Construction and Maintenance of Wireless Mobile Backbone Networks

Anand Srinivas, Gil Zussman, and Eytan Modiano

***Abstract*—We study a novel hierarchical wireless networking approach in which some of the nodes are more capable than others. In such networks, the more capable nodes can serve as Mobile Backbone Nodes and provide a backbone over which end-to-end communication can take place. Our approach consists of controlling the mobility of the Backbone Nodes in order to maintain connectivity. We formulate the problem of minimizing the number of backbone nodes and refer to it as the Connected Disk Cover (CDC) problem. We show that it can be decomposed into the Geometric Disk Cover (GDC) problem and the Steiner Tree Problem with Minimum Number of Steiner Points (STP-MSP). We prove that if these subproblems are solved separately by γ- and δ-approximation algorithms, the approximation ratio of the joint solution is $\gamma+\delta$. Then, we focus on the two subproblems and present a number of distributed approximation algorithms that maintain a solution to the GDC problem under mobility. A new approach to the solution of the STP-MSP is also described. We show that this approach can be extended in order to obtain a joint approximate solution to the CDC problem. Finally, we evaluate the performance of the algorithms via simulation and show that the proposed GDC algorithms perform very well under mobility and that the new approach for the joint solution can significantly reduce the number of Mobile Backbone Nodes.**

***Index Terms*—Wireless networks, Controlled mobility, Distributed algorithms, Approximation algorithms, Disk cover.**

I. Introduction

WIRELESS Sensor Networks (WSNs) and Mobile Ad Hoc Networks (MANETs) can operate without any physical infrastructure (e.g. base stations). Yet, it has been shown that it is sometimes desirable to construct a *virtual backbone* on which most of the multi-hop traffic will be routed [4]. If all nodes have similar communication capabilities and similar limited energy resources, the virtual backbone may pose several challenges. For example, bottleneck formation along the backbone may affect the available bandwidth and the lifetime of the backbone nodes. In addition, the virtual backbone cannot deal with network partitions resulting from the spatial distribution and mobility of the nodes.

Alternatively, if some of the nodes are more capable than others, these nodes can be dedicated to providing a backbone over which reliable end-to-end communication can take place. A novel hierarchical approach for a *Mobile Backbone Network* operating in such a way was recently proposed and studied by Rubin et al. (see [24] and references therein) and by Gerla et al. (e.g. [10],[29]). In this paper, we develop and analyze novel algorithms for the construction and maintenance (under node mobility) of a Mobile Backbone Network. Our approach is somewhat different from the previous works, since we focus on *controlling the mobility* of the more capable nodes in order to maintain network connectivity and to provide a backbone for reliable communication.

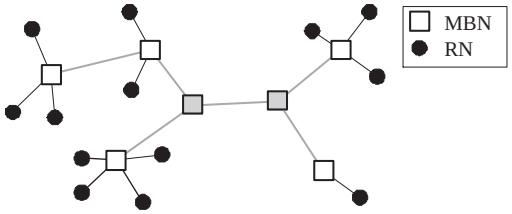

Fig. 1. A Mobile Backbone Network in which every Regular Node (RN) can directly communicate with at least one Mobile Backbone Node (MBN). All communication is routed through a connected network formed by the MBNs

A Mobile Backbone Network is composed of two types of nodes. The first type includes static or mobile nodes (e.g. sensors or MANET nodes) with limited capabilities. We refer to them as *Regular Nodes* (RNs). The second type includes mobile nodes with superior communication, mobility, and computation capabilities as well as greater energy resources (e.g. Unmanned-Aerial-Vehicles). We refer to them as *Mobile Backbone Nodes* (MBNs). The main purpose of the MBNs is to provide a mobile infrastructure facilitating network-wide communication. We specifically focus on minimizing the number of MBNs needed for connectivity. Yet, the construction of a Mobile Backbone Network can improve other aspects of the network performance, including node lifetime and Quality of Service as well as network reliability and survivability.

Fig. 1 illustrates an example of the architecture of a Mobile Backbone Network. The set of MBNs has to be placed such that (i) every RN can directly communicate with at least one MBN, and (ii) the network formed by the MBNs is connected. We assume a *disk* connectivity model, whereby two nodes can communicate if and only if they are within a certain communication range. We also assume that the communication range of the MBNs is significantly larger than the communication range of the RNs.

We term the problem of placing the *minimum* number of MBNs such that both of the above conditions are satisfied as the *Connected Disk Cover (CDC)* problem. While related problems have been studied in the past [2],[4],[13],[15],[27] (see Section II for more details), this paper is one of the first attempts to deal with the CDC problem.

The authors are with the Communications and Networking Research Group, MIT, Cambridge, MA 02139 (e-mail: {anand3, gilz, modiano}@mit.edu).

This research was supported by NSF grant CCR-0325401, by ONR grant N000140610064, and by a grant from Draper Laboratory. It was also supported by a Marie Curie International Fellowship within the 6th European Community Framework Programme.

A preliminary version of this work appeared in ACM MOBIHOC'06, Florence, May, 2006

Our first approach is based on a framework that *decomposes* the CDC problem into two subproblems. We view the CDC problem as a two-tiered problem. In the first phase, the minimum number of MBNs such that all RNs are *covered* (i.e. all RNs can communicate with at least one MBN) is placed. We refer to these MBNs as *Cover MBNs* and denote them in Fig. 1 by white squares. In the second phase, the minimum number of MBNs such that the MBNs' network is connected is placed. We refer to them as *Relay MBNs* and denote them in Fig. 1 by gray squares. In the first phase, the Geometric Disk Cover (GDC) problem [15] has to be solved, while in the second phase, a Steiner Tree Problem with Minimum Number of Steiner Points (STP-MSP) [19] has to be solved. We show that if these subproblems are solved separately by γ- and δ-approximation algorithms[1], the approximation ratio of the joint solution is $\gamma+\delta$.

We then focus on the Geometric Disk Cover (GDC) problem. In the context of static points (i.e. RNs), this problem has been extensively studied in the past (see Section II). However, much of the previous work is either (i) centralized in nature, (ii) too impractical to implement (in terms of running time), or (iii) has poor average or worst-case performance. Recently, a few attempts to deal with related problems under node mobility have been made [6],[13],[16].

We attempt to develop algorithms that do not fall in any of the categories above. Thus, we develop a number of practically implementable *distributed algorithms* for covering mobile RNs by MBNs. We assume that all nodes can detect their position via GPS or a localization mechanism. This assumption allows us to take advantage of location information in designing distributed algorithms. We obtain the worst case approximation ratios of the developed algorithms and the average case approximation ratios for two of the algorithms. Finally, we evaluate the performance of the algorithms via simulation, and discuss the tradeoffs between the complexities and approximation ratios.

Regarding the STP-MSP, [19] and [2] propose 3- and 4-approximation algorithms based on finding a Minimum Spanning Tree (MST). However, when applied to the STP-MSP, such *MST-based* algorithms may overlook efficient solutions. We present a *Discretization Approach* that can potentially provide improved solutions. In certain practical instances the approach can yield a 2 approximate solution for the STP-MSP.

We extend the Discretization Approach and show that it can obtain a solution to the *joint CDC* problem in a centralized manner. Even for the CDC problem, using this approach enables a *2-approximation* for specific instances. Due to the continuous nature of the CDC problem, methods such as integer programming cannot yield an optimal solution. Thus, for specific instances this approach provides the lowest known approximation ratio. It is shown via simulation that this is also the case in practical scenarios.

To conclude, our first main contribution is a decomposition result regarding the CDC problem. Other major contributions are the development and analysis of *distributed approximation algorithms* for the GDC problem in a mobile environment, as well as the design of a *novel Discretization Approach* for the solution of the STP-MSP and the CDC problem.

This paper is organized as follows. In Section II we review related work and in Section III we formulate the problem. Section IV presents the decomposition framework. Distributed approximation algorithms for placing the Cover MBNs are presented in sections V and VI. A new approach to placing the Relay MBNs is described in Section VII. A joint solution to the CDC problem is discussed in Section VIII. In Section IX we evaluate the algorithms via simulation. We summarize the results in Section X. Due to space constraints, some of the proofs are omitted and can be found in [26].

II. Related Work

Several problems that are somewhat related to the CDC problem have been studied in the past. For simplicity, when describing these problems we will use our terminology (RNs and MBNs). One such problem is the Connected Dominating Set problem [4]. Unlike the CDC problem, in this problem there is no distinction between the communication ranges of RNs and MBNs. Additionally, MBNs' locations are restricted to RNs' locations. Similarly, the Connected Facility Location problem [27] also restricts potential MBN locations. Furthermore, this problem implies a cost structure that is not directly adaptable to that of the CDC problem. Lu et al. [20] study a Connected Sensor Cover problem [12], where the objective is to cover discrete targets while maintaining overall network connectivity and maximizing network lifetime. The set of constraints in this problem can be mapped to the CDC problem. However, the objective function and algorithmic approach are different.

We note that Tang et al. [28] have recently independently formulated and studied the CDC problem (termed in [28] as the Connected Relay Node Single Cover). A centralized 4.5-approximation algorithm for this problem is presented in [28]. In section IV, we will show that our approach provides a *centralized* 3.5-approximation for the CDC problem.

We propose to solve the CDC problem by decomposing it into two NP-Complete subproblems: the Geometric Disk Cover (GDC) problem and the Steiner Tree Problem with Minimum number of Steiner Points (STP-MSP). Hochbaum and Maass [15] provided a Polynomial Time Approximation Scheme (PTAS) for the GDC problem. However, their algorithm is impractical for our purposes, since it is centralized and has a high running time for reasonable approximation ratios. Several other algorithms have been proposed for the GDC problem (see the review in [5]). For example, Gonzalez [9] presented an algorithm based on dividing the plane into strips. In [5] it is indicated that this is an 8-approximation algorithm. We will show that by a simple modification, the approximation ratio is reduced to 6.

Problems related to the GDC problem under mobility are addressed in [6],[13],[16]. In [16], a 4-approximate centralized algorithm and a 7-approximate distributed algorithm are presented. Hershberger [13] presents a *centralized* 9-approximation algorithm for a slightly different problem: the mobile geometric *square* cover problem. We build upon his

[1] A γ-approximation algorithm for a minimization problem always finds a solution with value at most γ times the value of the optimal solution.

TABLE I
NOMENCLATURE

Symbol	*Definition*
$N = \{1, \ldots, n\}$	Set of RNs
$M = \{d_1, \ldots, d_m\}$	Set of MBNs
i_x, i_y	x- and y-coordinates of RN i
r, R	Communication ranges of RN and MBN
D	Communication diameter of an RN ($D = 2r$)
P_{d_i}	Set of RNs covered by MBN d_i
q_{ALGO}	Strip width for strip-based algorithm ALGO
α	Strip width as a fraction of D ($q = \alpha D$)
d_i^L, d_i^R	Leftmost, rightmost RNs covered by MBN d_i
M_{cover} (resp. M_{relay})	Set of MBNs that explicitly cover (resp. do not cover) at least one RN

approach in order to develop a *distributed* algorithm for the GDC problem. Clustering nodes to form a hierarchical architecture has been extensively studied in the context of wireless networks (e.g. [1],[4],[8]). However, the idea of deliberately controlling the motion of specific nodes in order to maintain some desirable network property (e.g. lifetime or connectivity) has been introduced only recently (e.g. [17],[21],[23]).

III. Problem Formulation

We consider a set of *Regular Nodes* (RNs) distributed in the plane and assume that a set of *Mobile Backbone Nodes* (MBNs) has to be deployed in the plane. We denote by N the collection of *Regular Nodes* $\{1, 2, \ldots, n\}$, by $M = \{d_1, d_2, \ldots, d_m\}$ the collection of MBNs, and by d_{ij} the distance between nodes i and j. The locations of the RNs are denoted by the $x - y$ tuples (i_x, i_y) $\forall i$.

We assume that the RNs and MBNs have both a communication channel (e.g. for data) and a low-rate control channel. For the communication channel, we assume the disk connectivity model. Namely, an RN i can communicate bi-directionally with another node j (i.e. an MBN) if the distance between i and j, $d_{ij} \leq r$. We denote by $D = 2r$ the diameter of the disk covered by an MBN communicating with RNs. Regarding the MBNs, we assume that MBN i can communicate with MBN j if $d_{ij} \leq R$, where $R > r$. For the control channel, we assume that both RNs and MBNs can communicate over a much longer range than their respective data channels. Since given a fixed transmission power, the communication range is inversely related to data rate, this is a valid assumption.

At this stage, we assume that the number of available MBNs is not bounded (e.g. if required, additional MBNs can be dispatched). Yet, in our analysis, we will try to minimize the number of MBNs that are actually deployed . Finally, we assume that all nodes can detect their position, either via GPS or by a localization mechanism. We shall refer to the problem of Mobile Backbone Nodes Placement as the Connected Disk Cover (CDC) problem and define it as follows.

Problem CDC: Given a set of RNs (N) distributed in the plane, place the smallest set of MBNs (M) such that:

1) For every RN $i \in N$, there exists at least one MBN $j \in M$ such that $d_{ij} \leq r$.
2) The undirected graph $G = (M, E)$ imposed on M (i.e. $\forall k, l \in M$, define an edge $(k, l) \in E$ if $d_{kl} \leq R$) is connected.

We will study both the case in which the nodes are static, and the case in which the RNs are mobile and some of the MBNs move around in order to maintain a solution the CDC problem. We assume that there exists some sort of MBN routing algorithm, which routes specific MBNs from their old locations to their new ones. The actual development of such an algorithm is beyond the scope of this paper.

We now introduce additional notation required for the presentation and analysis of the proposed solutions (Table I includes some of the notation used throughout the paper). A few of the proposed algorithms operate by dividing the plane into strips. When discussing such algorithms, we assume that the RNs in a strip are ordered from left to right by their x-coordinate and that ties are broken by the RNs' identities (e.g. MAC addresses). Namely, $i < j$, if $i_x < j_x$ or $i_x = j_x$ and the ID of i is lower than ID of j. We note that in property (1) of the CDC problem it is required that every RN is connected to at least one MBN. We assume that even if an RN can connect to multiple MBNs, it is actually assigned to exactly one MBN. Thus, we denote by P_{d_i} the set of RNs connected to MBN d_i. We denote by d_i^L and d_i^R the leftmost and rightmost RNs connected to MBN d_i (their x-coordinates will be denoted by $(d_i^L)_x$ and $(d_i^R)_x$). Similarly to the assumption regarding the RNs, we assume that the MBNs in a strip are ordered left to right by the x-coordinate of their leftmost RN ($(d_i^L)_x$).

In order to evaluate the performance of the *distributed algorithms*, we define the following standard performance measures. We define the *Time Complexity* as the number of communication rounds required in reaction to an RN movement. We assume that during each round a node can exchange errorless control messages with its neighbors. We define the *Local Computation Complexity* as the complexity of the computation that may be performed by a node in reaction to its (or another node's) movement. We assume that the nodes maintain an ordered list of their neighbors. Hence, the Local Computation Complexity refers to the computation required to maintain this list as well as to make algorithmic decisions.

IV. Decomposition Framework

In this section we obtain an upper bound on the performance of an approach that solves the CDC problem by decomposing it and solving each of the two subproblems separately. The first subproblem is the problem of placing the minimum number of *Cover MBNs* such that all the RNs are connected to at least one MBN. In other words, all the RNs have to satisfy only property (1) in the CDC problem definition. This problem is the Geometric Disk Cover (GDC) problem [15] which is formulated as follows:

Problem GDC: Given a set N of RNs (points) distributed in the plane, place the smallest set M of Cover MBNs (disks) such that for every RN $i \in N$, there exists at least one MBN $j \in M$ such that $d_{ij} \leq r$.

The second subproblem deals with a situation in which a set of *Cover MBNs* is given and there is a need to place

the minimum number of *Relay MBNs* such that the formed network is connected (i.e. satisfies only property (2) in the CDC problem definition). This subproblem is equivalent to the Steiner Tree Problem with Minimum Number of Steiner Points (STP-MSP) [19] and can be formulated as follows:

Problem STP-MSP: Given a set of Cover MBNs (M_{cover}) distributed in the plane, place the smallest set of Relay MBNs (M_{relay}) such that the undirected graph $G = (M, E)$ imposed on $M = M_{cover} \cup M_{relay}$ (i.e. $\forall k, l \in M$, define an edge (k, l) if $d_{kl} \leq R$) is connected.

We now define a *Decomposition Based CDC Algorithm* and bound the worst case performance of such an algorithm.

Definition 1: A *Decomposition Based CDC Algorithm* solves the CDC problem by using a γ-approximation algorithm for solving the GDC problem, followed by using a δ-approximation algorithm for solving the STP-MSP.

Theorem 1: For $R \geq 2r$, the Decomposition Based CDC Algorithm yields a $(\gamma + \delta)$-approximation for the CDC problem.

Proof: Define $ALGO$ as the solution obtained by the Decomposition Based CDC Algorithm. Also, define $ALGO_{cov}$ and $ALGO_{rel}$ as the set of *Cover* and *Relay* MBNs in $ALGO$. Specifically, an MBN a_i is a *Cover* MBN if it covers at least 1 RN (i.e. $P_{a_i} \neq \emptyset$). Otherwise, a_i is a Relay MBN. Next, define OPT_{CDC} as the overall optimal solution similarly broken up into OPT^{cov}_{CDC} and OPT^{rel}_{CDC}. Thus we have that,

$$\begin{aligned} |ALGO| &= |ALGO_{cov}| + |ALGO_{rel}| \\ &\leq \gamma \cdot |OPT_{cov}| + \delta \cdot |OPT_{ALGO-cov-rel}| \end{aligned}$$

where OPT_{cov} represents the optimal GDC of the RNs, and $OPT_{ALGO-cov-rel}$ represents the optimal STP-MSP solution connecting the Cover MBNs placed by the γ approximate GDC algorithm, $ALGO_{cov}$.

Next, we make use of the fact that given $ALGO_{cov}$ as the input Cover MBNs, a candidate STP-MSP solution can be constructed by placing MBNs in the positions defined by the nodes in OPT_{CDC}. This is a valid STP-MSP solution, since $ALGO_{cov}$ is a valid GDC for the RNs, and therefore, every MBN in $ALGO_{cov}$ is at most a distance r away from *some* RN. Since OPT^{cov}_{CDC} is also a valid GDC, it follows that every MBN in $ALGO_{cov}$ is at most a distance $2r$ from *some* MBN in OPT^{cov}_{CDC}. Therefore, as long as $R \geq 2r$, the MBNs in $ALGO_{cov} \cup OPT_{CDC}$ form a connected network. Finally, since $OPT_{ALGO-cov-rel}$ represents an STP-MSP solution that must have a lower cost than this candidate solution, we have that,

$$\begin{aligned} |ALGO| &\leq \gamma \cdot |OPT_{cov}| + \delta \cdot (|OPT^{cov}_{CDC}| + |OPT^{rel}_{CDC}|) \\ &\leq (\gamma + \delta) \cdot |OPT^{cov}_{CDC}| + \delta \cdot |OPT^{rel}_{CDC}| \\ &\leq (\gamma + \delta) \cdot |OPT_{CDC}|, \end{aligned}$$

where the second line follows from the fact that the optimal GDC for the RNs is of lower cost than OPT^{cov}_{CDC}. ■

According to Theorem 1, even if the two subproblems are solved optimally (i.e. with $\gamma = \delta = 1$), this yields a 2-approximation to the CDC problem. A tight example of this fact is illustrated in Fig. 2. Fig. 2-a shows an n node instance

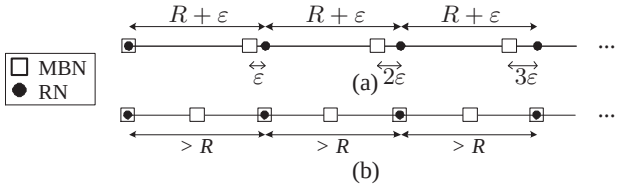

Fig. 2. Tight example of the approximation ratio of the decomposition algorithm: (a) optimal solution and (b) decomposition algorithm solution

of the CDC problem, where $\varepsilon << r$ refers to a sufficiently small constant. Also shown is the optimal solution with cost n MBNs. Fig. 2-b shows a potential solution obtained by using the decomposition framework (with $\gamma = \delta = 1$), composed of an optimal disk cover and an optimal STP-MSP solution. The cost is $n + n - 1 = 2n - 1$ MBNs. This example highlights the fact that under the Decomposition Framework, the cover MBNs are placed without considering the related problem of placing the relay MBNs.

Although in the next sections we mainly focus on distributed algorithms, we note that if a centralized solution can be tolerated, the approximation ratio of the GDC problem can be very close to 1 (e.g. using a PTAS [15]). Similarly, the lowest known approximation ratio of the STP-MSP solution (obtained by a centralized algorithm) is 2.5 [3]. Therefore, by Theorem 1, the *framework immediately yields a centralized 3.5-approximation algorithm for the solution of the CDC problem*.[2] This improves upon the centralized 4.5-approximation algorithm, recently presented in [28]. Since both algorithms use a PTAS, their respective complexities are quite high. The key point with respect to our Decomposition Framework is that *any future improvement to the approximation ratio of the STP-MSP will directly reduce the CDC approximation ratio*.

V. Placing the Cover MBNs - Strip Cover

Hochbaum and Maass [15] introduced a method for approaching the GDC problem by (i) dividing the plane into equal width strips, (ii) solving the problem locally on the points within each strip, and (iii) taking the overall solution as the union of all local solutions. Below we present algorithms that are based on this method. These algorithms are actually two different versions of a single generic algorithm. The first version locally covers the strip with rectangles encapsulated in disks while the second version locally covers the strip directly with disks. We then generalize (to arbitrary strip widths) the effects of solving the problem locally in strips and use this extension to provide approximation guarantees. Finally, since we are interested in distributed algorithms, we briefly discuss the distributed implementation.

A. Centralized Algorithms

For simplicity of the presentation, we start by describing the centralized algorithms. The two versions of the Strip Cover algorithm (*Strip Cover with Rectangles* - SCR and *Strip Cover with Disks* - SCD) appear below. In line 6, the

[2]When we use *our distributed* algorithms (presented in the following sections) within the framework, the overall approximation ratio is higher.

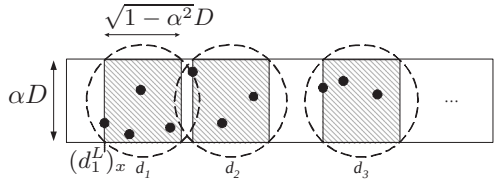

Fig. 3. An example illustrating step 9 of the SCR algorithm

first version (SCR) calls the *Rectangles* procedure and the second one (SCD) calls the *Disks* procedure. The input is a set of points (RNs) $N = \{1, 2, \ldots, n\}$ and their (x, y) coordinates, $(i_x, i_y)\,\forall i$. The output includes a set of disks (MBNs) $M = \{d_1, d_2, \ldots, d_m\}$ and their locations such that all points are covered. The first step of the algorithm is to divide the plane into K strips of width $q_{SC} = \alpha D$ (recall that $D = 2r$). The values of q_{SC} that guarantee certain approximation ratios will be derived below. We denote the strips by S_j and the set of MBNs in strip S_j by M_{S_j}.

Algorithm 1 Strip Cover with Rectangles/Disks (SCR/SCD)

1: **divide** the plane into K strips of width $q_{SC} = \alpha D$
2: $M_{S_j} \leftarrow \emptyset, \forall j = 1, \ldots, K$
3: **for** all strips $S_j, j = 1, \ldots, K$ **do**
4: **while** there exist uncovered RNs in S_j **do**
5: **let** i be the leftmost uncovered RN in S_j
6: **call Rectangles(i)** or **call Disks(i)**
7: $M_{S_j} \leftarrow M_{S_j} \cup d_k$
8: **return** $\bigcup_j M_{S_j}$

Procedure Rectangles(i)

9: **place** an MBN d_k such that it covers all RNs in the rectangular area with x-coordinates $[i_x, i_x + \sqrt{1-\alpha^2}D]$
10: **return** d_k

Procedure Disks(i)

11: $P_{d_k} \leftarrow \emptyset$ {set of RNs covered by the current MBN d_k}
12: **while** $P_{d_k} \cup i$ coverable by a single MBN (disk) **do**
13: $P_{d_k} \leftarrow P_{d_k} \cup i$
14: **if** there are no more RNs in the strip **then**
15: **break**
16: **let** i be the next leftmost uncovered RN in S_j not currently in P_{d_k}
17: **place** MBN (disk) d_k such that it covers the RNs P_{d_k}
18: **return** d_k

An example of the SCR algorithm and in particular of step 9 in which disks are placed such that they compactly cover all points in the rectangular area with x-coordinate range i_x to $i_x + \sqrt{1-\alpha^2}D$ is shown in Fig. 3.

As mentioned above, Gonzalez [9] presented an algorithm for covering points with unit-squares. It is based on dividing the plane into equal width strips and covering the points in each of the strips separately. In [5] it was indicated that when the same algorithm is applied to covering points with unit disks, the approximation ratio is 8. The SCR algorithm is actually a slight modification to the algorithm of [9]. Unlike in [9], in our algorithm we allow the selection of the strip width. This will enable us to prove that the approximation ratio for covering points with unit disks is actually 6.

The SCD algorithm requires to answer the following question (in Step 12): can a set of points $P_{d_k} \cup i$ be covered by a single disk of radius r? This is actually the decision version of the *1-center problem*[3]. Many algorithms for solving this problem exist, an example being an $O(n \log n)$ algorithm due to [14]. We will show that solving the 1-center problem instead of compactly covering rectangles (as done in the SCR algorithm) provides a lower approximation ratio.

The computational complexity of the SCR algorithm is $O(n \log n)$, resulting from sorting the points by ascending x-coordinate. In the SCD algorithm the 1-center subroutine may potentially need to be executed as many as $O(n)$ times for each of the $O(n)$ disks placed. Therefore, the computation complexity is $O(C(n)n^2)$, where $C(n)$ is the running time of the 1-center subroutine used in steps 12 and 17. By using a binary search technique to find the maximal P_{d_k}, we can lower the complexity to $O(C(n)n \log n)$.

B. Worst Case Performance Analysis

Let algorithm A denote the local algorithm within a strip, and let $|A_{S_j}|$ denote the cardinality of the GDC solution found by algorithm A covering only the points in strip S_j. Let algorithm B represent the overall algorithm, which works by running algorithm A locally within each strip and taking the union of the local solutions as the overall solution. In our case algorithm B is either the SCR or SCD algorithm and algorithm A is composed of steps 4-7 within the for loop.

Let $|OPT|$ represent the cardinality of an optimal solution of the GDC problem in the plane and $|OPT_{S_j}|$ the cardinality of an optimal solution for points exclusively within strip S_j. Note that $OPT \neq \bigcup_{S_j} OPT_{S_j}$, since OPT can utilize disks covering points across multiple strips. Finally, let Z_A denote the worst case approximation ratio of algorithm A. Namely, Z_A is the maximum of $|A_{S_j}|/|OPT_{S_j}|$ over all possible point-set configurations in a strip S_j. Similarly, let Z_B denote the worst case approximation ratio of algorithm B.

We characterize Z_B as a function of Z_A. Namely, if $q \leq D$, the cardinality of the solution found by algorithm B is at most $(\lceil \frac{D}{q} \rceil + 1)Z_A$ times that of the optimal solution, $|OPT|$. The proof can be found in [26].

Observation 1: If the strip width is $q \leq D$, a single disk can cover points from at most $(\lceil \frac{D}{q} \rceil + 1)$ strips.

Lemma 1: If the strip width is $q \leq D$, $Z_B = (\lceil \frac{D}{q} \rceil + 1)Z_A$.

We now show that in the SCR algorithm, $Z_A = 2$. This approximation ratio is tight, as illustrated in Fig. 4-a. We provide an inductive proof, since a similar proof methodology will be used in order to obtain the approximation ratios of the other GDC algorithms.

Lemma 2: If the strip width $q_{SC} \leq \frac{\sqrt{3}D}{2}$, steps 4-7 of the SCR algorithm provide a 2-approximation algorithm for the GDC problem within a strip.

Proof: Consider some strip S. Let $OPT_S = \{d_1, d_2, \ldots, d_{|OPT_S|}\}$ and $ALGO_S = \{a_1, a_2, \ldots, a_{|ALGO_S|}\}$ denote an optimal in-strip solution and SCR in-strip subroutine (steps 4-7) solution, respectively. Recall that we assume that the MBNs of both OPT_S and $ALGO_S$ are ordered from left to right by x-coordinate of the leftmost covered point (i.e. $i < j$ if $(d_i^L)_x \leq (d_j^L)_x$). Finally, define a_{b_m} as the b_m^{th} algorithm

[3]The 1-center problem for a set of points P is to find the location of the *center* from which the maximum distance to any point in P is minimized.

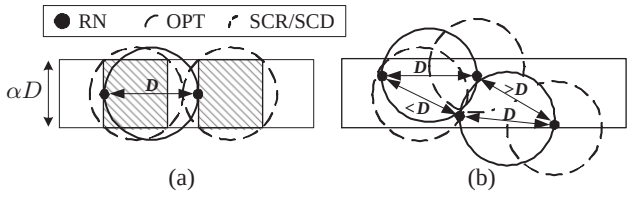

Fig. 4. Tight examples of the 2 and 1.5 approximation ratios obtained by the in-strip subroutines of the (a) SCR and (b) SCD algorithms

disk (from the left) corresponding to the disk that covers the rightmost point covered by the m^{th} OPT_S disk d_m.

Let $q_{SC} = \alpha D, \alpha < 1$. We now prove by induction that if $\alpha \leq \sqrt{3}/2$, the in-strip subroutine has approximation ratio of 2, i.e. $|ALGO_S| = b_{|OPT_S|} \leq 2|OPT_S|$.

Base Case: The *area* covered by d_1 (the leftmost optimal disk) is bounded by a rectangle with x-coordinate range $(d_1^L)_x$ (the x-coordinate of the leftmost point) to $(d_1^L)_x + D$. The *minimum* area covered by two SCR algorithm disks whose leftmost point is $(d_1^L)_x$ is a rectangle with x-coordinate range $(d_1^L)_x$ to $(d_1^L)_x + 2\sqrt{1-\alpha^2}D$. Thus, if $2\sqrt{1-\alpha^2}D \geq D$, $b_1 \leq 2$. This condition is met if $q_{SC} \leq \sqrt{3}D/2$.

Inductive Step: Assume that the in-strip algorithm uses no more than $2m$ disks to cover all the points covered by $d_1, \ldots, d_m$ (i.e. $b_m \leq 2m$). Consider the number of additional disks it takes for the algorithm to cover the points covered by $d_1, \ldots, d_m, d_{m+1}$. Since all of the points up to the rightmost point of d_m are already covered, by the same argument as the base case, the algorithm will use at most 2 extra disks to cover the points covered by d_{m+1}. It thus follows that if $q \leq \sqrt{3}D/2$, $b_{m+1} \leq b_m + 2 \leq 2m + 2 = 2(m+1)$. ■

By combining the results of lemmas 1 and 2, we obtain the approximation ratio of the SCR algorithm.

Theorem 2: If $\frac{D}{2} \leq q_{SC} \leq \frac{\sqrt{3}D}{2}$, the SCR algorithm is a 6-approximation algorithm for the GDC problem.

Proof: Define algorithm A as the in-strip subroutine of the SCR algorithm (steps 4-7) and algorithm B as the SCR algorithm. From Lemma 2, for $q \leq \sqrt{3}D/2$, $Z_A = 2$. From Lemma 1, $Z_B \leq Z_A(\lceil D/q \rceil + 1)$, the minimum value of which (for $q < D$) is $3Z_A$. This is attained when $q \geq D/2$. ■

In the lemma below we show that for the SCD algorithm $Z_A = 1.5$. The proof (omitted for brevity and can be found in [26]) follows from an inductive argument very similar to that of Lemma 2. The key difference is that given a leftmost RN covered by an OPT disk d_i, if either (i) d_i is the rightmost OPT disk or (ii) $(d_i^R)_x < (d_{i+1}^L)_x$, then the SCD algorithm will only use 1 disk to cover the RNs covered by d_i. In contrast, in such a case the SCR algorithm may still use 2 disks.

Lemma 3: If $q_{SC} \leq \frac{\sqrt{3}D}{2}$, steps 4-7 of the SCD algorithm provide a 1.5-approximation algorithm for the GDC problem within a strip.

Combining this result with Lemma 1 (similarly to the derivation of Theorem 2), we obtain the approximation ratio of the SCD algorithm. The approximation ratio for the in-strip subroutine of the SCD algorithm is tight, as shown in Fig. 4-b. For the problem instance presented in the figure, the optimal solution requires 2 disks, whereas the SCD algorithm always places 3 disks.

Theorem 3: If $\frac{D}{2} \leq q_{SC} \leq \frac{\sqrt{3}D}{2}$, the SCD algorithm is a 4.5-approximation algorithm for the GDC problem.

C. Average Case Performance Analysis

Up to now we discussed the *worst case* performance. We now wish to bound the approximation ratios in the *average case*. We assume that the RNs are randomly distributed according to a two dimensional Poisson process of density λ $nodes/unit^2$. A key property of such a distribution is that when the number of RNs is given, their positions are independent and each is *uniformly distributed* in the plane. Due to the random locations of the RNs, the number of MBNs placed by an optimal algorithm, $|OPT|$ is a random variable. Similarly, we define $|SCR|$ and $|SCD|$ as random variables corresponding to the number of disks placed by the SCR and the SCD algorithms. We define the *average approximation ratios* β_{SCR} and β_{SCD} as,

$$\beta_{SCR} = \frac{E[|SCR|]}{E[|OPT|]},\ \beta_{SCD} = \frac{E[|SCD|]}{E[|OPT|]}.$$

It should be noted that β_{SCR} differs from the expected value of the approximation ratio (e.g. $E[|SCR|/|OPT|]$). Yet, it provides a good measure of the average performance.

The following theorem and corollary bound the average approximation ratios of both the SCR and SCD algorithms (since SCD always outperforms SCR). The proof of the theorem is by combining the results of the following lemmas. The proofs of the lemmas and the corollary can be found in the Appendix.

Theorem 4: Given RNs distributed in the plane according to a two dimensional Poisson process with density λ,

$$\beta_{SCD} \leq \beta_{SCR} \leq \frac{D^2\lambda + 2D\sqrt{\lambda} + 1}{\alpha\sqrt{1-\alpha^2}D^2\lambda + 1}. \quad (1)$$

Corollary 1: If $q = \frac{D}{\sqrt{2}}$, then $\beta_{SCD} \leq \beta_{SCR} \leq 3$.

The consequence of the above is that although the worst case approximation ratios of the SCR/SCD algorithms are 6 and 4.5 (respectively), selecting a specific strip width results in an average approximation ratio which is bounded by 3. It is interesting to note that this strip-width lies in the range required for the worst case analysis of theorems 2 and 3.

Lemma 4: Given a strip width $q = \alpha D$, and an $L \times K\alpha D$ planar area,

$$E[|SCR|] \leq \frac{\lambda\alpha DLK}{\lambda\alpha\sqrt{1-\alpha^2}D^2 + 1}.$$

Lemma 5: Given an $L \times K\alpha D$ planar area,

$$E[|OPT|] \geq \frac{KL\alpha D}{D^2 + \frac{1}{\lambda} + \frac{2D}{\sqrt{\lambda}}}.$$

Finally, note that for a large number of RNs, the assumption that they are uniformly distributed is perhaps not realistic. In general, the RNs may tend to cluster together, resulting in nodes concentrated within single strips (rather than spread across a large number of strips). This will result in a better average case performance, since the strip-based algorithms are most effective when covering RNs within a single strip. Thus, β_{SCD} and β_{SCR} derived in this section are actually upper bounds on realistic average approximation ratios.

D. Distributed Implementation

The SCR and SCD algorithms can be easily implemented in a distributed manner. The algorithms are executed at the RNs and operate within the strips. The SCR algorithm executed at an RN i is described below. Recall that we denote the RNs within a strip according to their order from the left (i.e. $i < j$ if $i_x \leq j_x$). Ties are broken by node ID.

Algorithm 2 Distributed SCR (at RN i)

Initialization
1: **let** G_i be the set of RNs j such that $j < i$ and $i_x - j_x \leq D$
2: **if** $G_i = \emptyset$ **then**
3: **call** Place MBN
Construction and Maintenance
4: **if** *MBN Placed* message received **then**
5: **call** Place MBN
6: **if** i is disconnected from its MBN or enters from a neighboring strip **then**
7: **if** there is at least one MBN within distance r **then**
8: **join** one of these MBNs
9: **else call** Place MBN
Procedure Place MBN
10: **let** i^R be the rightmost RN s.t. $(i^R)_x \leq i_x + \sqrt{1-\alpha^2}D$
11: **place** MBN d_k covering RNs j, where $j_x \in [i_x, (i^R)_x]$
12: **if** $(i^R + 1)_x - (i^R)_x \leq D$ **then**
13: **send** an *MBN Placed* message to $i^R + 1$

Every RN that has no left neighbors within distance D initiates the disk placement procedure that propagates along the strip. The propagation stops once there is a gap between nodes of at least D. If an RN arrives from a neighboring strip or leaves its MBN's coverage area, it initiates the disk placement procedure that may trigger an update of the MBN's locations within the strip. Notice that MBNs only move when a recalculation is required. Although the responsibility to place and move MBNs is with the RNs, simple enhancements would allow the MBNs to reposition themselves during the maintenance phase. If after a recalculation, an MBN is not repositioned, then it is not required and can be used elsewhere. The time complexity (i.e. number of rounds) is $O(n)$. The computation complexity is $O(\log n)$. Control information has to be transmitted between RNs over a distance $D = 2r$.

The distributed SCD algorithm is similar to the distributed SCR algorithm. The main difference is that in Step 10 of *Place MBN*, i^R is defined as the rightmost coverable point (by a single disk of radius r), given that i is the leftmost point. Finding this point requires solving 1-center problems. Then, in Step 11 a disk that covers all the points between i and i^R should be placed. The time complexity of the distributed SCD algorithm is again $O(n)$. The local computation complexity is $O(C(n) \log n)$ to calculate the value of i^R, where $C(n)$ is the running time of the 1-center subroutine used.

VI. Placing the Cover MBNs - Mobile Cover

A. MObile Area Cover (MOAC) Algorithm

In the SCR and SCD algorithms, an RN movement may change the allocation of RNs to MBNs along the whole strip. Thus, although they may operate well in a relatively static environment, it is desirable to develop algorithms that are more tailored to frequent node movements. In particular, it is desirable to develop algorithms that are *adaptive*, i.e. require only local updates in response to local node movements. In this section we present such an algorithm which builds upon ideas presented in [13]. Hershberger [13] studied the problem of covering moving points (e.g. RNs) with mobile unit-squares (e.g. MBNs). Since the d-dimensional *smooth maintenance scheme* proposed in [13] does not easily lend itself to distributed implementation, we focus on the *simple 1-D algorithm* proposed there.

Applied to our context, the Simple 1-D algorithm covers mobile RNs along the strip with length D rectangles (MBNs). The key feature is that point transfers between MBNs are *localized*. Namely, changes do not propagate along the strip. According to [13], the algorithm has a worst case performance ratio of 3.[4] Extending the *Simple 1-D algorithm* of [13] to diameter D disks is not straightforward. We will first show that an attempt to simply use rectangles encapsulated in disks without any additional modifications results in a 4-approximation to the GDC problem within a strip. Then, we will present the MObile Area Cover (MOAC) algorithm which reduces the approximation ratio to 3.

We define the strip width as $q_{MOAC} = \alpha D$. We reduce disks to the rectangles encapsulated in them and use these rectangles to cover points within the strip, as was depicted in Fig. 3. The rectangles cover the strip width (αD) and their length is *at most* $\sqrt{1-\alpha^2}D$. We set $D = 1$ and $\alpha = \sqrt{5}/3$ (resulting in $\sqrt{1-\alpha^2}D = 2/3$). These are arbitrary values selected for the ease of presentation. Yet, the algorithm and the analysis are applicable to any $1/2 \leq \alpha \leq \sqrt{5}/3$. We restate the set of rules from [13] using our terminology and assuming (unlike [13]) that the rectangles' lengths are at most 2/3.

Algorithm 3 Simple 1-D [13] with $\sqrt{1-\alpha^2}D = 2/3$

0 **initialize** the cover greedily {using the SCR algorithm}
1 **maintain** the leftmost RN and rightmost RN of each MBN rectangle
2 **if** two adjacent MBN rectangles come into contact **then**
exchange their outermost RNs
3 **If** a set of RNs covered by an MBN becomes too long {the separation between its leftmost and rightmost RNs becomes greater than 2/3} **then**
split off its rightmost RN into a singleton MBN
check whether rule 4 applies
4 **if** two adjacent MBN rectangles fit in a 2/3 rectangle **then**
merge the two MBNs

The following lemma provides the performance guarantee of this algorithm. The proof follows a similar inductive methodology as that of Lemma 2, with the key observation that at most 5 algorithm MBNs can cover RNs covered by a single optimal MBN. Notice that since the changes are kept local, the approximation ratio holds at all time (i.e. there is no need to wait until the changes propagate).

Lemma 6: The *Simple 1-D algorithm [13]* with $\sqrt{1-\alpha^2} = 2/3$ is *at all times* a 4-approximation algorithm for the GDC

[4]We note that using the same inductive proof methodology, used for Lemma 2, one can show that the simple 1-D algorithm actually maintains a 2-approximation at all times.

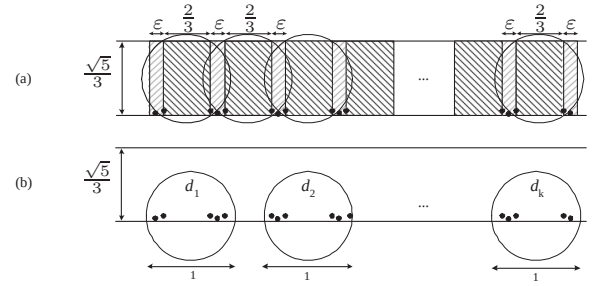

Fig. 5. Worst case example for the performance of the Simple 1-D algorithm: (a) algorithmic solution and (b) optimal solution. The number of optimal MBNs is denoted by k.

problem within a strip.

Recall that the overall solution to the GDC problem in the plane involves combining the solutions obtained in every strip. Due to lemmas 1 and 6, if implemented simultaneously in every strip, the algorithm provides a 12-approximation for the GDC problem in the plane, which is relatively high. We now focus on enhancements that reduce the approximation ratio while maintaining the desired locality property.

Fig. 5 presents an example which shows that the approximation ratio described in Lemma 6 is tight. It is shown that optimal MBN d_1 can cover RNs that are covered by 4 algorithm MBNs. Two of these algorithm MBNs cover RNs that are within $2/3$-length rectangles, while the two other cover a single RN. Similarly, d_2 covers RNs from 4 additional algorithm MBNs, and so on. Optimal MBN d_k covers RNs from exactly 3 algorithm MBNs that have not been covered by optimal MBNs $\{d_1, \ldots, d_{k-1}\}$. The resulting approximation ratio is $(4(k-1)+3)/k \approx 4$. One of the sources of inefficiency is the potential presence of ε-length MBNs (e.g. covering a single RN) that cannot merge with their $2/3$-length neighbor MBNs. Thus, up to 5 MBNs deployed by the Simple 1-D algorithm may cover points which are covered by a single optimal MBN (e.g. d_2 in Fig. 5). As long as such narrow MBNs can be avoided, a better approximation can be achieved. We now modify the Simple 1-D algorithm to yield the MOAC algorithm in which ε-length MBNs cannot exist.

Before describing the algorithm, we make the following definitions. For MBN d_i, in addition to its leftmost and rightmost RNs, defined earlier, as d_i^L and d_i^R, we also define L_i and R_i as the x-coordinates of its left and right *domain* boundaries. The interpretation of MBN d_i's domain is that any point in the x-range of $[L_i, R_i]$ will automatically become a member point of MBN d_i. Recall that by definition MBN d_i is to the left of MBN d_j if $(d_i^L)_x < (d_j^L)_x$.

The MOAC algorithm operates within strips and maintains the following *invariants* in each strip (in order of priority) at all times, for every MBN d_i:

1) Domain definition: $L_i \leq (d_i^L)_x \leq (d_i^R)_x \leq R_i$.
2) Domain length[5]: $\frac{1}{3} \leq |R_i - L_i| \leq \frac{2}{3}$.
3) Domain disjointness: $[L_i, R_i] \bigcap [L_j, R_j] = \emptyset, \forall d_j \in M$.
4) Domain influence: $\forall p \in N$, $L_i \leq p_x \leq R_i \Leftrightarrow p_x \in P_{d_i}$.

[5]The upper bound is the coverage length of a MOAC MBN (here arbitrarily chosen as $\sqrt{1-\alpha^2}D = 2/3$). To maintain the algorithm's properties, the lower bound should be half of the upper bound and their sum should be at least one. In addition, due to Lemma 1, $\alpha \geq 0.5$ has to hold.

Algorithm 4 MObile Area Cover (MOAC)

Initialization
1: **cover** the RNs with MBNs using the SCR algorithm
2: **for** all MBNs i **do**
3: $L_i \leftarrow d_i^L$; $R_i \leftarrow d_i^L + \frac{2}{3}$
4: $P_{d_i} \leftarrow$ all RNs within $[L_i, R_i]$
Maintenance
5: **if** an RN $p \in P_{d_i}$ moves *right* such that $p_x > R_i$ **then**
6: **if** $L_j \leq p_x \leq R_j$, $j \neq i$ $\{p$ in d_j's domain$\}$ **then**
7: **remove** p from P_{d_i}
8: **else if** $|p_x - L_i| \leq \frac{2}{3}$ **then**
9: **stretch** L_i and R_i to maintain invariant (1) by setting $R_i \leftarrow p_x$ and $L_i \leftarrow max(L_i, p_x - \frac{2}{3})$
10: **else** $\{p$ not in the immediate domain of any MBN$\}$
11: **remove** p from P_{d_i}
Disconnection
12: **if** at any time there exists an uncovered RN p **then**
13: **if** for some MBN d_j, $L_j \leq p_x \leq R_j$ **then**
14: $P_{d_j} \leftarrow P_{d_j} \cup p$
15: **else if** for some MBN d_j, L_j and R_j can be *stretched* to include p while maintaining invariant (2) **then**
16: $P_{d_j} \leftarrow P_{d_j} \cup p$
17: **stretch** L_j and R_j to maintain invariants (1) and (2)
18: **else** $\{p$ cannot be covered by an existing MBN$\}$
19: **let** d_{j-1} and d_{j+1} represent the MBNs to the left and right of p
20: **if** $|L_{j+1} - R_{j-1}| \geq \frac{1}{3}$ {i.e. enough "open space" to maintain invariant (2)} **then**
21: **create** MBN d_j with $P_{d_j} = p$ and $|R_j - L_j| \geq \frac{1}{3}$ while maintaining invariant (3)
22: **else** $\{< \frac{1}{3}$ space around $p\}$
23: **shrink** MBN d_{j-1} such that $R_{j-1} = p_x - \frac{1}{3}$
24: **create** MBN d_j with $L_j = p_x - \frac{1}{3}$ and $R_j = p_x$
25: $P_{d_{j-1}} \leftarrow$ all points in $[L_{j-1}, R_{j-1}]$
26: $P_{d_j} \leftarrow$ all points in $[L_j, R_j]$
Merge
27: **if** there exists MBN d_j such that $|(d_j^R)_x - (d_i^L)_x| \leq \frac{2}{3}$ or $|(d_i^R)_x - (d_j^L)_x| \leq \frac{2}{3}$ **then**
28: **merge** d_j into d_i

The MOAC algorithm is described below. It consists of rules regarding construction and maintenance of the MBNs' domains. In particular, the *Initialization* phase that places the MBNs and constructs their domains is described in lines 1-4. In order to initially cover all the RNs, the MBNs are placed according to the SCR algorithm. Then, for each MBN, the left and right domain boundaries (L_i and R_i) are set as the x-coordinates of the leftmost RN covered by the MBN and the rightmost edge of the rectangle generated by SCR (recall the example in Fig. 3). In line 4 all the RNs within the boundaries are associated with the MBN. Since for q_{MOAC}, SCR generates $2/3$-length rectangles, at the end of the phase all the invariants hold.

The *Maintenance* phase (lines 5-11) takes care of a situation in which an RN leaves its MBN's domain boundary.[6] If the RN moves into a domain of another MBN, it is removed from the set of RNs covered by the MBN. The Disconnection phase will immediately take care of assigning it to the new MBN. Otherwise, the algorithm tries to move the right boundary such that the RN will be covered and the MBN's domain will be at most 2/3 (we refer to such an operation as stretching R_i).

[6]For brevity, we only state the operations when an RN moves to the *right* of the boundary (there are analogous operations for a leftward movement).

Finally, if the RN cannot be covered by stretching R_i, it is removed from the set of points covered by the MBN. The Disconnection phase will immediately create a new MBN for it.

The Disconnection phase takes care of cases in which an RN is disconnected from its MBN (as described above) and cases in which an RN enters from a neighboring strip. In the simplest cases, such an RN joins an existing MBN whose boundaries may have to be stretched in order to cover it. In other cases, a new MBN is created in order to cover the RN. It has to be carefully created such that its domain length is at least 1/3. Note that the operations in lines 22-26 can always be accomplished without violating invariant (2). This is due to the fact that an MBN d_j is created for point p only if $|p_x - L_{j-1}| > 2/3$ (otherwise MBN d_{j-1} would have been stretched to cover p), which implies there is enough space for two MBNs of size greater or equal to 1/3 to coexist.

Finally, in the *Merge* phase, two neighboring MBNs have to be merged since all their RNs are within a 2/3-long interval. It can be initiated by movements of some of the RNs or immediately following the previous phases. Following the merge in line 28, the MBN should update its L_i and R_i such that the domain will include all RNs and will satisfy invariant (2). This is always possible, since the two merged MBNs satisfy the invariants prior to their merger. We note that the algorithm can be implemented in distributed manner by applying some of the rules at the MBNs and some of them at disconnected (i.e. uncovered) RNs (it should be clear from the context where each rule should be applied).

The following lemma provides the performance guarantee of the MOAC algorithm within the strip. Its proof is almost identical to that of Lemma 6. The main difference is that due to the enforced Domain invariants, at most 4 algorithm MBNs can cover RNs covered by a single optimal MBN. From Lemma 1 it follows that if MOAC is simultaneously executed in all strips, it is a 9-approximation algorithm.

Lemma 7: The MOAC algorithm is a 3-approximation algorithm *at all times* for the GDC problem within a strip.

The time complexity of the MOAC algorithm is $O(1)$, since all node exchanges are local. The local computation complexity is potentially $O(\log n)$, due to the operation in line 23. The only assumption required is that MBNs and disconnected RNs have access to information regarding L_j, d_j^L, d_j^R and R_j of their immediate neighbors to the right and left (as long as they are less than $2D$ away). Thus, in terms of complexity, MOAC is the best of the distributed algorithms.

B. Merge-and-Separate (MAS) Algorithm

The relatively high approximation ratio of the MOAC algorithm results from the fact that it reduces disks into rectangles, thereby losing about 35% of disk coverage area. The difficulty in dealing with disks is that there are no clear *borders* and that even confined to a single strip, many disks can overlap although they cover disjoint nodes.

On average any algorithm with a merge rule should perform well. However, just having a merge rule is not sufficient in the rare but possible case where many mutually pairwise *non-mergeable* MBNs move into the same area. Based on this premise, we present the Merge-And-Separate (MAS) algorithm, as an algorithm which merges pairwise disks where possible (similar to the MOAC algorithm) and separates disks, if too many mutually non-mergeable disks concentrate in a small area. As will be shown, the MAS algorithm retains some of the localized features of the MOAC and obtains a better performance ratio. However, this comes at a cost of increased local computation complexity.

We define the strip-widths as $q_{MAS} = \alpha D$ and set $D = 1$, $\alpha = \sqrt{5}/3$, $\sqrt{1-\alpha^2} = 2/3$. These are arbitrary values selected for the ease of presentation, the algorithm and the analysis are applicable to any $0.5 \leq \alpha < \sqrt{3}/2$. Let $x_{R_{\{i,j,k\}}}$ and $x_{L_{\{i,j,k\}}}$ be the x-coordinates of the rightmost and leftmost points of $\{P_{d_i} \cup P_{d_j} \cup P_{d_k}\}$. The algorithm is initialized by covering the nodes within a strip with MBNs by using the distributed SCR algorithm. The algorithm that then operates at an MBN d_i is described below. We note that as in the previous algorithms, most of the operations are performed in reaction to an RN movement. However, in order to maintain the locality of the algorithm, the Separation operation is performed periodically at each MBN. Fig. 6 demonstrates the Separation done at lines 8-11. For correctness of the algorithm, we assume that both the merge and separate operations can be executed atomically (i.e. without any interrupting operation).

Algorithm 5 Merge-and-Seperate (MAS)

Initialization
1: **cover** the RNs with MBNs using the SCR algorithm
2: $P_{d_i} \leftarrow$ all RNs within $[L_i, R_i]$
Merge
3: **for** all MBNs d_k within $2D$ of d_i **do**
4: **if** $\{P_{d_i} \bigcup P_{d_k}\}$ can be covered by a single MBN **then**
5: **merge** d_i and d_k
Separation
6: **for** all MBN pairs d_j, d_k within $2D$ of d_i **do**
7: **if** $|x_{R_{\{i,j,k\}}} - x_{L_{\{i,j,k\}}}| \leq 2D$ **then**
8: **separate and reassign** MBNs and RNs such that
9: $P_{d_i} \leftarrow$ all RNs in $[x_{L_{\{i,j,k\}}}, x_{L_{\{i,j,k\}}} + \frac{2}{3}]$
10: $P_{d_j} \leftarrow$ all RNs in $[x_{L_{\{i,j,k\}}} + \frac{2}{3}, x_{L_{\{i,j,k\}}} + \frac{4}{3}]$
11: $P_{d_k} \leftarrow$ all RNs in $[x_{L_{\{i,j,k\}}} + \frac{4}{3}, x_{R_{\{i,j,k\}}}]$
Creation
12: **if** an RN p enters from a neighboring strip or an RN $p \in P_{d_i}$, moves s.t. MBN d_i cannot cover P_{d_i} **then**
13: **create** a virtual MBN for p
14: **if** the virtual MBN cannot be *merged* with any of its neighbors **then**
15: **create** a new MBN to cover p

Define *steady state* as any point in time in which there are no merge or separate actions currently possible. Below we describe the performance of the MAS algorithm.

Lemma 8: In steady state, the MAS algorithm is a 2-approximation algorithm for the GDC problem within a strip.

Since point transfers are local (i.e. only take place between adjacent MBNs), the time complexity is $O(1)$. The computation complexity is $O(C(n))$ to evaluate the merge and the create rules, where $C(n)$ is the running time of the 1-center subroutine used. In order to make the required decisions, we assume that an MBN has access to all nearby (i.e. within a distance of $3D$) MBNs' point-sets and locations.

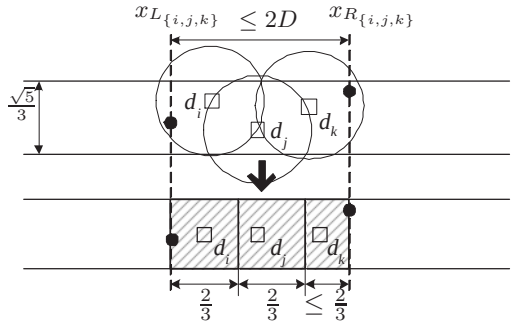

Fig. 6. The Separation rule of the MAS algorithm

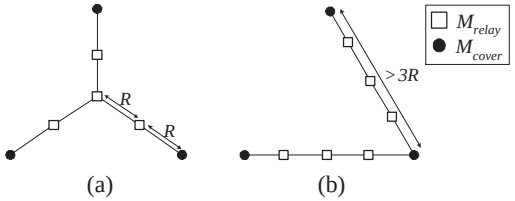

Fig. 7. STP-MSP solutions: (a) Optimal (4 Relay MBNs) and (b) MST-based (6 Relay MBNs)

VII. Placing the Relay MBNs

Recall that in Section IV we showed that the CDC problem can be decomposed into two subproblems. In this section, we focus on the second subproblem that deals with a situation in which a set of nodes (*Cover MBNs*) is given and there is a need to place the minimum number of nodes (*Relay MBNs*) such that the resulting network is connected. Recall that the distance between connected MBNs cannot exceed R. This problem is equivalent to the Steiner Tree Problem with Minimum number of Steiner Points (STP-MSP) [19].

In [19] a 4-approximation algorithm that places nodes along edges of the Minimum Spanning Tree (MST) which connects the Cover MBNs was proposed. In [2] an improved MST-based algorithm with an approximation ratio of 3 was proposed. These algorithms are simple and perform reasonably well in practice. However, their main limitation is that they only find *MST-based* solutions. Namely, since the Relay MBNs are in general placed along the edges of the MST, these algorithms cannot find solutions in which a Relay MBN is used as a central junction that connects multiple other Relay MBNs. An example demonstrating this inefficiency appears in Fig. 7.

We now present and analyze a *Discretization Approach* which provides a theoretical footing towards applying the vast family of discrete and combinatorial approaches that can potentially rectify the above inefficiency. The approach transforms the STP-MSP from an Euclidean problem to a discrete problem on a graph. Although the transformed problem does not admit a constant factor approximation, in many practical cases it can be solved optimally. We will show that if such a solution is obtained, it is a 2-approximation for the STP-MSP.

Our approach is based on an idea used by Provan [22] for dealing with the continuous analog of the STP-MSP problem, the Euclidean Steiner Minimal Tree (ESMT) problem [7]. In [22] it was proposed to discretize the plane and to solve a Network Steiner Tree problem [7] on the induced graph,

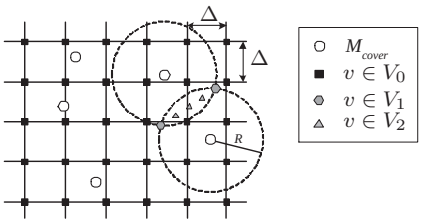

Fig. 8. An example of the sets V_0, V_1, and V_2 in the Discretization Approach

yielding an efficient approximate solution for the ESMT. We present a somewhat similar approach for solving the STP-MSP problem. Our approach is quite different from the approach of [22], since the STP-MSP problem is more sensitive to discretizing the plane than the ESMT problem.

Define V_0 as the lattice of points in the plane generated by gridding the plane with horizontal/vertical spacing Δ, the exact value of which will be derived later. Next, define V_1 as the set of pairwise intersection points of radius R circles drawn around each of the Cover MBNs. For the intersection region of any two circles, add three equally spaced points along the line between the two intersection points. Let V_2 denote the set of these points. The sets V_0, V_1 and V_2 are illustrated Fig. 8. Finally, define $conv(M_{cover})$ as the convex hull of the of Cover MBNs. We can now define

$$V = \left\{ (V_0 \cup V_1 \cup V_2 \cup M_{cover}) \cap^* conv(M_{cover}) \right\}.$$

where we define a special intersection operator $\cap^*$ to ensure that we pick enough points to be in V such that $conv(V) \supseteq conv(M_{cover})$.

For all $u, v \in V$, if $d_{uv} \leq R$, we define an edge (u, v). We denote the set of edges by E and the *induced graph* by $G = (V, E)$. Let the node weights be denoted by w_v. We now state the Node-Weighted Steiner Tree (NWST) problem [11],[18],[25], which has to be solved as part of our Discretization algorithm, presented below.

Problem NWST: Given a *node-weighted* undirected graph $G = (V, E)$ with zero-cost edges and a terminal set $M_{cover} \subseteq V$, find a minimum weight tree $T \subseteq G$ spanning M_{cover}.

Algorithm 6 Discretization

1: **create** the sets V_0, V_1, V_2, and V {Δ derived below}
2: $w_v \leftarrow 1 \, \forall v \in V - M_{cover}$
3: $w_v \leftarrow 0 \, \forall v \in M_{cover}$
4: **create** the set E
5: **find** a minimum weight NWST on $G = (V, E)$

The set of nodes selected in step 5 correspond to the Relay MBNs in the STP-MSP solution. We assume that step 5 is performed by a β_{NWST}-approximation algorithm. The following theorem provides the performance guarantee of the Discretization algorithm.

Theorem 5: If $\Delta \leq \frac{R}{7}$, the Discretization algorithm is a $2\beta_{NWST}$-approximation algorithm for the STP-MSP.

Our methodology in proving the theorem is as follows. We start by assuming the optimal STP-MSP tree is known, and describe a method to construct a candidate Steiner tree T in

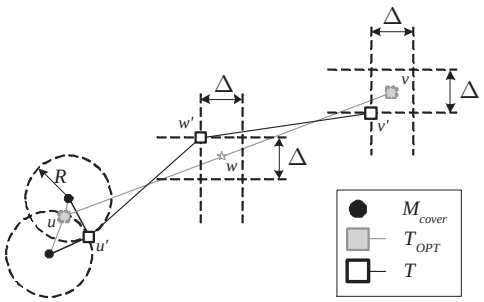

Fig. 9. An example of the construction of the candidate tree T from the optimal STP-MSP tree T_{OPT}

G from this optimal tree. We then use the definition of T in order to bound the ratio between an *approximate solution* to the Node-Weighted Steiner Tree (NWST) problem in G to the *optimal solution* of the STP-MSP in the plane.

Recall that the set of terminals/Cover MBNs M_{cover} is given as input to the problem. Define $T_{OPT} = (M^*, E^*)$ as the optimal solution to the STP-MSP. The node set M^* is composed of the Cover MBNs M_{cover} and the optimal set of Relay MBNs denoted by M^*_{relay}. Below, we present an algorithm for the construction of a candidate tree $T = (M^T, E^T)$ in the graph $G = (V, E)$. T is constructed such that it is a feasible STP-MSP solution. An example of steps 4-5, 7, and 12-14 of the algorithm is illustrated in Fig. 9.

Algorithm 7 Construction of a Feasible STP-MSP (CFS)

1: $M^T \leftarrow M_{cover}$
2: $E^T \leftarrow$ edges $(i,j) \in E^*$ where both $i, j \in M_{cover}$
3: **for** all $u \in M^*_{relay}$ that have edges (in E^*) to a set of Cover MBNs (in M_{cover}) **do**
4: **add** to M^T a Relay MBN $u' \in V$ located at the nearest point to u that can be directly connected to the same set of Cover MBNs
5: **add** to E^T edges connecting u' and the set of Cover MBNs
6: **for** all $u \in M^*_{relay}$ that do not have edges (in E^*) to any Cover MBNs in M_{cover} **do**
7: **add** to M^T a Relay MBN $u' \in V$ located at the nearest point to u
8: **for** all Relay MBNs $u, v \in M^*_{relay}$ s.t. $(u,v) \in E^*$ **do**
9: **if** $d_{u'v'} \leq R$ **then**
10: **add** to E^T an edge (u', v')
11: **else**
12: $w \leftarrow$ midpoint of the line segment $\overrightarrow{(u,v)}$
13: **add** to M^T a Relay MBN $w' \in V$ located at the nearest point to w
14: **add** to E^T edges $(u', w'), (w', v')$

In the following lemma we show that T is also a feasible solution to the NWST problem in G.

Lemma 9: If $\Delta \leq \frac{R}{7}$, then T, constructed by the CFS algorithm, is a Steiner tree in G.

Proof outline: We have to show that T connects all the nodes from M_{cover} by a tree whose nodes are in V and that the edges added to E^T are valid edges in E. The nodes of T (i.e. M^T) are by definition in V, since they are selected from V. A node in V satisfying the condition in step 4 always exists, since V includes the intersections of radius R circles drawn around each of the Cover MBNs.

Regarding the edges of T (i.e. E^T), note that those added in steps 2 and 5 must be valid in E, since by definition they represent edges between nodes in V that are less than R apart. The final part involves showing that edges added between new Relay MBNs (i.e. in steps 10 and 14) are of length at most R. This is done by using the triangle inequality and the definition of the set V. A detailed proof appears in [26]. ■

The next lemma shows that the number of Relay MBNs in T, i.e. $|M^T_{relay}| = |M^T| - |M_{cover}|$, is less than twice the number of Relay MBNs in the optimal solution of the STP-MSP (T_{OPT}). The proof of the Lemma and that of Theorem 5 can be found in the appendix.

Lemma 10: In T, constructed by the CFS algorithm, $|M^T_{relay}| < 2|M^*_{relay}|$.

It was shown in [18] that the NWST problem does not admit a constant factor approximation algorithm and that the best theoretically achievable approximation ratio is $\ln k$, where k is the number of terminals (in our formulation $k = |M_{cover}|$). Indeed, for the case in which all node weights are equal, [11] presented a $(\ln k)$-approximation algorithm. Thus, in general, the Discretization algorithm yields a worst case approximation ratio of $2\ln|M_{cover}|$. However, in some cases the NWST problem can be solved optimally by discrete methods such as integer programming [25]. Since in such cases $\beta_{NWST} = 1$, the approximation ratio will be 2. Notice that it is likely that the Discretization algorithm will have better average performance than the MST-type algorithms, due to the use of Relay MBNs as central junctions.

Since the Discretization algorithm takes care of placing only the *Relay MBNs*, it might be feasible to implement it in a centralized manner, as described above. Yet, if there is a need for a distributed solution, one of the MST-based algorithms [2],[19] should be used. Since these algorithms do not deal very well with the mobility of Cover MBNs, the development of distributed algorithms for the STP-MSP that take into account mobility remains an open problem.

VIII. Joint Solution

Using the decomposition framework presented in Section IV, the overall approximation ratio of the CDC problem is the sum of the approximation ratios of the algorithms used to solve the subproblems. Hence, this framework yields a *centralized* 3.5-approximation algorithm. We note that the Discretized algorithm developed in the previous section can be applied towards solving the CDC problem. Accordingly, in *specific instances* when the NWST problem can be solved optimally (e.g. using integer programming), a centralized 2-approximate solution for the CDC problem can be obtained.

The key insight is that the CDC problem can be viewed as an extended variant of the STP-MSP. Namely, given a set of RNs (terminals) distributed in the plane, place the smallest set of MBNs (Steiner points) such that the RNs and MBNs form a connected network. Additionally, RNs must be *leaves in the tree*, and edges connecting them to the tree must be of length at most r. The other edges in the tree must be of at most R.

For the Discretization algorithm to apply, we need to make the following modifications. First, in the definition of the

TABLE II
TIME COMPLEXITY (# OF ROUNDS), LOCAL COMPUTATION COMPLEXITY, AND APPROXIMATION RATIO OF THE DISTRIBUTED GDC ALGORITHMS ($C(n)$ IS THE COMPLEXITY OF A DECISION 1-CENTER ALGORITHM).

Algorithm	*Time Complexity*	*Local Computation Complexity*	*In-Strip Approximation Ratio*
MOAC	$O(1)$	$O(\log n)$	3
SCR	$O(n)$	$O(\log n)$	2
MAS[7]	$O(1)$	$O(C(n))$	2
SCD	$O(n)$	$O(C(n)\log n)$	1.5

vertex set V, M_{cover} should be replaced with the set of RNs, N. Second, V_1 and V_2 should now be defined with respect to the pairwise intersections of radius r circles drawn around each of the RNs. Finally, in the definition of the edge set E, RNs should only have edges to vertices in V within distance r, and no two RNs should have an edge between them. With these modifications, it can be shown that if $R \geq 2r$ and $\Delta \leq R/6$, the Discretization algorithm is a $2\beta_{NWST}$-approximation algorithm for the overall CDC problem.

IX. PERFORMANCE EVALUATION

We now briefly discuss the tradeoffs between the complexities and approximation ratios of the GDC algorithms, and evaluate via simulation the GDC algorithms in a mobile environment. We also focus on the CDC problem and compare results obtained by the Discretization Approach to those obtained by decomposing the problem. Additional results can be found in [26].

Table II shows the complexities and approximation ratios of the *distributed* GDC algorithms. It can be seen that there are clear tradeoffs between decentralization and approximation. These tradeoffs are further demonstrated by simulation. Fig. 10 presents simulation results for a network with mobile RNs. The Random Waypoint mobility model is used, wherein RNs continually pick a random destination in the plane and move there at a random speed in the range $[V_{min}, V_{max}]$, where $V_{min} = 10m/s$ and $V_{max} = 30m/s$. We used a plane of dimensions $600m \times 600m$ and set $r = 100m$. The figure shows the average number of MBNs used over a $500s$ time period as a function of the number of RNs. Each data point is an average of 10 instances (each instance was simulated over $1000s$ from which the first $500s$ were discarded).

Next we compare solutions of the CDC problem obtained by the decomposition framework to joint solutions obtained by the Discretization algorithm. Fig. 11 depicts a random example of 10 RNs distributed in a $1000m \times 1000m$ area.[8] The communication ranges of the RNs and the MBNs are $r = 100m$ and $R = 200m$, respectively. In the decomposition framework, we used an optimal disk cover (obtained by integer programming) and the 3-approximation STP-MSP algorithm from [2]. The Discretization algorithm uses the NWST approximation algorithm from [18]. In this example, the joint

[7]The approximation ratio of the MAS algorithm holds when the algorithm is in steady state.

[8]We deliberately selected a small number of RNs in order to demonstrate a partitioned network that requires Relay MBNs.

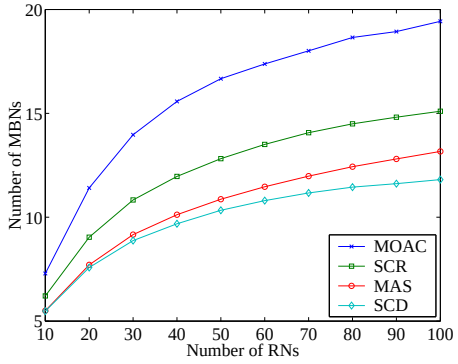

Fig. 10. The average number of Cover MBNs used by GDC algorithms over a time period of $500s$

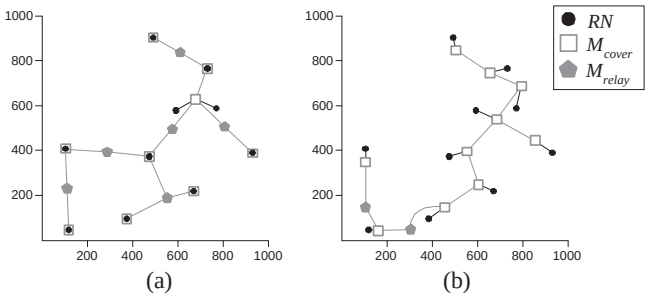

Fig. 11. An example comparing solutions obtained by (a) an optimal Disk Cover and the STP-MSP algorithm from [2], and (b) the Discretization algorithm using an NWST algorithm [18]

solution requires 12 MBNs while the decomposition based solution requires 15 MBNs .

Fig. 12 presents similar results for a general case with the same parameters (area, r, and R). The Decomposition framework used the SCD algorithm along with the MST algorithm [19] and along with the Modified MST-based algorithm [2]. Each data point is averaged over 10 random instances. It can be seen that the joint solution provides a significant performance improvement (about 25% for a large number of RNs). Yet, while the decomposition framework uses distributed algorithms, the joint solution is centralized. Thus, a reasonable compromise could be to place the Cover MBNs in a distributed manner and to place the Relay MBNs by a centralized Discretization Approach.

X. CONCLUSIONS

The architecture of a hierarchical Mobile Backbone Network has been presented only recently. Such an architecture can significantly improve the performance, lifetime, and reliability of MANETs and WSNs. In this paper, we concentrate on placing and mobilizing backbone nodes, dedicated to maintaining connectivity of the regular nodes. We have formulated the Mobile Backbone Nodes Placement problem as a Connected Disk Cover problem and shown that it can be decomposed into two subproblems. We have proposed a number of *distributed* algorithms for the first subproblem (Geometric Disk Cover), bounded their worst case performance, and studied their performance under mobility via simulation. As a byproduct, it has been shown that the approximation ratios of algorithms presented in [9] and [13] are 6 and 2. A new approach for the

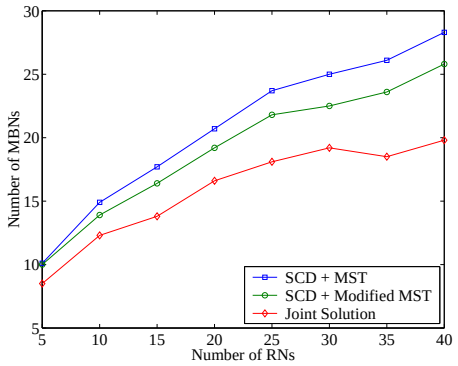

Fig. 12. Number of MBNs as a function of the number of RNs computed by: (i) the decomposition approach using the SCD with the MST-based [19] algorithms, (ii) the decomposition approach using the SCD with the modified MST-based [2] algorithm, and (iii) the Discretization algorithm

solution of the second subproblem (STP-MSP) and of the joint problem (CDC) has also been proposed. We have demonstrated via simulation that when it is used to solve the CDC problem in a centralized manner, the number of the required MBNs is significantly reduced.

This work is the first approach towards the design of distributed algorithms for construction and maintenance of a Mobile Backbone Network. Hence, there are still many open problems to deal with. For example, moving away from the strip approach may be beneficial. Moreover, there is a need for distributed algorithms for the STP-MSP, capable of dealing with Cover MBNs mobility. A major future research direction is to generalize the model to other connectivity constraints and other objective functions. Finally, an important future research direction is to address the problem when the number of available MBNs is fixed.

APPENDIX

Proof of Lemma 4: Consider a single strip S, whose width is αD. Since the RNs are distributed in the plane according to a two dimensional Poisson process, the horizontal (x-coordinate) distance between RNs is exponentially distributed with mean $1/\lambda\alpha D$. Thus, the expected distance to the location of the first disk is $E[T_1] = 1/\lambda\alpha D$ (see Fig. 13). Once a disk is placed, the expected distance between the end of its coverage and the start of the next disk is denoted by $E[T']$. Due to the memoryless property of the exponential random variable, $E[T'] = 1/\lambda\alpha D$. Therefore, the expected number of disks used within a strip (denoted by $E[|SCR|_S]$) is the total length of the strip (less the initial space) divided by the expected distance between the start of one disk and the start of another. Namely, assuming that $L >> 1/\lambda\alpha D$

$$E[|SCR|_S] = \frac{L - \frac{1}{\lambda\alpha D}}{\sqrt{1-\alpha^2}D + \frac{1}{\lambda\alpha D}} \approx \frac{\lambda\alpha D L}{\lambda\alpha\sqrt{1-\alpha^2}D^2 + 1},$$

The expected number of disks used in the plane is $E[|SCR|_S]$ multiplied by the number of strips. ■

Proof of Lemma 5: In order to lower bound the expected number of disks required by the optimal solution, we divide the plane into horizontal strips of width q separated by vertical distances D. We first find a lower bound on the expected number of optimal disks required to cover RNs in a q-width strip S (denoted by $E[|OPT|_S]$). Within such a strip, the area covered by each OPT disk is at most a rectangle of size $q \times D$. Using a similar argument to that of Lemma 4, it can be seen that once a disk is placed, the expected distance between the end of its coverage area and the start of the next disk is $1/\lambda q$. Assuming that $L >> 1/\lambda q$, $E[|OPT|_S]$ is at least the strip length divided by the expected distance between the start of one disk and the start of another. Namely, $E[|OPT|_S] \geq L/(D + 1/(\lambda q))$.

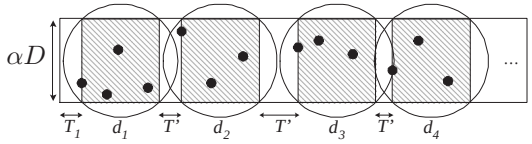

Fig. 13. Probabilistic analysis of the SCR algorithm within a strip

Since the distance between q-width strips is D, it is impossible for an OPT disk to cover RNs from multiple strips. Moreover, since there may be RNs between the strips, there will be a need for more OPT disks than the ones used to cover the q-width strips. Therefore, the expected number of OPT disks required in order to cover only RNs in the q-width strips is a lower bound on the expected number of OPT disks required for the whole plane. Such a bound can be found by multiplying $E[|OPT|_S]$ by the number of q-width strips, i.e.,

$$E[|OPT|] \geq \left(\frac{L}{D + \frac{1}{\lambda q}}\right) \cdot \left(\frac{K\alpha D}{D+q}\right). \quad (2)$$

To find the tightest possible lower bound, we select q so as to maximize $E[|OPT|]$. Setting $q = \sqrt{1/\lambda}$ achieves this, and yields the result. ■

Proof of Corollary 1: We derive the maximum value of (1) by differentiating with respect to λ, obtaining

$$\beta_{SCR}\,|_{\lambda=\lambda_{max}} \leq \frac{\alpha\sqrt{1-\alpha^2}+1}{\alpha\sqrt{1-\alpha^2}}. \quad (3)$$

For $\frac{1}{2} \leq \alpha < 1$, (3) is minimized when $\alpha = 1/\sqrt{2}$, at which point it attains a value of 3. ■

Proof of Lemma 10: In the CFS algorithm, each Relay MBN u in T_{OPT} is replaced by a Relay MBN u' in T (steps 4 and 7). For each edge connecting a pair of Relay MBNs in T_{OPT}, at most one additional MBN is added in T (w' in step 13). Since T_{OPT} is a tree, there can be at most $|M^*_{relay}| - 1$ such edges. Therefore, the total number of Relay MBNs in T is, $|M^T_{relay}| \leq |M^*_{relay}| + |M^*_{relay}| - 1 < 2|M^*_{relay}|$. ■

Proof of Theorem 5: Let the number of Relay MBNs in T_{OPT} and T be $|T_{OPT}| = |M^*_{relay}|$ and $|T| = |M^T_{relay}|$, respectively. Recall that in the Discretization algorithm, the Cover MBNs in G were assigned a weight of 0 and the other nodes were assigned a weight of 1. Let T^{NWST}_{OPT} be the optimal (minimum weight) Node-Weighted Steiner Tree (NWST) in G and denote its weight by $|T^{NWST}_{OPT}|$. Due to Lemma 9 when $\Delta \leq R/7$, T is a feasible solution to the NWST problem in

G. Therefore, and due to Lemma 10,

$$|T_{OPT}^{NWST}| \leq |T| \leq 2|T_{OPT}|. \quad (4)$$

In Step 5 of the Discretization algorithm, the NWST problem in G is solved by a β_{NWST} approximation algorithm. We denote the obtained solution by T_{ALGO} and denote the number of Relay MBNs in this solution by $|T_{ALGO}|$. From (4) we get that $|T_{ALGO}| \leq \beta_{NWST}|T_{OPT}^{NWST}| \leq 2\beta_{NWST}|T_{OPT}|$. ■

REFERENCES

[1] D.J. Baker and A. Ephremides, "The architectural organization of a mobile radio network via a distributed algorithm," *IEEE Trans. Commun.*, vol. 29, no. 11, pp. 1694-1701, 1981.

[2] D. Chen, D.-Z. Du, X.-D. Hu, G.-H. Lin, L. Wang, and G. Xue, "Approximations for steiner trees with minimum number of steiner points," *J. Global Optimization*, vol. 18, no. 1, pp. 17-33, 2000.

[3] X. Cheng, D.-Z. Du, L. Wang, and B. Xu, "Relay sensor placement in wireless sensor networks," To appear in *ACM/Springer WINET*, 2008.

[4] B. Das, R. Sivakumar, and V. Bharghavan, "Routing in ad-hoc networks using a virtual backbone," in *Proc. IEEE ICCCN'97*, Sept. 1997.

[5] M. Franceschetti, M. Cook, and J. Bruck, "A geometric theorem for approximate disk covering algorithms," *Technical Report ETR035*, Caltech, 2001.

[6] J. Gao, L. Guibas, J. Hershberger, L. Zhang, and A. Zhu, "Discrete mobile centers," *Disc. Comput. Geom.*, vol. 30, no. 1, pp. 45-63, 2003.

[7] M.R. Garey and D.S. Johnson, *Computers and Intractibility*. Freeman, 1979.

[8] M. Gerla and J. T. Tsai, "Multicluster, mobile, multimedia radio network," *ACM/Kluwer Wireless Networks*, vol. 1, no. 3, pp. 255-265, 1995.

[9] T. Gonzalez, "Covering a set of points in multidimensional space," *Inf. Proc. Letters*, vol. 40, no. 4, pp. 181-188, 1991.

[10] D. Gu, G. Pei, M. Gerla, and X. Hong, "Integrated hierarchical routing for heterogeneous multi-hop networks," in *Proc. IEEE MILCOM'00*, 2000.

[11] S. Guha and S. Khuller, "Improved methods for approximating node weighted steiner trees and connected dominating sets," *Inf. Comput.*, vol. 150, no. 1, pp. 57-74, 1999.

[12] H. Gupta, S. Das, and Q. Gu, "Connected sensor cover: self organization of sensor networks for efficient query execution," in *Proc. ACM MOBIHOC'03*, June 2003.

[13] J. Hershberger, "Smooth kinetic maintenance of clusters," in *Proc. ACM SoCG'03*, June 2003.

[14] J. Hershberger and S. Suri, "Finding tailored partitions," *J. Algorithms*, vol. 12, no. 3, pp. 431-463, 1991.

[15] D.S. Hochbaum and W. Maass, "Approximation schemes for covering and packing problems in image processing and VLSI," *J. ACM*, vol. 32, no. 1, pp. 130-136, 1985.

[16] H. Huang, A. Richa, and M. Segal, "Approximation algorithms for the mobile piercing set problem with applications to clustering in ad-hoc networks," *ACM/Kluwer MONET*, vol. 9, no. 2, pp. 151-161, 2004.

[17] A. Kansal, M. Rahimi, D. Estrin, W.J. Kaiser, G. J. Pottie, and M. B. Srivastava, "Controlled mobility for sustainable wireless sensor networks," in *Proc. IEEE SECON'04*, Oct. 2004.

[18] P. Klein and R. Ravi, "A nearly best-possible approximation algorithm for node-weighted steiner trees," *J. Algorithms*, vol. 19, no. 1, pp. 104-115, 1995.

[19] G. Lin and G. Xue, "Steiner tree problem with minimum number of steiner points and bounded edge-length," *Inf. Proc. Letters*, vol. 69, no. 2, pp. 53-57, 1999.

[20] M. Lu, J. Wu, M. Cardei, and M. Li, "Energy-efficient connected coverage of discrete targets in wireless sensor networks," in *Proc. ICCNMC'05*, Aug. 2005.

[21] J. Luo and J.-P. Hubaux, "Joint mobility and routing for lifetime elongation in wireless sensor networks," in *Proc. IEEE INFOCOM'05*, 2005.

[22] J.S. Provan, "Convexity and the steiner tree problem," *Networks*, vol. 18, no. 1, pp. 55-72, 1988.

[23] R. Rao and G. Kesidis, "Purposeful mobility for relaying and surveillance in mobile ad hoc sensor networks," *IEEE Trans. Mob. Comp.*, vol. 3, no. 3, pp. 225-231, 2004.

[24] I. Rubin, A. Behzad, R. Zhang, H. Luo, and E. Caballero, "TBONE: a mobile-backbone protocol for ad hoc wireless networks," in *Proc. IEEE Aerospace Conf.*, Mar. 2002.

[25] A. Segev, "The node-weighted steiner tree problem," *Networks*, vol. 17, no. 1, pp. 1-17, 1987.

[26] A. Srinivas, G. Zussman, and E. Modiano, "Construction and maintenance of a mobile backbone for wireless networks," *MIT/LIDS Technical Report #2676*, Oct. 2005. available at http://web.mit.edu/modiano/www/papers/T11.pdf.

[27] C. Swamy and A. Kumar, "Primal-dual algorithms for connected facility location problems," *Algorithmica*, vol. 40, no. 4, pp. 245-269, 2004.

[28] J. Tang, B. Hao, and A. Sen, "Relay node placement in large scale wireless sensor networks," *Computer Commun.*, vol. 29, no. 4, pp. 490-501, 2006.

[29] K. Xu, X. Hong, and M. Gerla, "Landmark routing in ad hoc networks with mobile backbones," *J. Parallel Distrib. Comput.*, vol. 63, no. 2, pp. 110-122, 2003.

PLACE PHOTO HERE

Anand Srinivas received his Ph.D. from MIT in 2007. He received his Masters of Science in EECS from MIT in 2004, and his Bachelors of Applied Science in Computer Engineering from the University of Toronto in 2001. In 2004 he also received a Masters of Science in Aerospace Engineering from MIT. His research interests include reliability and energy-efficiency in wireless ad-hoc networks, routing and network optimization, graph theory, and the design of efficient algorithms. He is currently working for Airvana Inc., Chelmsford, MA.

PLACE PHOTO HERE

Gil Zussman received the B.Sc. degree in Industrial Engineering and Management and the B.A. degree in Economics from the Technion – Israel Institute of Technology in 1995 (both *summa cum laude*). He received the M.Sc. degree (*summa cum laude*) in Operations Research from Tel-Aviv University in 1999 and the Ph.D. degree in Electrical Engineering from the Technion in 2004. Between 1995 and 1998, he served as an engineer in the Israel Defense Forces and between 2004 and 2007 he was a Postdoctoral Associate at MIT. He is currently an Assistant Professor at the Department of Electrical Engineering in Columbia University. His research interests are in the area of wireless networks. Gil received the Marie Curie Outgoing International Fellowship, the Fulbright Fellowship, the IFIP Networking 2002 Best Student Paper Award, and the OPNETWORK 2002 and ACM SIGMETRICS/IFIP Performance 2006 Best Paper Awards.

PLACE PHOTO HERE

Eytan Modiano received his B.S. degree in Electrical Engineering and Computer Science from the University of Connecticut at Storrs in 1986 and his M.S. and Ph.D. degrees, both in Electrical Engineering, from the University of Maryland, College Park, MD, in 1989 and 1992 respectively. He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Post Doctoral Fellow during 1992-1993. Between 1993 and 1999 he was with the Communications Division at MIT Lincoln Laboratory. He joined the MIT faculty in 1999, where he is an Associate Professor in the Department of Aeronautics and Astronautics and in the Operations Research Center. His research is on communication networks and protocols with emphasis on satellite, wireless, and optical networks. He is an Associate Editor for Communication Networks for IEEE Transactions on Information Theory and for The International Journal of Satellite Communications. He served as a guest editor for IEEE JSAC, Computer Networks, and the Journal of Communications and Networks. He had been the Technical Program co-chair for Wiopt 2006, Infocom 2007, and MOBIHOC 2007.