Discrete Mobile Centers*

Jie Gao,[1] Leonidas J. Guibas,[1] John Hershberger,[2] Li Zhang,[3] and An Zhu[1]

[1]Department of Computer Science, Stanford University,
Stanford, CA 94305, USA
{jgao,guibas,anzhu}@cs.stanford.edu

[2]Mentor Graphics, 8005 S.W. Boeckman Road,
Wilsonville, OR 97070, USA
john_hershberger@mentor.com

[3]Compaq Systems Research Center, 130 Lytton Avenue,
Palo Alto, CA 94301, USA
l.zhang@compaq.com

Abstract. We propose a new randomized algorithm for maintaining a set of clusters among moving nodes in the plane. Given a specified cluster radius, our algorithm selects and maintains a variable subset of the nodes as cluster centers. This subset has the property that (1) balls of the given radius centered at the chosen nodes cover all the others and (2) the number of centers selected is a constant-factor approximation of the minimum possible. As the nodes move, an event-based kinetic data structure updates the clustering as necessary. This kinetic data structure is shown to be responsive, efficient, local, and compact. The produced cover is also smooth, in the sense that wholesale cluster re-arrangements are avoided. This clustering algorithm is distributed in nature and can enable numerous applications in ad hoc wireless networks, where mobile devices must be interconnected to perform various tasks collaboratively.

1. Introduction

Collaborating mobile devices are of interest in diverse applications, from wireless networking to sensor nets to robot exploration. In these applications there are mobile nodes

that need to communicate as they move so as to accomplish the task at hand. These tasks can vary from establishing an ad-hoc multi-hop network infrastructure that allows point-to-point communication, to aggregating and assimilating data collected by distributed sensors, to mapping an unknown environment collaboratively. A challenge common to all these tasks is that communication is usually accomplished using low-power radio links or other short-range technologies. As a result only nodes sufficiently close to each other can communicate directly and therefore the communication topology of the network is strongly affected by node motion (as well as obstacle interference, etc.). The mobile networking community has been especially active in studying such problems in the context of networking protocols allowing the seamless integration of devices such as PDAs, mobile PCs, phones, pagers, etc., that can be mobile as well as switch off and on at arbitrary times. An example of such an effort is the recent *Bluetooth* specification [12].

A principle that has been discussed a number of times for enabling such collaborative tasks is the organization of the mobile nodes into *clusters* [3], [6], [10], [22]. Clustering allows hierarchical structures to be built on the mobile nodes and enables more efficient use of scarce resources, such as bandwidth and power. For example, if the cluster size corresponds roughly with the direct communication range of the nodes, much simpler protocols can be used for routing and broadcasting within a cluster; furthermore, the same time or frequency division multiplexing can be re-used across non-overlapping clusters. Clustering also allows the health of the network to be monitored and misbehaving nodes to be identified, as some nodes in a cluster can play watchdog roles over other nodes [18].

Motivated by these issues, in this paper we study the problem of maintaining a clustering for a set of n moving points or nodes in the plane. There is, of course, a huge literature on clustering, as the problem in many variations has been studied by several different communities, including operations research, statistics, and computational geometry. In our setting we assume that all the nodes are identical and each can communicate within a region around itself, which we take to be an L_p ball. For most of the paper we focus on a ball in the L_∞ metric, that is an axis-aligned square whose side is of length r, as this makes the analysis the simplest. We call two nodes *visible* to each other if they are within the communication range of each other. We seek a minimal subset of the n nodes, the *centers*, such that every node is visible to at least one of the centers. In the mobile device setting, unlike the general facilities location context, it is appropriate to insist that the centers are located at the nodes themselves, as these are the only active elements in the system; thus we are interested in "discrete center" problems. We survey the literature on the static version of this problem in Section 2. The problem is known to be NP-complete and most of the work has focused on approximation algorithms.

Much less is known, however, about maintaining a clustering on mobile nodes. There have been a few papers in the mobile networking community [3], [6], [10], [22] proposing and simulating a number of distributed algorithms for cluster maintenance, but to our knowledge there has been very little prior work on a theoretical analysis of the problem. In particular, existing algorithms for the static version cannot be adapted to the mobile case efficiently. Many static algorithms utilize space partition methods, i.e., they partition space into smaller sub-regions and solve for each region separately. For instance, one can design a simple constant approximation algorithm by choosing one center out of every pre-fixed grid square of length $r/2$. Algorithms of such flavor totally ignore the underlying topology of the node set and, as a result, suffer from many unnecessary

solution changes during node motion. For example, if nodes are traveling together with the same velocity, then in fact there is no need to change the solution. A "good" algorithm should only undergo solution changes that are necessary. Another desirable property is that the algorithm can be implemented in a distributed manner on nodes with modest capabilities, so as to be useful in the mobile ad hoc network setting. As nodes move around, we need efficient ways of incrementally updating the solution, based on local information as much as possible. We formalize such properties in our study.

In this paper we present a new randomized clustering algorithm that provides a set of centers that is an $O(1)$ approximation to the optimal discrete center solution. Our algorithm uses $O(\log \log n)$ rounds of a "center nomination" procedure in which each node nominates another node within a certain region around itself to be a center; a round of the nomination procedure can be implemented in $O(n \log n)$ time. Furthermore, we show how this approximately optimal clustering can be maintained as the nodes move continuously. The goal here is to exploit the continuity of the motion of the nodes to avoid recomputing and updating the clustering as much as possible. We employ the framework of *Kinetic Data Structures* (KDS) [4], [11] to provide an analysis of our method. For this analysis we assume that nodes follow posted flight plans, though they may change them at any moment by appropriately notifying the data structure. The correctness of the clustering is certified by a set of conditions, or *certificates*, whose predicted failure times are inserted as events into an event queue. At each certificate failure the KDS certification repair mechanism is invoked to repair the certificate set and possibly the clustering as well. We show that the proposed structure is responsive, efficient, local, and compact. Certificate failures and flight-plan updates can be processed in expected time $O(\log^{3.6} n)$ and $O(\log n \log \log n)$, respectively. Under the assumption of pseudo-algebraic motions for the nodes, we show that our structure processes at most $O(n^2 \log \log n)$ events (certificate failures). We also give a construction showing that for any constant $c > 1$, there is a configuration of n points moving linearly on the real line so that *any* c-approximate set of centers must change $\Omega(n^2/c^2)$ times. Thus, even though an approximate clustering is not a canonical structure [1], we can claim efficiency for our method.

To summarize, our clustering algorithm has a number of attractive properties:

- We can show that the clustering produced is an $O(1)$ approximation.
- The clustering generated by the algorithm is *smooth* in the sense that a point's movement causes only local clustering changes. This is in contrast to the optimal clustering solution, which may undergo a complete rearrangement upon small movements of even a single point.
- In the KDS setting the algorithm also supports dynamic insertion and deletion of nodes, with the same update bound as for a certificate failure, in addition to the mentioned properties of our KDS.
- The algorithm can be implemented in a distributed fashion: each node only reasons about the nodes visible to it in order to carry out the clustering decisions. In fact, the algorithm can be implemented without any knowledge of the actual positions of the nodes—only knowledge of distances to a node's visible neighbors are necessary.

Because of these properties, our algorithm has many applications to ad hoc wireless networks. We defer the detailed discussion to Section 5.3.

The remainder of the paper is organized as follows. Section 2 summarizes previous work on discrete centers and related problems. Section 3 introduces the basic algorithm and analyzes the approximation factors for the clusterings it produces. Section 4 describes a hierarchical version of the algorithm and proves the constant approximation bound. Section 5 shows how this clustering can be maintained kinetically under node motion and analyzes the performance of the algorithm in both centralized and distributed settings. Finally Section 6 concludes with some directions for future research.

2.   Previous Work

There is little prior work on this specific mobile clustering problem. The static version of the problem is known to be NP-complete [9] and to admit a PTAS (polynomial time approximation scheme). It is equivalent to finding the minimum dominating set in the intersection graph of unit disks. The dominating set problem is defined as follows. Given a graph $G = (V, E)$, find a minimum size subset V' of vertices, such that every vertex in $V \setminus V'$ is adjacent to some node in V'. For our problem we build a graph G on all the points and create an edge between two points if a disk of size r centered at one point contains the other point. The goal is to find the minimum dominating set in G.

The dominating set problem on general graphs is NP-complete and hard to approximate as well. In fact, no algorithm with approximation factor better than $(1 - \varepsilon) \ln n$ exists unless NP $\subset$ DTIME($|V|^{\log \log |V|}$) [8]. A greedy algorithm can construct a solution of size $k^* \log n$, where k^* is the size of the optimal solution (this follows from a reduction to the set cover problem). For the dominating set in an intersection graph, Hunt et al. [14] gave a PTAS, providing a solution of size no more than $(1 + \varepsilon)k^*$, for any constant $\varepsilon > 0$. The basic idea of the PTAS comes from an algorithm by Hochbaum and Maas [13] for the continuous variant, in which centers can be arbitrary points on the plane. Roughly speaking, the method in [14] divides the space into strips of a certain width, and a sub-problem is formed by grouping several consecutive strips together and proceeding recursively.

The networking community has developed many protocols to deal with changing network topologies. However, no theoretical bounds have been derived for many of these heuristics. We note that our basic algorithm is similar to the Lowest-ID Cluster Algorithm proposed by Ephremides et al. [7]. Experiments show that this scheme works well in practice. A similar idea leads to the Max-Min D-clustering scheme that was proposed by Amis et al. [2]. For the connected dominating set problem, Wu and Li proposed a distributed algorithm that performs badly in the worst case ($O(n)$-approximation) but works well in simulation [26].

3.   Basic Algorithm

Before presenting the algorithms, we first give some formal definitions. A d-cube with size r is a d-dimensional axis-aligned cube with *side length r*. When $d = 1$ or 2, a d-cube is also called an interval or a square, respectively. For two points p and q, p is said to be *r-covered* by or *r-visible* from q if p is inside the cube with size r centered at q. For

a set of n points (nodes) $P = \{p_1, p_2, \ldots, p_n\}$ in the d-dimensional space, a subset of P is called an r-*cover* of P if every point in P is r-covered by some point in the subset. The points in a cover are also called (*discrete*) *centers*. A minimum r-cover of P is an r-cover that uses the minimum number of points. We denote by $\alpha_P(r)$ (or $\alpha(r)$ if P is clear from the context) the number of points in a minimum r-cover of P. An r-cover is called a c-approximate cover of P if it contains at most $c \cdot \alpha_P(r)$ points. When r is not mentioned, we understand it to be 1. In this paper we are interested in computing and maintaining $O(1)$-approximate covers for points moving in the space. For the sake of presentation, we discuss our algorithms for points in one and two dimensions, but our techniques generally can be extended to higher dimensions. In the rest of the paper, log is understood to be $\log_2$, and ln to be $\log_e$, unless otherwise specified in the context. This distinction is important because in a few places, log appears in exponents, and we have to make the base explicit in order to give precise asymptotic bounds.

In the following we first present the algorithms for the static version of the problem and later describe their implementation for moving points.

3.1. *Description of the Basic Algorithm*

The algorithm, which is distributed in nature, is the following: we impose a random numbering (a permutation of $1, 2, \ldots, n$) onto the n points, so that point p_i has an index N_i. In most situations in practice each mobile node is given a unique identifier (UID) at set-up time, and these UIDs can be thought of as providing the random numbering (either directly or via a hash function on the UIDs). Each point p_i nominates the largest indexed point in its visible range to be a center (note that a point can nominate itself if there is no other point with larger index inside its range). All points nominated are the centers in our solution. A cluster is formed by a selected center and all the points that nominated it.

First, we note that randomization is essential for the performance of our scheme. Without randomization, the only approximation bound that holds, even in the one-dimensional case, is the trivial $O(n)$ bound. For example, consider the one-dimensional case in which n points are equally spaced along a unit interval, with their indices increasing monotonically from left to right. Each point in the left half of the set has a different center, which is the rightmost point within distance $\frac{1}{2}$ of it. Thus the number of centers produced by the algorithm is $n/2$, even though the optimal covering uses only a single center.

In the following we are able to show that for *any* configuration, if the ordering is assigned randomly, the basic algorithm yields a sub-linear approximation ($\log n$ in one dimension, and $\sqrt{n}$ in higher dimensions) with high probability.

3.2. *Analysis for the Basic Algorithm*

3.2.1. *Analysis for the One-Dimensional Case.* As a warm-up, we first present the analysis for this algorithm in the one-dimensional case, where points lie along the real line and the unit square corresponds to the unit interval.
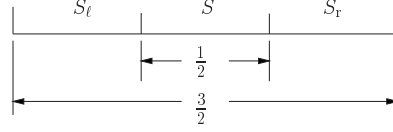
Fig. 1.   Visible range in one dimension.

Lemma 3.1.   *If V' is a subset of the points that are mutually visible to each other, then there is at most one point in V' nominated by points in V'.*

Let the optimal centers be O_i, $i = 1, 2, \ldots, k$. We partition each unit interval U_i centered at O_i into two sub-intervals with O_i as the dividing point. We define the *visible range* of an interval to be all the points on the line that are visible to at least one of the nodes in the interval and call nodes in the visible range the *visible set* for that interval.

Theorem 3.2.   *The basic algorithm has an approximation factor of $4 \ln n + 2$ in expectation.*

Proof.   It suffices to show that, for each sub-interval S, the number of centers nominated by points in S is at most $2 \ln n + 1$. The visible range of S is contained in an interval of size $\frac{3}{2}$ as shown in Fig. 1. We use S_ℓ to denote the portion of the interval to the left of S and S_r for the right portion. Note that the points in S are mutually visible. Lemma 3.1 shows that all the points in S nominate at most one center in S.

Now we calculate the expected number of centers in S_r that are nominated by points in S. Let $x = |S|$ and $y = |S_r|$ be the number of nodes in the respective sub-intervals. Scan all points from left to right in S_r. The ith point in S_r can be nominated by a point in S only if it has the largest index compared with all points to its left in $S \cup S_r$. Therefore, the expected number of centers in S_r is no more than $\sum_{i=1}^{y} (1/(x+i)) < \ln n$. A similar argument works for S_ℓ, and we can conclude that all points in S nominate at most $2 \ln n + 1$ centers.   □

We remark that the approximation bound is asymptotically tight. Consider the following situation in Fig. 2: the unit interval centered at p is divided into two sub-intervals S_ℓ and S_r. S_ℓ contains $\sqrt{n}$ evenly distributed points, each of which can see $\sqrt{n}$ more points in S_r from left to right. In this configuration, with probability $1/2$, the leftmost point q in S_ℓ nominates a point in the first group of $\sqrt{n}$ points in S_r. This is because q sees $2\sqrt{n}$ points ($\sqrt{n}$ in S_ℓ and another $\sqrt{n}$ in S_r). Under a random numbering, the point with the
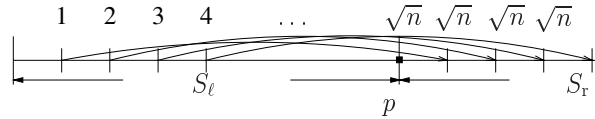

Fig. 2.   Lower bound for the one-dimensional case.

maximum rank falls in S_r with probability $1/2$. In general, a point in the ith group of S_r is nominated by the ith point in S_ℓ with probability $1/(i+1)$. Thus the expected number of centers (in S_r alone) is $\sum_{i=1}^{\sqrt{n}}(1/(i+1)) = \Omega(\log n)$. However, a single center at p covers all the points.

We can also prove that the $O(\log n)$ upper bound holds with high probability. This fact is useful in our hierarchical algorithm, which achieves a constant approximation factor, and in our kinetic maintenance algorithms.

Theorem 3.3. *The probability that the basic algorithm selects more than $ck^* \ln n$ centers is $O(1/n^{\Theta(c^2)})$, where k^* is the optimal number of centers.*

Proof. We divide the optimal intervals in the same way as in the proof of Theorem 3.2. Consider a sub-interval S and its right portion S_r. We look for the fraction of random numberings such that points in S nominate not too many centers in S_r. We sort all points in $S \cup S_r$ according to their coordinates from left to right into a sequence of m points. The sequence of their indices can be viewed as a random permutation on numbers $1, 2, \ldots, m$. Each center in S_r must have a bigger index than all the other points to its left. Thus, to guarantee that points in S nominate no more than s centers in S_r, it suffices to ensure that the total number of left-to-right maximal indices in the sequence is no more than s. The number of permutations with s left-to-right maxima is known as the Stirling number $C(m, s)$, which is asymptotically equal to $m!\, e^{-\theta^2/2}/\sqrt{2\pi}$, for $s = \ln m + \theta\sqrt{\ln m}$, as $m \to \infty$ and $\theta/m \to 0$ [24]. Let x be the random variable of the number of left-to-right maxima in this permutation. Then we have

$$\text{Prob}(x \geq s) = \int_s^\infty P(l)\, dl \leq \int_s^\infty \frac{C(m, l)}{m!}\, dl.$$

If we set $s = c \ln n$, this formula becomes

$$\begin{aligned}
\text{Prob}(x \geq c \ln n) &\leq \int_{(c-1)\sqrt{\ln n}}^\infty \frac{e^{-\theta^2/2}}{\sqrt{2\pi}}\sqrt{\ln m}\, d\theta \\
&\leq \sqrt{\frac{\ln m}{\pi}} \int_{(c-1)(\sqrt{\ln n}/\sqrt{2})}^\infty e^{-x^2}\, dx \\
&\leq \frac{n^{-(c-1)^2/2}}{\sqrt{2\pi}(c-1)} \leq n^{-\Theta(c^2)}.
\end{aligned}$$

For $O(k^*)$ sub-intervals, since each needs to be considered only twice for its left and right points, the probability that there are more than $ck^* \ln n$ centers is less than $\Theta(n)\, n^{-\Theta(c^2)}$, which is $O(n^{-\Theta(c^2)})$. $\qquad\square$

3.2.2. *Analysis for the Two-Dimensional Case.*   Unfortunately Theorem 3.3 does not extend to higher dimensions. We will show that in two (and higher) dimensions, the method above produces a $\Theta(\sqrt{n}\log n)$ approximate cover with high probability. The analysis is similar to the one-dimensional case. Again, we consider the sub-squares with side length $\frac{1}{2}$. For such a square S, suppose that L is the visible range of S. Clearly, L
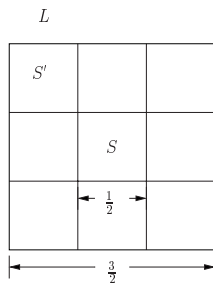
Fig. 3.   Visible range in two dimensions.

is a square of side length $\frac{3}{2}$ and can be partitioned into nine sub-squares where S is the center one (Fig. 3). Now, we have the following lemma:

Lemma 3.4.   *Suppose that* $|L| \leq m$. *Then the number of centers nominated inside* S *is* $O(\sqrt{m})$ *in expectation. Furthermore, the probability that* S *contains more than* $8\sqrt{m} \ln m + 1$ *centers is bounded by* $O(1/m^{\ln m})$.

Proof.   We need to consider only those points inside L. It suffices to bound the number of centers nominated by points in each sub-square S' of L. If $S' = S$, since all the points are mutually visible in S, there can be at most one point nominated. For $S' \neq S$, suppose that $x = |S|$, $y = |S'|$. A point $p \in S$ can be nominated by a point $q \in S'$ if q finds that p has the largest index in its visible range. Since q sees all points in S', p must have rank higher than all the points in S'. Thus, the probability that p can be nominated is at most $1/(1 + y)$. Thus, in expectation, there are at most $x/(1 + y)$ points nominated. On the other hand, since there are only y points in S', there can be at most y centers nominated by points in S'. The expected total number of centers is therefore no more than $\min(y, x/(1 + y)) \leq \sqrt{x + y + 1} - 1 < \sqrt{m}$.

Furthermore, if $y < \sqrt{m} \ln m$, then we know that S' cannot nominate more than $\sqrt{m} \ln m$ points. Otherwise, S' contains $y > \sqrt{m} \ln m$ points. In order to nominate s points in S, S must contain at least s points with higher ranks than all the points in S'. That is, S must contain the s highest ranked points in $S \cup S'$.

The probability for this to happen is

$$\binom{x}{s} \bigg/ \binom{x + y}{s} = \frac{x!\,(x + y - s)!}{(x + y)!\,(x - s)!}$$

$$< \left(\frac{x}{x + y}\right)^s < \left(1 - \frac{y}{m}\right)^s < \left(1 - \frac{\sqrt{m} \ln m}{m}\right)^s.$$

Thus, if $s > \sqrt{m} \ln m$, we have

$$\mathrm{Prob}(x \geq s) < \left(1 - \frac{\ln m}{\sqrt{m}}\right)^{\sqrt{m} \ln m} < \left(\frac{1}{e}\right)^{\ln^2 m} = O\left(\frac{1}{m^{\ln m}}\right).$$

Summing over all the nine sub-squares, we see that the expected number of centers nominated in S is bounded by $O(\sqrt{m})$, and, with high probability, the number of centers nominated is bounded by $O(\sqrt{m} \ln m)$. $\square$
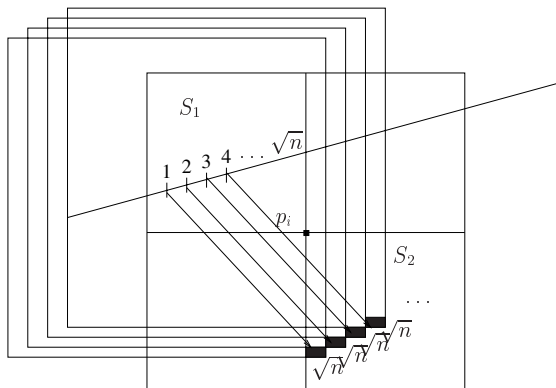
Fig. 4. Lower bound for the two-dimensional case.

By Lemma 3.4, it is easy to obtain

Theorem 3.5. *For points in the plane, the algorithm has an approximation factor of $O(\sqrt{n})$ in expectation. Further, the probability that there are more than $\sqrt{n} \ln n \cdot k$ centers is $O(1/n^{\ln n - 1})$, where k is the optimal number of centers.*

Proof. Consider an optimal covering U_i, $1 \le i \le k$. We partition each U_i in the optimal solution into four quadrant sub-squares and apply Lemma 3.4 to each sub-square. Since there are at most $O(n)$ sub-squares, the high probability result also holds. $\qquad\square$

Again, this bound is asymptotically tight. Consider the configuration in Fig. 4: the upper left sub-square S_1 has $\sqrt{n}$ points, each of which can see a distinct set of $\sqrt{n}$ points in the lower right sub-square S_2. Each point in S_1 will nominate a point in S_2 with probability $1/2$. Thus the expected number of centers in S_2 is $\Omega(\sqrt{n})$. We remark that in this analysis, the use of the unit square and the dimensionality is not essential. It is easy to extend the analysis to any centrally symmetric covering shape in any dimension; the constant factors, however, depend on the covering shape and the dimensionality.

Note also that the worst-case examples in Theorems 3.2 and 3.5 require a significantly non-uniform distribution of the points. The distributions encountered in practice are much less skewed, and the basic algorithm returns much better results, as experiments show [10].

4.   Hierarchical Algorithms for Clustering

The basic algorithm is simple, but it achieves only an $O(\sqrt{n})$ approximation for points in the plane. To obtain a constant-factor approximation, we use a hierarchical algorithm that proceeds in a number of rounds. At each round we apply the basic algorithm to the *centers* produced by the previous round, using a larger covering cube. Suppose that $\delta_i = 2^i/\log n$, for $i > 0$. Initially, set P_0 to be P, the input set of points. At the ith step, for $1 \le i < \log \log n$, we apply the basic algorithm using squares with side length δ_i to the

set P_{i-1} and let P_i be the output. The final output of the algorithm is $P' = P_{\log\log n - 1}$. (To make our analysis fully rigorous, we would need to use $\lfloor \log n \rfloor$ and $\lfloor \log\log n \rfloor$ instead of $\log n$ and $\log\log n$; however, in the interest of readability, we omit the floor functions from this paper.) We claim that:

Lemma 4.1.   *P' is a 1-cover of P.*

Proof.   We actually prove a stronger statement: P_i is a $(2^{i+1}/\log n)$-cover of P.

We proceed by induction. The assertion is clearly true when $i = 0$. Suppose that it is true for i, i.e., every point $p \in P$ can be covered by a size $2^{i+1}/\log n$ square centered at a point $q \in P_i$. If q is also in P_{i+1}, then p is covered. Otherwise, there must be a q' so that q nominates q' at the $(i+1)$th step. Thus, p is covered by q' with a square with side length $2^{i+1}/\log n + \delta_{i+1} = 2^{i+2}/\log n$. That is, P_{i+1} is a $(2^{i+2}/\log n)$-cover of P. $\square$

In the following we bound the approximation factor for P'. To explain the intuition, we first consider the situation when P admits a single center, i.e., there is a unit square that covers all the points in P. Recall that $\alpha(x)$ denotes the number of centers of an optimal covering of P by using squares with side length x. First, we observe that

Lemma 4.2. $\alpha(x) \leq 4/x^2$.

Proof.   We uniformly divide the unit square into $4/x^2$ small squares of size $x/2$. We then pick one point from each non-empty small square, which gives an x-cover with at most $4/x^2$ centers. $\square$

According to Theorem 3.5, the expected size of P_{i+1} is at most $c\sqrt{|P_i|}\alpha(\delta_{i+1})$, for some constant $c > 0$. Denote by n_i the size of P_i. We have the following recursive relation:

$$n_0 = n, \qquad n_{i+1} \leq c\sqrt{n_i}\alpha(\delta_{i+1}) \leq c\sqrt{n_i}\frac{4\log^2 n}{2^{2i+2}}.$$

By induction, it is easy to verify that

$$n_i \leq \frac{(c^2 \log^4 n)n^{1/2^i}}{4^{2i-4}}.$$

Thus $|P'| = n_{\log\log n - 1} \leq c^2 2^{14} = O(1)$.

We cannot apply this argument directly to the general case because $\alpha(x)$ can be as large as $\Theta(n)$. In order to establish a similar recursive relation, we consider points restricted to lie in squares of a certain size. For any square S with side length δ_i, let $m_i(S)$ denote the expected value of $|P_i \cap S|$. Further, let m_i denote the maximum of $m_i(S)$ over all the squares S with size δ_i. We then have the following relation between m_i's.

Lemma 4.3. $m_{i+1} \leq c\sqrt{m_i}$, *for some constant $c > 0$ and any $0 \leq i < \log\log n - 1$.*

Proof. Consider a square S of side length δ_{i+1}. Its visible region L, with respect to side length δ_{i+1}, is a square with side length $2\delta_{i+1} = 4\delta_i$. Thus L can be covered by $4^2 = 16$ squares with side length δ_i. That is, $|P_i \cap L| \leq 16m_i$ in expectation. By Lemma 3.4, we know that the expected number of points inside S that survive after the $(i+1)$th step of the algorithm is $O(\sqrt{|P_i \cap L|}) = O(\sqrt{m_i})$. Thus, we have $m_{i+1} \leq c\sqrt{m_i}$, for some constant $c > 0$. $\qquad\square$

Now, we can prove that

Theorem 4.4. *P' is a constant approximation to the optimal covering of P with unit squares in expectation.*

Proof. Clearly, $m_0 \leq n$. Solving the recursive relation in Lemma 4.3, we find that $m_i \leq O(c^2 n^{1/2^i})$. Setting $i = \log\log n - 1$, we have $m_{\log\log n - 1} = O(1)$, i.e., for a square S with side length $\frac{1}{2}$, the expected number of points of P' inside S is $O(1)$.

Now, suppose that an optimal cover uses k unit squares. We can then cover all the points by $4k$ squares with side length $\frac{1}{2}$. Since each of these squares contains $O(1)$ points in P' in expectation, the total number of points in P' is bounded by $O(k)$. $\qquad\square$

In addition, we have:

Corollary 4.5. *For a modified version of the hierarchical algorithm, i.e., we stop the center election process as soon as m_i drops below $\log n$, then the number of centers generated is an $O(\log^3 n)$ approximation to the optimal cover of P, with probability $1 - o(1)$.*

Proof. From Lemma 3.4, at round i, $m_{i+1} \leq 8\sqrt{m_i} \ln m_i + 1$, with probability $1 - O(1/m_i^{\ln m_i})$. So $m_{i+1} \leq c' m_i^{1/(2-\delta)}$ for some constant c' and $0 < \delta < 1$. In this corollary we change the base of the log function from 2 to $2 - \delta$, so $m_i \leq c'^{(2-\delta)/(1-\delta)} n^{1/(2-\delta)^i}$. To obtain an $O(\log^3 n)$ approximation, we could stop the center election process as soon as m_i drops below $\log n$. For a square of side length 1, the total number of centers inside is $O(\log^3 n)$, because the size of squares at level i is at least $1/\log n$. We achieve this bound with probability bigger than $(1 - O(1/(\log n)^{\ln\log n}))^{\log\log n - 1} \geq 1 - o(1)$. $\qquad\square$

5. Kinetic Discrete Clustering

To kinetize the algorithm, we place a half-size square centered over each point. If two such squares intersect, we know the corresponding points are mutually visible. In this section when we say "squares," we refer to these half-size squares.

5.1. *Standard KDS Implementation*

The intersection relation between two squares can change only at discrete times. If two squares of the same size intersect with each other, one square must have a corner inside the other square. Therefore, we can maintain the left and right extrema of squares in x-sorted order and the top and bottom extrema of squares in y-sorted order. The certificates of the KDS are the ordering certificates for the x- and y-sorted lists of square extrema. We maintain the lists containing the extrema of active squares for each level of the hierarchy. An event is a certificate failure. When an event happens, we first check whether it is a "real" event, i.e., whether it causes two squares to start/stop intersecting. When two squares S_1, S_2 start intersecting, we will need to check the square with the lower rank, say S_1, to see if its nomination has a lower rank than S_2. If so, we need to change S_1 to point to S_2. If S_1, S_2 stop intersecting, we need to check if S_1 nominated S_2. If so, we need to find another overlapping square with the highest rank. To answer this query efficiently, we maintain a standard range search tree [19] for the n points. For our purpose, the internal nodes of the second-level binary trees in the range tree are augmented with the maximum index of the points stored at descendants of each node. This will let us find the points within a query square that are larger than some query index in $O(\log^2 n)$ time. To maintain the range search trees kinetically, we keep sorted lists of the x- and y-coordinates of the points themselves, in addition to the sorted lists containing the extrema of the squares on each level. A range tree can be updated by deleting a point and re-inserting it in the right place [5].

For the hierarchical algorithm, we need to maintain these structures for each level. In addition, we also need to insert or delete a point to or from a level, as a consequence of an event happening at a lower level. This requires the sorted lists and range search trees used in the basic algorithm above to be dynamic. These requirements can be easily satisfied by maintaining balanced binary search trees and dynamic range search trees.

5.2. *Kinetic Properties*

This KDS has most of the properties of a good KDS [4]. We assume the points have bounded-degree algebraic motion in the following arguments.

To analyze the efficiency, i.e., the number of events, of our algorithms, we first give some lower bound constructions.

Lemma 5.1. *The number of changes of the optimal cover for n points in motion is* $\Theta(n^3)$ *in the worst case.*

Proof. Consider the graph G in which each vertex represents a point and each edge joins a visible pair of points. Clearly, the minimum discrete covering of the points is exactly the same as the minimum dominating set of the graph. The graph can change only when two points become or cease to be visible to each other. For bounded degree algebraic motions, this can happen only $O(n^2)$ times. For each such event, the change to
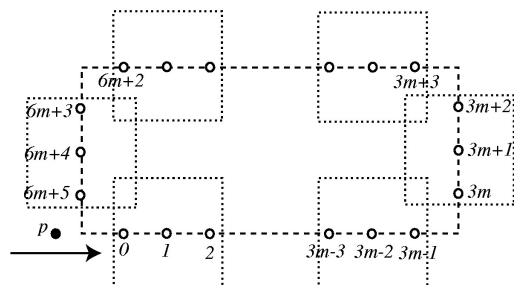
Fig. 5.   Lower bound for optimal coverings.

the minimum covering is at most $O(n)$. Thus, in the worst case, the number of changes is $O(n^3)$.

We now construct an example in which any optimal cover must change $\Theta(n^3)$ times. The construction uses $6m + 6$ static points along the perimeter of a rectangle $[0, R] \times [0, 1.6]$, where $R = 0.4(3m + 1)$. The left and right sides of the rectangle have three points apiece, located at $(0, 0.4i)$ and $(R, 0.4i)$ for $i = 1, 2, 3$. The top and bottom sides of the rectangle have $3m$ points apiece, located at $(0.4i, 0)$ and $(0.4i, 1.6)$, for $i = 1, \ldots, 3m$. We label the points counterclockwise from 0 to $6m + 5$ as shown in Fig. 5. In this configuration, each point i can see the points $i - 1, i + 1$ (modulo $6m + 6$) and no other points. Thus, an optimal cover contains $2m + 2$ centers and can be realized in one of three ways by using points $3i$, $3i + 1$, or $3i + 2$, respectively, which we call types 0, 1, and 2, respectively. Clearly, to change from one type to another, we need to make $\Theta(m)$ changes to the cover.

Now consider what happens when a single point p moves linearly along the x-axis. For any i, suppose that q_j is the middle point between the pair $3i + j, 3i + j + 1$, for $0 \le j \le 2$. When p is located at q_j, the only points p can see are $3i + j$ and $3i + j + 1$. Thus, an optimal cover has to use either $3i + j$ or $3i + j + 1$ as a center. In other words, an optimal cover has to be of type j or $j + 1$. It is easily verified that when p moves from q_0 to q_2, an optimal cover has to change its type. Therefore, an optimal cover undergoes $\Theta(m)$ changes when p moves from q_0 to q_2. When p moves from $(0, 0)$ to $(R, 0)$, the number of changes is $\Theta(m^2)$. We repeat this procedure by sending m points along the x-axis, passing through the interval $[0, R]$ one at a time. This causes a total of $\Theta(m^3)$ changes to optimal covers. The total number of points is $n = 7m + 6$, so the total number of center changes is $\Theta(n^3)$.   □

While the optimal cover in this construction changes $\Omega(n^3)$ times, a 2-approximate cover does not change at all—we can simply use an optimal cover for the static points and assign each moving point to be a center. However, in the following, we will show that for any constant c, there is a set of moving points that forces any c-approximate cover to change $\Omega(n^2/c^2)$ times.

Theorem 5.2.   *For any constant $c > 1$, there exists a configuration of n points moving linearly on the real line so that any c-approximate cover undergoes $\Omega(n^2/c^2)$ changes.*
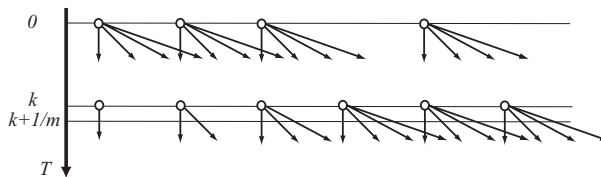
Fig. 6.   Lower bound approximate coverings.

Proof.   In the following we assume that c is an integer and $n = 2cm$, where $m > 2c$ is an integer. We group n points into m groups, each containing $2c$ points. We label each point by (i, j) where $0 \leq i < m$ is the group number, and $0 \leq j < 2c$ is the numbering within each group. Initially, all the points in the ith group are located at $i \cdot 2m$, and the speed of the point (i, j) is $j \cdot 2m$. To summarize, we consider points $p(i, j, t)$ defined as $p(i, j, t) = (i + jt) \cdot 2m$, for $0 \leq i < m$, $0 \leq j < 2c$, and $t \geq 0$.

Whenever $t = k + 1/m$, for some integer $k < m$, $p(i, j, t) = (i + jk + j/m) \cdot 2m = 2(i + jk)m + 2j$. For any two distinct points (i, j) and (i', j'), if $i + jk \neq i' + j'k$, then $|p(i, j, t) - p(i', j', t)| > 2m - 4c \geq 2$; if $i + jk = i' + j'k$, since (i, j) and (i', j') are distinct, $j' \neq j$ and $|p(i, j, t) - p(i', j', t)| \geq 2$. Thus, at time t, no two points are within distance 1. In other words, any covering has to have n centers (Fig. 6).

On the other hand, at time $t = k$ for an integer $k < m$, since $p(i, j, k) = (i + jk) \cdot 2m$ where $0 \leq i < m$, $0 \leq j < 2c$, and $k < m$, each point has position $2sm$ for some $0 \leq s < m + 2ck$. That is, at $t = k$, the minimum covering has at most $m + 2ck$ centers (Fig. 6). Thus, a c-approximate cover may have at most $c(m + 2ck)$ centers. Therefore, between times k and $k + 1/m$, there are at least $n - c(m + 2ck) = n/2 - 2c^2 k$ changes to any c-approximate covering. In total, for all $0 \leq t < K$, the number of changes is at least $\sum_{0 \leq k < K} (n/2 - 2c^2 k) > Kn/2 - c^2 K^2$. Setting $K = n/4c^2 < m$, we have established that the total number of changes is $\Omega(n^2/c^2)$.   □

Lemma 5.3.   *The number of events in our basic algorithm is $O(n^2)$.*

Proof.   An event is the failure of an ordering certificate in an x- or y-sorted list of square side coordinates or point coordinates. Since the points have bounded-degree algebraic motion, each pair of points can cause $O(1)$ certificate failures.   □

Theorem 5.4.   *The number of events processed by our hierarchical KDS is at most $O(n^2 \log \log n)$, and hence the KDS is efficient.*

Proof.   We maintain x- and y-ordering certificates on each of $\log \log n$ levels. As in Lemma 5.3, each pair of points can cause $O(1)$ certificate failures on each level. In addition, in the hierarchical KDS, we need to consider the events for maintaining the range search tree. Those events can happen when two points swap their x- or y-ordering. Such an exchange requires possible updates of the range trees on all levels where the exchanging pair is present. Again, there are $O(n^2)$ exchange events at each level.   □

We now proceed to examine the cost of processing the kinetic events.

Theorem 5.5. *The expected update cost for one event is $O(\log^{3.6} n)$. Hence the KDS is responsive in an expected sense.*

Proof.   When two points exchange in x- or y-order, only the relevant range search trees need to be updated. We need $O(\log^2 n)$ time to update each of $\log \log n$ range trees.

When two points p_i, p_j start/stop being mutually visible at any level of the hierarchy, we can update the centers involved with p_i, p_j in $O(\log^2 n)$ time, since we may need to search for a replacement center in the range tree. One new center may appear and one old center may disappear; these changes bubble up the hierarchy.

On hierarchy levels above the bottom, we divide the changes into two kinds, those caused by the motion of the points in that level and those caused by insertion or deletion of points bubbled up from lower levels. The number of changes of the first kind per event is a constant.

Let us consider the insertion of point p. The only points that may change their centers are those in p's visible range S. We divide S into four quadrants S_i, each with k_i ($i = 1, 2, 3, 4$) points. If there is some point in S_i that nominates p to be its center, the index of p must be bigger than the indices of all the k_i points. The probability of this occurring is $1/(k_i + 1)$. Therefore, the expected number of point-center changes caused by the appearance of p is at most

$$\frac{k_1}{k_1 + 1} + \frac{k_2}{k_2 + 1} + \frac{k_3}{k_3 + 1} + \frac{k_4}{k_4 + 1} + 1 \le 5.$$

Assuming that p becomes a center, how many centers does it replace? For a given quadrant S_i, suppose the number of centers its points nominate is $m_i \le k_i$. At most one of these centers is inside S_i. If m' points are outside S_i, the probability that p replaces j of them is at most $1/(m' + 1)$. Hence the expected number of centers replaced in a single quadrant is upper bounded by either

$$\frac{1}{k_i + 1}\left(1 + \frac{1 + \cdots + m_i - 1}{m_i}\right) = \frac{m_i + 1}{2(k_i + 1)} \le \frac{1}{2},$$

if one of the centers is inside S_i, or by

$$\frac{1}{k_i + 1}\left(\frac{1 + \cdots + m_i}{m_i + 1}\right) = \frac{m_i}{2(k_i + 1)} \le \frac{1}{2}$$

if none of the centers is inside S_i. Each replaced center may stop being a center at this level of the hierarchy if it is nominated by no points outside S. Thus the expected number of centers created/destroyed in this level (inserted/deleted at higher levels) due to the appearance of p is at most $4 \times \frac{1}{2} + 1 = 3$.

We can make a similar argument for the disappearance of a point. So the expected total number of point-center changes at all levels of the hierarchy per event is at most

$$5 \times (3^{\log \log n} + 3^{\log \log n - 1} + \cdots + 1)$$

which is $O(3^{\log \log n}) = O(\log^{1.6} n)$.

Since insertion or deletion in a range search tree costs $O(\log^2 n)$, the total expected update cost is $O(\log^{3.6} n)$. □

Theorem 5.6. *The KDS uses $O(n \log n \log \log n)$ storage. Each point participates in at most $O(\log \log n)$ ordering certificates. Therefore, the KDS is compact and local.*

Proof. Range trees take $O(n \log n)$ space per level. All other data structures use less space. Each point participates in at most $O(1)$ ordering certificates in each level. □

5.3. *Distributed Implementation*

The KDS implementation requires a central node to collect all the information and perform the hierarchical clustering algorithm. For wireless mobile ad hoc networks [23], all the wireless nodes are homogeneous and a central node is not available most of the time. Furthermore, the cost of communicating all the data to such a node can be prohibitive. Our hierarchical algorithm can be implemented in a distributed manner, making it appropriate for mobile networking scenarios. Specifically, the nodes in ad hoc wireless networks use omnidirectional antennas with a broadcast nature, i.e., a single transmission can be received by all the nodes within a certain neighborhood. In the most popular power-attenuation model [20], a node can send out some messages whose signal power drops as $1/r^\alpha$ where r is the distance from the transmitter antenna and α is a constant between 2 and 4. We assume each node can adjust its transmitting power so that only the nodes within a certain range can receive the messages.

To implement the hierarchical clustering algorithm, we first describe how to obtain range information about internode distances. Each node broadcasts a "who is there" message and waits for replies. Each node that hears the request responds. The hierarchy can be implemented by having nodes broadcast with different power for each level or by other local positioning mechanisms. When the nodes move around, each node broadcasts these "Hello" beacons periodically, with a time interval dependent on its moving speed. Therefore, each point keeps track of its neighborhood within different size ranges. When a neighbor enters or leaves any of the $\log \log n$ ranges of a node, each node involved checks whether it needs to update its center. When it nominates a center that is not nominated by any other node, the center will also be added to a higher level and may cause updates in that level. If a node ceases to be pointed to by any node, then it also has to be deleted from higher levels in the hierarchy. Clearly, all of these operations can be done locally without centralized control.

Notice that by the power-attenuation model, the energy consumed at each node is kept low, since each node transmits only within a small neighborhood. We emphasize here that range information (inter-node distances) is sufficient for each node to select a center for each level. This information can often be obtained using the node radios themselves or acoustic range-finding ultrasonic devices [25]. No global positioning information is needed to implement our clustering algorithm.

This contrasts with many algorithms proposed for mobile ad hoc networks [16], [15], [17], which require that the exact location of each wireless node be known. Obtaining

global position information is expensive—a GPS receiver per node is a costly addition. Providing only a modest number of nodes with GPS and computing the positions of the others by multi-lateration techniques [21] is possible for static nodes, but quickly breaks down when most nodes are moving, as the convergence of the method is slow. A final drawback of GPS is that it does not work for indoor situations.

In the distributed implementation described above, the total storage needed is $O(sn)$, where s is the maximum number of nodes inside a node's range. In the worst case, this can be $\Theta(n^2)$, but in practice, s is often small. Furthermore, we can restrict the storage of each node to be n^ε, where $0 < \varepsilon < 1$, and still get a constant approximation. If we let each node keep up to $O(n^\varepsilon)$ neighbors and select a center among them, then we have the following:

Lemma 5.7. *In two-dimensional space, the number of centers nominated inside a unit-size square S by the space-restricted one-level algorithm is $O(n^{\max(1-\varepsilon,1/2)})$ in expectation.*

Proof. We use the same notation as Lemma 3.4. Suppose L is the visible range of S. L can be divided into nine sub-squares of size $\frac{1}{2}$. Consider any such sub-square S', and suppose $m = |L|$, $x = |S|$, $y = |S'|$. For points in S' such that the number of neighbors is $\le n^\varepsilon$, from Lemma 3.4, the expected number of centers in S that are nominated by those nodes in S' is bounded by $\sqrt{m} \le \sqrt{n}$. The rest of the points in S' store only n^ε neighbors. A point p in S can be nominated by those points only if p has the largest index among n^ε points. So the probability is at most $1/n^\varepsilon$. The expected number of centers in S nominated by points in S' is no more than $x/n^\varepsilon \le n^{1-\varepsilon}$. $\square$

Let c_ε denote $\max(1 - \varepsilon, \frac{1}{2})$.

Theorem 5.8. *The expected number of centers generated by the space-restricted hierarchical algorithm is a constant-factor approximation to the optimum.*

Proof. Recall the notations in Section 4. From the previous lemma we established the recurrence $m'_{i+1} \le cm_i^{c_\varepsilon}$. So $m_i \le c^{1/(1-c_\varepsilon)}n^{c_\varepsilon^i}$. We simply change the base of the log function from 2 to $1/c_\varepsilon$ in the proof of Theorem 4.4. After $\log \log n - 1$ rounds, we again obtain a constant approximation. $\square$

6. Summary and Future Work

Our randomized hierarchical algorithm can be extended to the case in which the ranges are any congruent convex shape and to higher dimensions. Most of the analysis for the two-dimensional case works for any dimension d, except that the constant approximation factor depends exponentially on d. Our algorithms also support efficient insertion or deletion of nodes.

This work also raises several open problems. In the standard KDS setting, where we have nearly linear space, our center updating algorithm exploits the fact that the ranges

are aligned congruent squares. Can we find a similar algorithm in a standard KDS setting with congruent disk ranges instead? Finally, our algorithm is randomized and it would be interesting to find a deterministic algorithm for the mobile centers problem.

We note that our algorithm clusters are based solely on the positions of the mobile nodes. It would be interesting to develop clustering strategies that utilize additional information about the node motions, say both position and velocity. Such clusterings may be far more stable under motion, although they may require more clusters. In fact, a trade-off between the quality and stability of a clustering needs to be investigated. Besides clustering, numerous other problems for ad hoc networks can be studied in the same style as the clustering problem, including network connectivity, route maintenance, node misbehavior detection, etc. We should also note that the constant approximation ratio in our analysis is quite large. That is partly because our analysis applies in the worst case. In practice, we can expect much better performance.

We believe that kinetic clustering is a fundamental problem for the organization of mobile devices and deserves further study. Motion models and quality measures for different application areas need to be studied and developed further.

Acknowledgments

References

1. P. K. Agarwal, J. Basch, M. de Berg, L. J. Guibas, and J. Hershberger. Lower bounds for kinetic planar subdivisions. In *Proceedings of the* 15*th ACM Symposium on Computational Geometry*, pages 247–254, 1999.
2. A. D. Amis, R. Prakash, T. H. P. Vuong, and D. T. Huynh. Max-Min D-cluster formation in wireless ad hoc networks. In *Proceedings of the* 19*th IEEE INFOCOM*, pages 32–41, March 1999.
3. S. Basagni. Distributed clustering for ad hoc networks. In *Proceedings of the '99 International Symposium on Parallel Architectures*, *Algorithms*, *and Networks* (*I-SPAN '99*), pages 310–315, June 1999.
4. J. Basch, L. Guibas, and J. Hershberger. Data structures for mobile data. *Journal of Algebra*, 31(1):1–28, 1999.
5. J. Basch, L. J. Guibas, and L. Zhang. Proximity problems on moving points. In *Proceedings of the* 13*th Annual ACM Symposium on Computational Geometry*, pages 344–351, 1997.
6. C. Chiang, H. Wu, W. Liu, and M. Gerla. Routing in clustered multihop, mobile wireless networks with fading channel. In *Proceedings of IEEE SICON '97*, pages 197–211, April 1997.
7. A. Ephremides, J. E. Wieselthier, and D. J. Baker. A design concept for reliable mobile radio networks with frequency hopping signaling. *Proceedings of IEEE*, 75(1):56–73, Jan. 1987.
8. U. Feige. A threshold of $\ln n$ for approximating set cover. In *Proceedings of the ACM Symposium on Theory of Computing*, pages 634–652, 1996.
9. R. J. Fowler, M. S. Paterson, and S. L. Tanimoto. Optimal packing and covering in the plane are NP-complete. *Information Processing Letters*, 12(3):133–137, 1981.
10. M. Gerla and J. Tsai. Multicluster, mobile, multimedia radio network. *ACM–Baltzer Journal of Wireless Networks*, 1(3):255–265, 1995.
11. L. J. Guibas. Kinetic data structures—a state of the art report. In P. K. Agarwal, L. E. Kavraki, and M. Mason, editors, *Proceedings of the Workshop on Algorithmic Foundations of Robotics*, pages 191–209. Peters, Wellesley, MA, 1998.

12. J. Haartsen, M. Naghshineh, J. Inouye, O. Joeressen, and W. Allen. Bluetooth: vision, goals, and architecture. *Mobile Computing and Communications Review*, 2(4):38–45, Oct. 1998.

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

14. H. B. Hunt, H. Marathe, V. Radhakrishnan, S. Ravi, D. Rosenkrantz, and R. Stearns. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *Journal of Algorithms*, 26(2):238–274, 1998.

15. 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.

16. B. Karp and H. Kung. GPSR: greedy perimeter stateless routing for wireless networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking* (*MobiCom*), pages 243–254, 2000.

17. W.-H. Liao, J.-P. Sheu, and Y.-C. Tseng. GRID: a fully location-aware routing protocol for mobile ad hoc networks. *Telecommunication Systems*, 18(1–3):37–60, 2001.

18. S. Marti, T. J. Giuli, K. Lai, and M. Baker. Mitigating routing misbehavior in mobile ad hoc networks. In *Proceedings of the* 6*th Annual International Conference on Mobile Computing and Networking*, pages 255–265, 2000.

19. F. P. Preparata and M. I. Shamos. *Computational Geometry*: *An Introduction*. Springer-Verlag, New York, 1985.

20. T. Rappaport. *Wireless Communications*: *Principles and Practice*. Prentice-Hall, Englewood Cliffs, NJ, 1996.

21. A. Savvides, C.-C. Han, and M. B. Strivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of the* 7*th ACM/IEEE International Conference on Mobile Computing and Networking* (*MobiCom*), pages 166–179, 2001.

22. J. Sharony. An architecture for mobile radio networks with dynamically changing topology using virtual subnets. *ACM–Baltzer Mobile Networks and Applications Journal*, 1(1):75–86, 1996.

23. C.-K. Toh. *Ad Hoc Mobile Wireless Networks*: *Protocols and Systems*. Prentice-Hall, Englewood Cliffs, NJ, 2002.

24. J. H. van Lint and R. M. Wilson. *A Course in Combinatorics*. Cambridge University Press, Cambridge, 1992.

25. A. Ward, A. Jones, and A. Hopper. A new location technique for the active office. *IEEE Personnel Communications*, 4(5):42–47, Oct. 1997.

26. J. Wu and H. Li. On calculating connected dominating set for efficient routing in ad hoc wireless networks. In *Proceedings of the* 3*rd International Workshop on Discrete Algorithms and Methods for Mobile Computing & Communications*, pages 7–14, 1999.