve the delivery rate. This implementation will involve some preprocessing of flooding the network from the landmarks but the routing paths of the long links are implicitly implied by the landmark distances. Thus the routing table size is improved to $O(\log^2 n \log \log^2 n)$, compared with $O(n^{1/\rho} \log n \log \log n)$ when the routes have to be explicitly stored on the nodes of the paths.

In summary, the augmentation of long links with spatial distribution to get $1 + \varepsilon$ stretch routing on an approximate distance oracle is favorable for its simplicity, generality and ‘blindness’ to any global state. It is a good example that global routing properties emerge from purely distributed and uncoordinated routing table design.

B. Related work

In this subsection we will survey related work and establish their connection to our results.

Spatial distribution in routing. The spatial distribution in selecting the long links in our paper coincides with the small-world model and decentralized search (essentially the greedy routing) proposed by Kleinberg [9], [10] to model Stanley Milgram’s famous experiment [11], [12]. Our results show that if each node chooses $O(\log n)$ long links, our distributed routing scheme (slightly different from the decentralized search [9], [10]) with long links has $O(\log n)$ jumps, and also a total travel distance at most $1 + \varepsilon$ of the distance between source and destination on the grid.

The spatial distribution has been explored in a number of other data delivery and information dissemination scenarios in sensor networks, e.g., for adding long communication wires to reduce power consumption [13], or, for gossip and locality-sensitive information exchange [8], [14].

Small state routing in sensor networks. To deal with the problem of local minimum in geographical forwarding, various techniques have been proposed to solve the problem of ‘routing around holes’. Earlier proposals assume unit disk graph model on the communication network and propose to planarize the network and apply face routing [1]–[3]. In virtual ring routing (VRR) [15], proposed by Caesar *et al.*, the nodes are ordered by their node IDs (or any other identifiers) on a ring and the paths for nearby nodes on the ring are stored in the routing tables of the nodes on these paths.

The small state and small stretch (S4) routing by Mao *et al.* [16] adopted the idea of compact routing schemes by Thorup and Zwick [6], [17]. The basic idea is to select about $O(\sqrt{n})$ landmarks. These landmarks flood the network and

other nodes record the hop count distance to these landmarks. In addition, a node p also maintains routing table entries to the nodes that are closer to p than their closest landmarks. The routing table size is about $O(\sqrt{n})$ and a greedy routing scheme is guaranteed to deliver the message to the destination with maximum stretch of 3. By exploiting the geometric properties of the sensor network deployment, we are able to get $1 + \varepsilon$ stretch and reduce both the number of landmarks and the routing table size to polylogarithmic in the network size.

Compact routing in general. From a theoretical aspect, compact routing that minimizes the routing table size while achieving low stretch routing has been studied extensively [18]. We do not have the space to survey this broad topic here, instead, we refer the reader to the relevant papers in [17], [19]–[24] and related references.

II. ROUTING WITH SPATIAL DISTRIBUTIONS

In this section we describe the idea of using spatial distribution to route with $1 + \varepsilon$ stretch in a suitable metric space \mathcal{M} . We assume that a node is able to get the approximate distance $d(p, q)$ from just the names of p, q . The implementation of this distance oracle is beyond the scope of this short abstract. Due to lack of space, we can only present the core intuitions and theorems here. A full version containing all proofs is available online [25].

Accurate distance oracle. To demonstrate the basic concept, we first consider the case in which the oracle is in fact accurate, that is, $d = \sigma$. The objective is to recursively build a route from s to t with the help of the long links. Suppose s takes a long link to node p , then we want $\sigma(s, p) + \sigma(p, t)$ to be not very large compared to $\sigma(s, t)$:

$$\sigma(s, p) + \sigma(p, t) \leq \gamma \cdot \sigma(s, t), \quad (1)$$

Where $\gamma \geq 1$ is a parameter depending on ε . Observe that inequality (1) defines an ellipse in \mathbb{R}^2 with s and t at foci. Now we impose an additional restriction that moving from s to p implies a certain progress in direction of t . In particular, p is closer to t by a factor of at least $0 \leq \beta \leq 1$:

$$\sigma(p, t) \leq \beta \cdot \sigma(s, t). \quad (2)$$

This describes a disk centered at t .

Next, we select γ and β such that the selection procedure enforced by inequalities (1) and (2) when applied recursively, produces a path of stretch at most $1 + \varepsilon$:

$$R(s, t) \leq (1 + \varepsilon) \cdot \sigma(s, t), \quad (3)$$

where R gives the length of the path created recursively.

A *forwarding region* $F_\varepsilon(s, t)$ is a set of points p in \mathcal{M} from which s can select p satisfying the relations above. The following lemma gives a precise relation:

Lemma 2.1. *Values of γ and β satisfying $\gamma + \varepsilon\beta \leq 1 + \varepsilon$ constitute the forwarding region, with the equality corresponding to the region boundary.*

It is easy to see that γ must lie in the interval $[1, \frac{2+3\varepsilon}{2+\varepsilon}]$ for a given ε . For each value of γ , we have a region $H_{\gamma, \varepsilon}(s, t) \subseteq \mathcal{M}$

which is the intersection of the ellipse bounded region and the disk. Thus, formally, the forwarding region is the union: $F_\varepsilon(s, t) = \cup_\gamma H_{\gamma, \varepsilon}(s, t)$. See Figure 1.

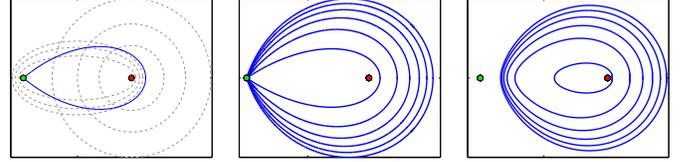


Fig. 1. (i) Boundary of F_ε as intersection of ellipses and circles. (ii) Forwarding regions for different values of ε from 0.2 to 2. (iii) Forwarding regions for different values of ε from 0.2 to 2 for approximate oracle. Observe that in this case the forwarding regions are smaller and source s is not in the forwarding region. This is due to inaccurate distance estimates and necessitates the use of *long links* - without which s cannot access the forwarding region.

Approximate distance oracle. For approximate distance oracles, it would be sufficient to guarantee the following inequalities (corresponding to relations (1)-(2) respectively):

$$\begin{aligned} \delta_2 d(s, p) + \delta_2 d(p, t) &\leq \gamma \delta_1 d(s, t) \\ \delta_2 d(p, t) &\leq \beta \delta_1 d(s, t) \end{aligned} \quad (4)$$

It can be verified that Lemma 2.1, with the same boundary condition holds here as well. Figure 1 shows the forwarding regions for the Euclidean case, with accurate and inaccurate oracles. Note that the shape of the forwarding regions is independent of the distance between the two nodes in question.

In a sensor graph setting, we use the (upper and lower) bounded growth rate model. If we place at most a constant number of sensor nodes inside any unit disk and the holes in the sensor networks are not very fragmenting, the number of nodes at k hops from a node p will be around $\Theta(k)$. In general, we consider a graph such that number of nodes in an r neighborhood $|N_r(p)| = \Theta(r^\rho)$. The forwarding region in a graph setting with source p and destination q is defined similarly as the nodes that satisfy the inequalities (4).

The analysis above suggests a natural routing scheme. Each node keeps the routing table entries for its immediate neighbors, as well as the long link neighbors it has selected. To route from s to t , a long link sp such that $p \in F_\varepsilon(s, t)$ has to exist. Our routing table design below shows how to select long links such that for any destination a neighbor in the forwarding region exists with high probability.

A. Routing table construction

To build the routing table, we use a spatial distribution (see [9]) of directed links. In particular, for nodes p and q separated by a distance r , the probability of a directed link pq being built by p is proportional to $1/r^\rho$.

Theorem 2.2. *From each node it is sufficient to select $O\left(\left(\frac{2}{\varepsilon}\right)^{O(\rho)} \ln n \ln \ln n\right)$ links, to guarantee a link in the forwarding region for every possible destination with probability $1 - 1/\log^{\Theta(1)} n$.*

The theorem above describes a guarantee for a suitable link to a forwarding region to exist. However, we still need to prove the existence of a path of $(1 + \varepsilon)$ stretch for a given routing request, that will take us to within a small constant distance of the destination. This is done by showing the existence of a short sequence of forwarding links.

Theorem 2.3. *It is sufficient to select $O\left(\left(\frac{2}{\varepsilon}\right)^{O(\rho)} \ln n \ln \ln n\right)$ long links per node to guarantee a path of stretch at most $1 + \varepsilon$ with probability at least $1 - \frac{1}{\rho \log^{\Theta(1)} n}$.*

And the routing table size is not too large.

Theorem 2.4. *The average routing table size of the scheme is bounded by $O\left(\left(\frac{2}{\varepsilon}\right)^{O(\rho)} n^{1/\rho} \ln n \ln \ln n\right)$.*

In the case of sensor networks in a plane ($\rho \approx 2$), for a given stretch ε , this amounts to a table size of $O(\sqrt{n} \ln n \ln \ln n)$ per node.

III. IMPLEMENTATION IN SENSOR NETWORKS

When geographical location is available and the Euclidean distance is a good approximate distance oracle, we can implement the selection of long links with a spatial geographical sampling as in [8]. The details are omitted here. Otherwise, we can use a landmark-based scheme. We select $m = O(\log n \log \log n)$ landmarks l_i uniformly randomly in the sensor network. The landmarks then flood the network and every other node p records a landmark-based distance vector, the vector of minimum hop count distance to all m landmarks. We used the centered distance measure proposed in [4] to approximate the graph distance.

Landmark-based sampling. To build the long links for a node p , we will use the landmarks to help with sampling. In particular, we select first randomly k out of the m landmarks. For each landmark l_i , we select from the distribution $1/(r \ln D)$ (D is the network diameter) a distance ξ . If $\xi \leq \sigma(p, l_i)$, we take the node q along the path from p to l_i with distance ξ from p as the long link partner. Otherwise we drop landmark l_i . Intuitively, we select along the path from p to l_i a node q with the spatial distribution restricted on this path. Since the landmarks are randomly selected, the probability that a landmark l_i is at distance r from p is proportional to r . Now the probability that for each landmark l_i we can obtain a valid long link is

$$\text{Prob}\{\xi < \sigma(p, l_i)\} = \int_0^D \int_1^\zeta \frac{1}{\xi \ln D} d\xi \frac{2\zeta}{D^2} d\zeta = 1 - \frac{1}{2 \ln D}.$$

Thus in expectation we obtain $k(1 - \frac{1}{2 \ln D})$ long links for each node. This means that choosing $m = O(\log n \log \log n)$ landmarks suffices to get enough long links for each node. At last we remark that although different nodes use the same set of landmarks to create their long links, the theoretical analysis in the previous section still holds – as the only requirement is that we have a sufficient number of independent long links for each individual node.

Landmark-based routing tables. With the long links constructed by the landmarks, the routing table size can be further reduced. In fact, a node p remembers in its routing table the long link partners and their landmark-based addresses. Different from the geographical case, the routes for the long links are implicitly implied by the landmark distances. The size of the routing table is therefore $O(\log^2 n \log \log^2 n)$, for $O(\log n \log \log n)$ landmarks/long link neighbors, and a storage of $O(\log n \log \log n)$ for storing the address of each long link neighbor.

IV. SIMULATIONS

We compare our approach with two recently proposed routing protocols, S4 [16] and VRR [15], on three important criteria - delivery rate, the size of routing table and routing stretch. In summary, our approach achieves high delivery rate (above 99%) and small stretch (about 1.03) with only a small number of long links, and a small routing table with modest preprocessing.

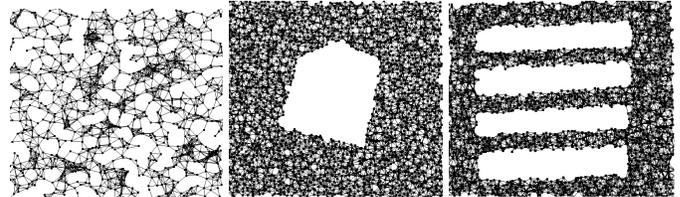


Fig. 2. Network topologies used in simulations. (i) Topology 1. Random network: 1000 nodes, avg. degree 7.2; (ii) Topology 2. Network with one hole: 2400 nodes, avg. degree 9.5; (iii) Topology 3. Network with multiple holes: 2000 nodes, avg. degree 10.6.

Simulation setup. We focus on evaluating the performance of all approaches at the routing layer. We adopt a lossy radio model used in the standard simulator TOSSIM [26] to determine direct communication links between nodes, and only consider links with sufficient low loss rate. We run simulations on three typical topologies as in Figure 2. Each simulation run is repeated 10 times. In each round, we randomly selected 10000 pairs of source and destination. All results are averaged over all pairs.

A. Geographic routing table

Delivery rate. To show the effect of long links on the delivery rate, we vary the number of long links each node maintains from 0 to 16. When the number of long links is set to 0, the routing protocol is essentially the geographical greedy routing based on the location information within one-hop neighborhood. Figure 3 (i) shows that greedy routing performs very poorly without long links. The delivery rate is only around 50%, 65% and 44% in Topology 1, 2 and 3 respectively. When the number of long links increases, the delivery rate reaches 99% with 6, 8, 7 long links per node in three different topologies, respectively. The results confirm that a small number of long links can significantly improve the delivery rate in most of typical network topologies. Since our scheme behaves similarly in various topologies, in the rest of this subsection, unless mentioned otherwise, we only present results on Topology 2 due to space limitation.

Routing table size. We compare the average routing table size of our scheme with VRR and S4. Our scheme uses much smaller routing table than VRR when maintaining the same number of long links.

Size of routing table	Our scheme	S4	VRR
Topology 1	26.08	68.83	41.52
Topology 2	39.02	105.85	62.48
Topology 3	37.28	90.62	63.82

TABLE I. Average size of routing table.

Table I shows the routing table size of three schemes with a set of fixed parameters. For comparisons, we use 50 landmarks for S4 and each node maintains routes to 4 virtual neighbors

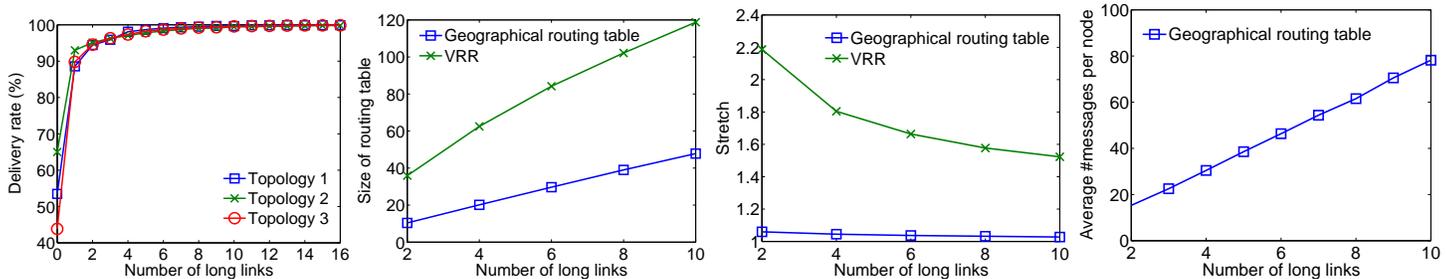


Fig. 3. (i) Delivery rate of geographical routing table with varying number of long links in different network topologies. (ii)-(iv) Performance of our scheme and VRR in Topology 2. (ii) The average size of routing table. (iii) Average stretch. (iv) Communication cost in preprocessing stage.

in VRR. We select those parameters since they give the best performance of S4 and VRR in terms of both routing table size and stretch. For our scheme, we use 6, 8, 7 long links in three topologies respectively to get above 99% delivery rate. We use the same set of parameters in other Tables. From Table I, S4 requires the largest routing table (in the order of $O(\sqrt{n})$). Our scheme has the smallest routing table size, but achieves comparable delivery rate.

Average stretch	Our scheme	S4	VRR
Topology 1	1.03	1.03	1.73
Topology 2	1.03	1.03	1.80
Topology 3	1.04	1.02	1.75

TABLE II. Average stretch.

Stretch. Figure 3(iii) shows the average stretch of our scheme and VRR with varying number of long links. The stretch of our scheme is always below 1.1 and decreases when the number of long links increases. With 6 long links, the stretch is only about 1.03. Table II compares the average stretch of three schemes. It shows that our scheme achieves similar stretch as S4 (but with smaller routing table) and is better than VRR.

V. CONCLUSION

This paper outlines a theory to build useful routing links in very general domains. The method is distributed and uncoordinated, but guarantees global properties such as routing with low stretch and compact routing tables. The use of spatial distribution ensures that the routing works well at all scales and distances.

REFERENCES

- [1] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia, "Routing with guaranteed delivery in ad hoc wireless networks," *Wireless Networks*, vol. 7, no. 6, pp. 609–616, 2001. [Online]. Available: citeseer.ist.psu.edu/bose01routing.html
- [2] 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)*, 2000, pp. 243–254.
- [3] F. Kuhn, R. Wattenhofer, Y. Zhang, and A. Zollinger, "Geometric ad-hoc routing: of theory and practice," in *PODC '03: Proceedings of the twenty-second annual symposium on Principles of distributed computing*, 2003, pp. 63–72.
- [4] 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)*, vol. 1, March 2005, pp. 339–350.
- [5] J. Bruck, J. Gao, and A. Jiang, "MAP: Medial axis based geometric routing in sensor networks," *Wireless Networks*, vol. 13, no. 6, pp. 835–853, 2007.
- [6] M. Thorup and U. Zwick, "Approximate distance oracles," in *Proc. ACM Symposium on Theory of Computing*, 2001, pp. 183–192.
- [7] J. Kleinberg, A. Slivkins, and T. Wexler, "Triangulation and embedding using small sets of beacons," in *Proc. 45th IEEE Symposium on Foundations of Computer Science*, 2004, pp. 444–453.
- [8] R. Sarkar, X. Zhu, and J. Gao, "Hierarchical spatial gossip for multi-resolution representations in sensor networks," in *Proc. of the International Conference on Information Processing in Sensor Networks (IPSN'07)*, April 2007, pp. 420–429.
- [9] J. Kleinberg, "The small-world phenomenon: an algorithm perspective," in *STOC '00: Proceedings of the thirty-second annual ACM symposium on Theory of computing*, 2000, pp. 163–170.
- [10] —, "Navigation in a small world," *Nature*, no. 406, p. 845, 2000.
- [11] S. Milgram, "The small world problem," *Psychology Today*, no. 1, 1967.
- [12] J. Travers and S. Milgram, "An experimental study of the small world problem," *Sociometry*, no. 32, 1969.
- [13] G. Sharma and R. Mazumdar, "Hybrid sensor networks: a small world," in *MobiHoc '05: Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing*, 2005, pp. 366–377.
- [14] D. Kempe, J. Kleinberg, and A. Demers, "Spatial gossip and resource location protocols," in *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, 2001, pp. 163–172.
- [15] M. Caesar, M. Castro, E. B. Nightingale, G. O'Shea, and A. Rowstron, "Virtual ring routing: network routing inspired by DHTs," in *SIGCOMM '06: Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, 2006, pp. 351–362.
- [16] Y. Mao, F. Wang, L. Qiu, S. S. Lam, and J. M. Smith, "S4: Small state and small stretch routing protocol for large wireless sensor networks," in *Proceedings of the 4th USENIX Symposium on Networked System Design and Implementation (NSDI 2007)*, April 2007.
- [17] M. Thorup and U. Zwick, "Compact routing schemes," in *SPAA '01: Proceedings of the thirteenth annual ACM symposium on Parallel algorithms and architectures*, 2001, pp. 1–10.
- [18] D. Peleg, *Distributed Computing: A Locality-Sensitive Approach*. SIAM Monographs on Discrete Mathematics and Applications, 2000.
- [19] L. J. Cowen, "Compact routing with minimum stretch," in *SODA '99: Proceedings of the tenth annual ACM-SIAM symposium on Discrete algorithms*, 1999, pp. 255–260.
- [20] T. Eilam, C. Gavoille, and D. Peleg, "Compact routing schemes with low stretch factor," *J. Algorithms*, vol. 46, no. 2, pp. 97–114, 2003.
- [21] I. Abraham and D. Malkhi, "Name independent routing for growth bounded networks," in *SPAA '05: Proceedings of the seventeenth annual ACM symposium on Parallelism in algorithms and architectures*, 2005, pp. 49–55.
- [22] G. Konjevod, A. W. Richa, and D. Xia, "Optimal-stretch name-independent compact routing in doubling metrics," in *PODC '06: Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, 2006, pp. 198–207.
- [23] H. T.-H. Chan, A. Gupta, B. M. Maggs, and S. Zhou, "On hierarchical routing in doubling metrics," in *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, 2005, pp. 762–771.
- [24] I. Abraham, C. Gavoille, A. V. Goldberg, and D. Malkhi, "Routing in networks with low doubling dimension," in *Proc. of the 26th International Conference on Distributed Computing Systems (ICDCS)*, July 2006.
- [25] R. Sarkar, X. Zhu, and J. Gao, "Spatial distribution in routing table design for sensor networks," <http://www.cs.sunysb.edu/~rik/papers/spatial-routing-journal.pdf>, 2008.
- [26] P. Levis, N. Lee, M. Welsh, and D. Culler, "TOSSIM: accurate and scalable simulation of entire tinyos applications," in *SenSys '03: Proceedings of the 1st international conference on Embedded networked sensor systems*, 2003, pp. 126–137.