Joint node placement and assignment for throughput optimization in mobile backbone networks

Anand Srinivas and Eytan Modiano
Communications and Networking Research Group
Massachusetts Institute of Technology
Cambridge, MA 02139
{anand3, modiano}@mit.edu

***Abstract*—We study the novel hierarchical architecture of Mobile Backbone Networks. In such networks, a set of Mobile Backbone Nodes (MBNs) are deployed to provide an end-to-end communications capability for the Regular Nodes (RNs). In this work, we address the joint problem of placing a fixed number K MBNs in the plane, and assigning each RN to exactly one MBN. We formulate and solve two problems under a general communications model. The first is the Maximum Fair Placement and Assignment (MFPA) problem in which the objective is to maximize the throughput of the minimum throughput RN. The second is the Maximum Throughput Placement and Assignment (MTPA) problem, in which the objective is to maximize the aggregate throughput of the RNs. Our main result is a novel optimal polynomial time algorithm for the MFPA problem for fixed K. For a restricted version of the MTPA problem, we develop an optimal polynomial time algorithm for $K \leq 2$. We also develop two heuristic algorithms for both problems, including an approximation algorithm for which we bound the worst case performance loss. Finally, we present simulation results comparing the performance of the various algorithms developed in the paper.**

I. Introduction

Wireless Sensor Networks (WSNs) and Mobile Ad Hoc Networks (MANETs) can operate without any physical infrastructure (e.g. base stations). However, it has been shown that it is sometimes desirable to construct a *backbone* over which reliable end-to-end communication can take place [6],[7]. In particular, if some of the nodes are more capable than others, these nodes can be dedicated to providing the backbone. Such networks, termed *Mobile Backbone Networks*, have been recently been studied in [15],[19],[16].

Based on [15] and [19], a Mobile Backbone Network was defined in [16] as composed of two types of nodes. The first type includes static or mobile nodes (e.g. sensors or MANET nodes) with limited capabilities. These nodes are refered to as *Regular Nodes* (RNs). The second type includes mobile nodes with superior communication, mobility, and computation capabilities as well as greater energy resources. These nodes are termed *Mobile Backbone Nodes* (MBNs). The main purpose of the MBNs is to provide a mobile infrastructure facilitating network-wide communication.

An implicit assumption in previous formulations of the Mobile Backbone Network construction problem is that an arbitrary number of MBNs are available for deployment, and the goal is to minimize the number actually deployed. Indeed, this problem formulation was given in [16] as the *Connected Disk Cover* (CDC) problem. Specifically, the CDC problem aims to place the *minimum* number of MBNs such that (i) All RNs are covered by at least one MBN, and (ii) The MBNs form a connected network. In many scenarios however, a more appropriate (and perhaps realistic) assumption would be that the number of available MBNs is *fixed* a-priori, and the objective is to do the "best we can" with these fixed resources.

Note however, that the CDC-type formulation for MBN placement arises very naturally given the assumption of a discrete communications model, such as the "disk" connectivity model. In such a model, two nodes can communicate if they are within some fixed range, and cannot otherwise. However, while the disk model is a good first-order communications model, a more realistic model would account for the fact that the data rate at which two nodes can reliably communicate is actually a continuous function of the received Signal-to-Interference-and-Noise Ratio (SINR). The SINR in turn, depends on the wireless channel conditions and underlying PHY/MAC protocols (i.e. the System model). In this paper and for the specific context of Mobile Backbone Networks, we distill these issues into the following general model: The "throughput" achieved by an RN transmitting to its assigned MBN is a *decreasing* function of (i) The distance between the RN and MBN, and (ii) The total number of RNs assigned to that MBN. The idea is that first factor models the loss due to wireless propagation, and the second models loss due to interference caused by multiple RNs trying to access a single MBN. We elaborate further on the mathematical specifics of the model, as well as provide examples in section III.

With the above communications model, we are able to reformulate the backbone construction problem in a manner significantly different from previous formulations, and thereby requiring significantly different solution methodologies. In particular, we consider the joint problem of *placing* a fixed number of MBNs, and *assigning* each RN to exactly one MBN, such that a throughput objective is maximized. We consider two objective functions, yielding two separate prob-

This work was supported by NSF grant CCR-0325401, by ONR grant N000140610064, and by a grant from Draper Laboratory.

lems. The first is to maximize the throughput of the minimum throughput RN, which we term the Maximum Fair Placement and Assignment (MFPA) Problem. The second is to maximize the aggregate system throughput (i.e. sum of the throughputs achieved by each RN), which we term the Maximum Throughput Placement and Assignment (MTPA) problem.

It should be noted that in contrast to previous backbone construction problem formulations, the MFPA/MTPA involve a non-trivial assignment component. Specifically, a solution needs to balance assigning RNs to their closest MBNs and not assigning too many RNs to any particular MBN. Thus for the overall problems, not only do K MBNs need to be placed at arbitrary locations on the plane, but once placed there are K^N different RN to MBN assignments, among which the optimal one must be chosen, where N is the number of RNs.

Despite this, we are able to develop an optimal polynomial time algorithm for the MFPA problem for fixed K. We also develop an optimal solution for a restricted version of the MTPA problem for $K \leq 2$. As will be described later, the key lies in exploiting certain geometric properties of the placement portion of the problem, and certain combinatoric structure for the associated assignment subproblem. We also develop approximation and heuristic algorithms for both problems.

As a final point, to our knowledge the joint placement and assignment problems considered in this paper have not been addressed before. Thus the primary goal of this paper is to provide a theoretical framework and develop basic optimal solutions. We leave the development of more efficient, distributed and mobility-handling algorithms for future work.

This paper is organized as follows. In Section II we review related work and in Sections III and IV we formulate the problem and give illustrative examples. Section V presents an optimal solution for the MFPA problem. In section VI, we discuss solutions for a restricted version of the MTPA problem. In section VII, we present approximation and heuristic algorithms for both problems. Finally, in section VIII we evaluate the performance of the algorithms via simulation.

II. Related Work

The Mobile Backbone Architecture was originally presented in [15],[19] (and references therein). In their work, they assume the RNs and MBNs are already placed, and a-priori form a connected network. Thus the focus of their work relates to developing system-level protocols for routing, scheduling, MBN election, etc. Our approach differs in that we focus on the more fundamental problem of, given a set of arbitrarily located RNs, how to place the MBNs and assign RNs to MBNs, such that a network performance objective is optimized.

Somewhat along these lines is the work of [16], in which the specific network performance objective is end-to-end connectivity. In particular, they formulate the *Connected Disk Cover* (CDC) problem, which aims to place the minimum number of MBNs such that (i) All RNs are covered by at least one MBN, and (ii) The placed MBNs form a connected network. They present various approximation algorithms towards solving the CDC problem. However, that an arbitrary number of MBNs

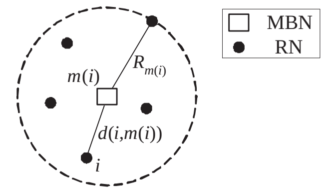

Fig. 1. Example of a Cluster.

are assumed to be available, as well as the fact that a disk connectivity model is assumed, makes both the CDC problem formulation as well as the associated solution methodology considerably different from that of this paper.

Given the more general communications model considered in this paper, the closest related work is actually in regards to base station selection/placement for cellular and indoor wireless systems, e.g. [4],[17],[18],[12],[10]. However, there are several aspects which differentiate our work from the work in this area. First, the major levers of optimization in our work are both the MBN (e.g. base station) placement *and* the RN to MBN assignments. By contrast, much of the cellular work use trivial solutions to the assignment problem (e.g. assign each RN to the nearest MBN) and optimize via base station placement/selection and/or power control. Another key difference is that practical considerations for cellular base station placement usually a-priori restricts the set of possible locations to a discrete set of candidates. This restriction typically results in solution methodologies along the lines of simple heuristics, or large scale optimization tools (e.g. MILP, GA, etc). In contrast, we develop optimal combinatorial algorithms for the joint node placement and assignment problems of this paper.

Finally, the basic idea of clustering nodes to form a hierarchical architecture has been extensively studied in the context of wireless networks (e.g. [5],[7]). Yet, 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. [16],[11],[14]).

III. Problem Formulation

We consider a set of N Regular Nodes (RNs), distributed in the plane and assume that a set of $K < N$ Mobile Backbone Nodes (MBNs) are to be deployed. We denote the set of RNs by $P = \{1, 2, \ldots, N\}$ and the set of MBNs by $M = \{m_1, m_2, \ldots, m_K\}$. For every RN i, let $m(i)$ denote the MBN to which i has been assigned, (e.g $m(i) = k$ if i is assigned to m_k), and let $d(i, m(i))$ represent the distance between them. In general, let $d(i, j)$ represent the distance between nodes i and j. Next, for every MBN m_k, let P_k denote the set of RNs assigned to it. We refer to the tuple of an MBN and its assigned RNs as a *cluster*. For cluster k corresponding to (m_k, P_k), we define the *cluster radius* R_k as, $R_k = max_{j \in P_k} d(j, m_k)$. The number of RNs assigned to MBN m_k, $|P_k|$, is refered to as the *cluster size*. An example of a cluster is shown in Fig. 1. For the communications model, we assume that the throughput of an RN i transmitting to its assigned MBN $m(i)$

is some *function* $H\big(d(i,m(i)),\ |P_{m(i)}|\big)$, that is *decreasing* in both it's arguments. As mentioned earlier, the dependence of $H()$ on $d(i,m(i))$ models wireless propagation loss, and the dependence on $|P_{m(i)}|$ reflects loss due to interference at MBN $m(i)$. Note that in this communications model we assume that RNs from different MBNs do not interfere with each other, e.g. different clusters operate on different frequencies.

To gain some intuition about the form $H()$ could take, consider the following two system examples: (i) Slotted Aloha-based, and (ii) CDMA-based. In the Slotted Aloha based model, we assume that all RNs assigned to an MBN m_k transmit within a slot with equal probability, $1/|P_k|$. Additionally, we associate a "distance penalty" proportional to $d^{-\alpha}$ for an RN located a distance d away from m_k, where α represents the path loss exponent. This could, for example, reflect extra coding that needs to be used in order to deal with the propagation loss. The resulting throughput of a node i in this system is therefore simply the probability that exactly one RN transmits in a slot, multiplied by the distance penalty, i.e.,

$$\begin{aligned} TP_{SA}(i) &= \tfrac{1}{|P_{m(i)}|}\left(1-\tfrac{1}{|P_{m(i)}|}\right)^{|P_{m(i)}|-1}\left(\tfrac{1}{d(i,m(i))^{\alpha}}\right) \\ &\approx \tfrac{1}{e\cdot|P_{m(i)}|\cdot d(i,m(i))^{\alpha}} \\ &\triangleq H_{SA}\Big(d(i,m(i)),|P_{m(i)}|\Big) \end{aligned} \tag{1}$$

where we have left out most of the constants for simplicity, and we use the approximation that $(1-1/x)^{x-1} \rightarrow 1/e$ even for small values of $x \geq 1$. Note that (1) is of the desired form for $H()$, i.e. decreasing in both $d(i,m(i))$ and $|P_{m(i)}|$. Next, consider a CDMA-based system in which power control is employed. Specifically, in order to combat the near-far problem, all RNs assigned to an MBN $m(i)$ equalize their received power (equal to 1, for simplicity) at $m(i)$ to that of the farthest away RN. Thus the throughput achieved by every RN within a cluster is the same, and is proportional to its Signal-to-Interference-and-Noise Ratio (SINR) at $m(i)$, i.e.,

$$\begin{aligned} TP_{cdma}(i) &= \tfrac{1/R^{\alpha}_{m(i)}}{(1/R^{\alpha}_{m(i)})(|P_{m(i)}|-1)+\eta} \\ &= \tfrac{1}{|P_{m(i)}|+\eta\cdot R^{\alpha}_{m(i)}-1} \\ &\triangleq H_{cdma}\Big(R_{m(i)},|P_{m(i)}|\Big) \end{aligned} \tag{2}$$

where η represents the noise at MBN $m(i)$, and $R_{m(i)}$ the radius of cluster $m(i)$. Again, note the form of the throughput function is as desired, since it is decreasing in both distance and cluster size. For the purpose of intuition, we will carry these two examples throughout the paper, whenever possible directly applying to them the general results that we derive.

We now give a precise formulation for the two problems that will be addressed in this paper: (i) The Maximum Fair Placement and Assignment (MFPA) Problem and (ii) Maximum Throughput Placement and Assignment (MTPA) problem.

Problem MFPA: Given a set of RNs (P) distributed in the plane, place K MBNs (M) and assign each RN i to exactly one MBN $m(i)$ such that the following is maximized:

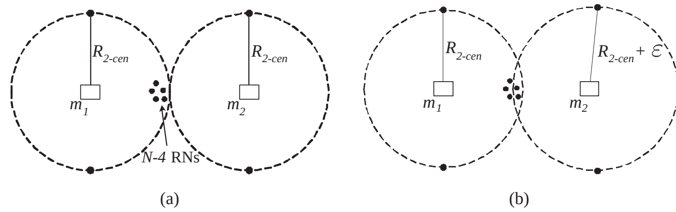

Fig. 2. $K = 2$ MFPA example. (a) 2-Center Solution. (b) Optimal Solution.

$$\min_{i\in P} TP(i) = \min_{i\in P}\left\{H\Big(d(i,m(i)),|P_{m(i)}|\Big)\right\} \tag{3}$$

Problem MTPA: Given a set of RNs (P) distributed in the plane, place K MBNs (M) and assign each RN i to exactly one MBN $m(i)$ such that the following is maximized:

$$\sum_{i\in P} TP(i) = \sum_{i\in P} H\Big(d(i,m(i)),|P_{m(i)|}\Big) \tag{4}$$

As a final point, we enforce the following additional conditions on the $H()$ function,

1) $H(R,X) > 0, \forall R \geq 0, X \geq 1$.
2) $H(R,X) < \infty, \forall R \geq 0, X \geq 1$ (only for MTPA)

Notice that condition (2) is needed for the general MTPA problem as stated above to be well defined. Otherwise, any solution in which an MBN is placed on top of an RN could yield infinite aggregate throughput (i.e. artificially exploiting the so-called "near-field" effect). Since $K < N$, this is not an issue for the MFPA problem, i.e. the worst case throughput RN cannot have an MBN on top of it.

IV. Illustrative Examples

In this section we attempt to give some additional intuition regarding the complexity of the joint placement and assignment problems addressed in this paper. To begin, consider a 1 MBN instance of the MFPA problem. With just one MBN, we immediately note that the assignment portion of the problem is trivial (i.e. all N RNs are assigned to the one MBN). Furthermore, the associated placement portion of the problem can be solved optimally by placing the single MBN so as to minimize the farthest distance from any RN. This is precisely the well known 1-center problem[1], for which several efficient polynomial time algorithms exist [1]. Applying one of these algorithms solves the 1 MBN MFPA problem optimally.

Next, consider the 2 MBN example illustrated in Fig. 2. Fig. 2(a) shows the MFPA solution if we simply apply a 2-center algorithm, and assign RNs to their nearest MBN. As shown, the worst case RN attains a throughput of $H(R_{2-cen}, n-2)$ in this case, where R_{2-cen} is the 2-center radius. However, by increasing the radius of the second cluster by a small amount, i.e. enough to enclose half of the $n-4$ RNs clustered

[1]In general, the K-center problem places K MBNs such that the farthest distance from any RN to its nearest MBN is minimized.

together, the optimal solution can potentially increase the worst case RNs' throughput to $H(R_{2-cen}+\epsilon, \frac{n}{2})$; this is shown in Fig. 2(b). Clearly depending on the exact form of $H()$, this improvement can be quite significant. As demonstrated in this simple example, even if we are given a placement of the MBNs, the assignment problem is non-trivial, as it may potentially be beneficial to assign RNs to farther away MBNs.

Thus the main difficulty of the MFPA and MTPA problems for $K > 1$ can be summarized as follows. First, there are an infinite number of potential locations for the MBNs (i.e. anywhere on the plane). Second, for any particular placement of K MBNs, there are K^N different assignments of RNs to MBNs (i.e. each RN can be assigned to one of K MBNs).

V. MFPA Solution

The key to our approach in solving the MFPA problem is to decouple the placement and assignment problems in a way that does not affect the optimality of the resulting decoupled solution. We start with the following observation and lemma. The observation applies to any feasible MFPA solution, and follows from the fact that the overall minimum throughput RN must be the minimum throughput RN in its own cluster.

Observation 1: Let RN i have minimum throughput among all RNs, and let $m(i)$ be its assigned MBN. Then, the throughput of i can be expressed as a function of its cluster's radius and size, i.e. $TP(i) = H(R_{m(i)}, |P_{m(i)}|)$.

Lemma 1: Let $P_1^*, P_2^*, \ldots, P_K^*$ represent the optimal MFPA assignments of RNs to MBNs $m_1, m_2, \ldots, m_K$ respectively. Then, there exists an optimal solution to the overall MFPA problem in which the MBNs are placed at the *1-center locations* of $P_1^*, P_2^*, \ldots, P_K^*$.

Proof: Consider an optimal solution to the MFPA problem in which the MBNs are *not* placed at the 1-center locations of $P_1^*, \ldots, P_K^*$. Next, consider the solution obtained by moving all of the MBNs to their respective 1-center locations. By definition of the 1-center, doing this never increases the radius of any of the K clusters. Therefore, since the cluster sizes $|P_1^*|, \ldots, |P_K^*|$ are fixed, then by observation 1 the throughput of the worst case throughput RN does not decrease. ■

The consequence of the above lemma is that for the placement problem, the finite space of 1-center locations contains at least one solution of optimal cost. Additionally, the associated cluster radii of each of the K clusters are by definition 1-center radii. Thus as a first step, we have reduced the search space from an infinite number of locations on the plane, to a finite set of 1-center locations (with associated 1-center radii).

At first glance, the total number of 1-center locations/radii might seem prohibitively large and thus our reduction of limited use. For example, every subset of RNs has an associated 1-center location and radius, and there are 2^N subsets. However, it turns out that all of these locations/radii come from a relatively small (i.e. polynomial in N) set of candidates. To show this, we need the following fact, illustrated in Fig. 3, regarding the 1-center of a set of RNs P [13],

Fact 1: The unique 1-center location and radius of a set of RNs P, denoted $1C(P)$ and $R(P)$, is *defined* by either:

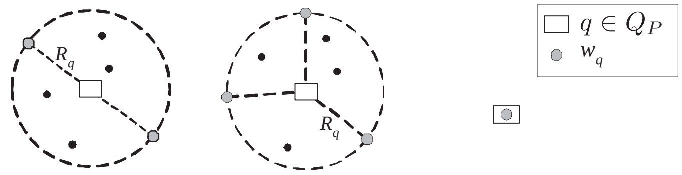

Fig. 3. Illustration of the forms of 1-center (location,radius) tuples. (a) Midpoint of a pair of points. (b) Circumcenter of a triplet of points. (c) On top of a single point.

1) A *pair* of RNs $i, j \in P$. If this is the case, then $1C(P)$ is situated at the *midpoint* of i,j, and $R(P) = d(i,j)/2$.
2) A *triplet* of RNs $i, j, k \in P$ that form an *acute triangle*. If this is the case, then $1C(P)$ is situated at the *circumcenter*[2] of $\{i, j, k\}$ and $R(P)$ is the *circumradius*.
3) A *single* RN $i \in P$. This is the degenerate case where $P = \{i\}$ is a singleton set, and $1C(P)$ is situated on i itself, and $R(P) = 0$.

Indeed, the actual 1-center $(1C(P), R(P))$ tuple has minimum $R(P)$ such that all RNs are within distance $R(P)$ of the location $1C(P)$. Let Q_P denote the full set of candidate 1-center locations, as described in fact 1 with respect to the original set of RNs P. Note that since each $q \in Q_P$ is defined by either 1,2 or 3 RNs in P, it follows that that Q_P has cardinality at most $\binom{N}{1}+\binom{N}{2}+\binom{N}{3}$. Additionally, as described in Fact 1 and shown in Fig. 3, for each $q \in Q_P$, we associate R_q to denote the *1-center radius* of a cluster whose 1-center location is q, and the set w_q to denote the set of *defining RNs* for q. Note that though several locations in the set Q_P may be coincident, all w_q's are distinct. We now state the following lemma, which follows by construction of Q_P and fact 1.

Lemma 2: The 1-center (location, radius) tuple of *any subset* $T \subseteq P$ corresponds to some (q, R_q) tuple, $q \in Q_P$.

Combining lemmas 1 and 2 and Fact 1, we can conclude that restricting our placements of MBNs to the set Q_P still allows us to find the optimal solution to the overall MFPA problem. Moreover, we can restrict ourselves to solutions whereby if an MBN m_k is placed at location $q \in Q_P$, all of the RNs assigned to it must be within distance R_q, i.e. $d(i, m_k) \leq R_q$, $\forall i \in P_k$. Otherwise, by Fact 1 q cannot be the unique 1-center location of P_k, i.e. there must exist some other location $q' \in Q_P$ that is the actual 1-center location of P_k, with corresponding 1-center radius $R_{q'}$. As per lemma 1, moving m_k to location q' cannot decrease the MFPA objective.

For clarity, we illustrate the exhaustive search over all placements among locations in Q_P as the high-level framework shown below. Let $m_1^*, \ldots, m_k^*$ denote the optimal locations of the K MBNs, $m^*(1), \ldots, m^*(N)$ the optimal RN to MBN assignments, and U^* the associated optimal cost.

Up to this point, we have not discussed the assignment subproblem, which we need to solve as a subroutine in step

[2]For a triplet of RNs, the *circumcenter* is the center of the circle that has all three RNs on its boundary. The radius of this circle is the *circumradius*.

Algorithm 1 High-Level Optimal MFPA Framework

1: **initialize** $U^* = -\infty$
2: **create** the set Q_P by enumerating over all defining subsets of size $1, 2$ and 3 of P.
3: **for** all $\binom{|Q_P|}{K}$ placements of K MBNs $m_1, \ldots, m_K$ **do**
4:     **if** all RNs are within R_j of at least 1 MBN m_j in current MBN placement **then**
5:         **calculate** the optimal MFPA assignments $m(i), \forall i \in P$, given the current MBN placement and subject to the constraint that $m(i) = k$ only if $d(i, m_k) \le R_k$. Let U represent the corresponding worst case RN throughput.
6:         **if** $U > U^*$ **then**
7:             **set** $U^* \leftarrow U$, **update** $m^*(i), m_k^*, \forall i \in P, k \in K$
8: **return** U^*, $m_1^*, \ldots, m_K^*$ and $m^*(1), \ldots, m^*(N)$

5 of the high-level framework. It turns out that the specific methodologies used to solve this problem for $K = 2$ and $K > 2$ are quite different, as we describe below.

A. $K = 2$ MFPA Assignment Subproblem

With the placement locations and radii fixed, for $K = 2$ the resulting MFPA assignment subproblem turns out to be easy to solve. In this situation, as depicted in Fig. 4(a), we define $C(1)$ and $C(2)$ as the sets of RNs that lie *exclusively* within radius R_1 and R_2 of MBNs m_1 and m_2 respectively. Similarly, let $C(1, 2)$ denote the "common set" of RNs that lie within the radii of both m_1 and m_2. The main idea is that since the radii are fixed, RNs in $C(1), C(2)$ *must* be assigned to m_1, m_2 respectively. Moreover, in assigning the remaining RNs in $C(1, 2)$, it is only the *number* assigned to each MBN that affects the MFPA objective. Thus we can search over the $|C(1, 2)| + 1$ different possibilities and pick the best one.

The worst case computational complexity of the overall MFPA algorithm for $K = 2$ is therefore $O(N^7)$. This follows from the fact that $|Q_P| \le N^3$ and we need to solve $\binom{|Q_P|}{2}$ assignment problems, each of which takes $O(N)$ time.

B. General K MFPA Assignment Subproblem

The MFPA assignment subproblem for $K > 2$ is significantly more difficult than for $K \le 2$. To get a sense of the additional difficulty, consider the 2 vs. 3 MBN example illustrated in Fig. 4. For 2 MBNs m_1, m_2, there is only one type of "common set" of RNs, i.e. $C(1, 2)$, yielding at most $O(N)$ ways to assign different numbers of RNs to each MBN.

For $K > 2$ MBNs, the number of ways to divide different numbers of RNs *within a single common set* generalizes to $O(N^{K-1})$. Yet, the real difficulty is that for $K > 2$, there can potentially be many types of common sets. For example, in Fig. 4(b), RNs in the set $C(1, 2, 3)$ can be assigned to any of the 3 MBNs, whereas RNs in $C(2, 3)$ can only be assigned to either m_2 or m_3. Thus, the total number of ways the RNs within all of these different common sets can be divided among K MBNs is $O((N^{K-1})^I)$, where I represents the number of distinct common sets. Observing that each MBN location and radius represents a circular region, we can actually bound I by K^2 [2]. This results in a total complexity of $O(N^{K^3})$ to enumerate all possible assignments. While still polynomial

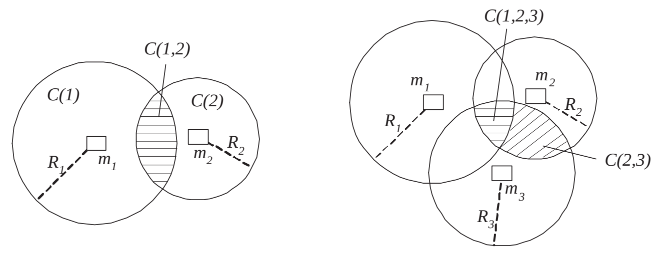

(a) (b)

Fig. 4. (a) $K = 2$ vs. (b) $K = 3$ examples of assignment subproblem.

in N, spending this complexity for *each* of the $O(N^{3K})$ assignment subproblems yields an overall algorithm definitely outside the realm of practicality (e.g. even for $K = 3$).

With a more practical solution desired, we now develop an optimal algorithm for the general K MFPA assignment subproblem that is polynomial in *both* K and N. To this end, we start by formulating the MFPA assignment subproblem using a mathematical programming notation. Define indicator variables x_{ij} to equal to 1 if RN i is assigned to MBN m_j. Next, define indicator constants z_{ij} to be equal to 1 if $d(i, m_j) \le R_j$. The resulting formulation can be written as,

$$\max \quad \min_{j \in M} \quad H\big(R_j, \sum_{i \in P} x_{ij}\big) \tag{5}$$

$$s.t. \quad \sum_{j \in M} x_{ij} = 1, \forall i \in P \tag{6}$$

$$x_{ij} \le z_{ij}, \forall i \in P, j \in M \tag{7}$$

$$x_{ij} \in \{0, 1\} \tag{8}$$

where constraints (6) ensure that every RN is assigned to exactly 1 MBN, constraints (7) that we only make valid assignments, and constraints (8) integrality of the final assignment. Defining the *increasing* function $F() = 1/H()$, since $H() > 0$, we can re-write the objective function in (5) as,

$$\min \quad \max_{j \in M} \quad F\big(R_j, \textstyle\sum_i x_{ij}\big) \tag{9}$$

Applying one more transformation, we have (avoiding re-writing constraints (6)-(8) for brevity),

$$\min \quad W \tag{10}$$

$$s.t. \quad \sum_{i \in P} x_{ij} \le g(W; R_j), \forall j \in M \tag{11}$$

where we have used the common trick of converting a *minimax* objective function into a simple *min* objective function by introducing an extra real valued variable W and moving the *max* part of the objective function into the constraints. We define $g(W; R_j)$ to be the *inverse* with respect to $\sum_i x_{ij}$ of $F(R_j, \sum_i x_{ij})$, i.e.,

$$g\Big(F(R_j, \textstyle\sum_i x_{ij}) \ ; \ R_j\Big) = \textstyle\sum_i x_{ij} \tag{12}$$

which we assume exists. This assumption is justified since $F()$ is monotonically increasing, and therefore constitutes a one-to-

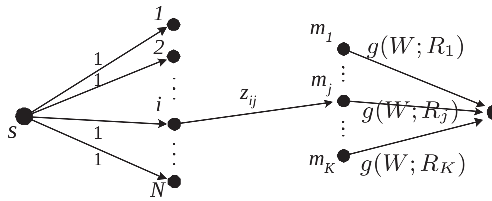

Fig. 5. Construction of the Flow Graph $G = (V, E, C)$ for a given W.

one (in $\sum_i x_{ij}$) function. As an example, for the Slotted Aloha $H()$ given in (1) we have that $g(W; R_j) = W/(e \cdot R_j^{\alpha})$.

At this point, we note that the above *optimization* problem can be solved by way of solving a series of *feasibility* problems (e.g. fix W, and see if there exist x_{ij}'s that satisfy constraints (11), (6)-(8)).. Consider the following lemma, which follows from the observations that: (i) $g(W; R_j)$ is monotonically increasing since it is the inverse of a monotonically increasing function, and (ii) Constraints (8) induce the right hand side of constraints (11) to be integral for an optimal solution.

Lemma 3: The optimal W^* must satisfy $g(W^*; R_j) \in \mathbb{Z}$. That is, $g(W^*; R_j)$ must be integral.

We can combine this lemma with the fact that there are at most $K \cdot N$ distinct integer feasible values for $g(W^*; R_j)$. Specifically, for each R_j (of which there are K), W^* can be one of $F(R_j, b), b = 1, \ldots, N$. Therefore, we can *exactly* find the optimal W^* by solving $K \cdot N$ feasibility problems.

The remaining question is: Given a value for W, how can we efficiently find (or not find) an assignment of x_{ij}'s that answers the feasibility question? To this end, we will now show that the feasibility problem can be transformed into a classical graph problem, *Integer Max-Flow*, for which several efficient *polynomial time* algorithms exist [3].

The Integer Max-Flow problem is defined as follows: We are given a *flow graph* $G = (V, E, C)$, where C defines an integer set of *capacities* c_{ij} on each edge $(i, j) \in E$, and a *source* vertex s and a *sink* vertex t, $s, t \in V$. The objective is to assign positive integer *flows* f_{ij} on each each edge $(i, j) \in E$ such that the aggregate flow from s to t, equal to $\sum_j f_{sj}$, is maximized. The f_{ij}'s must obey the following constraints:

1) $f_{ij} \leq c_{ij}, \forall (i, j) \in E$ (capacity constraints)
2) $\sum_i f_{ij} - \sum_k f_{jk} = 0, \forall j \in V \backslash \{s, t\}$ (flow conservation)
3) $\sum_i f_{is} = \sum_j f_{tj} = 0$ (source and sink property)

Returning to our problem, we start by constructing a flow graph $G = (V, E, C)$ in the following manner, depicted in Fig. 5. Let $P \in V$ represent a set of vertices corresponding to each RN, and similarly $M \in V$ for the MBNs. Next, define source and sink vertices $s, t \in V$. Next, define N source edges (s, i) with capacities $c(s, i) = 1$, $\forall i \in P$. Next, define edges between nodes (i, j) with capacities $c(i, j) = z_{ij}$, $\forall i \in P, j \in M$. Finally, define K sink edges (j, t) with capacities $c(j, t) = g(W; R_j)$, $\forall j \in M$. At this point we run a Max-Flow algorithm to find the maximum (integral) flow between s and t in G. Given the Max-Flow solution, we interpret a non-zero flow $f_{ij} = 1$ on an edge of type $(i, j), i \in P, j \in M$ to mean that in the assignment solution RN i should be assigned to MBN j. Our main result lies in the following lemma, which we only prove for one direction; the converse holds by construction.

Lemma 4: For a given W, the MFPA assignment subproblem is feasible if and only if the Max-Flow from s to t has value equal to N.

Proof: Assume an integer max-flow of value N is found. To show this corresponds to a feasible solution to the MFPA assignment subproblem, it suffices to show that all of the constraints (11), (5), (6) are satisfied. Constraints (5) are satisfied since if the max-flow is equal to N, it must mean that all source edges carry a flow of 1. Thus by flow conservation, each RN (at the endpoint of each of the source edges) must be assigned to exactly 1 MBN. Next, note that constraints (6) are satisfied since if edge $(i, j), i \in P, j \in M$ has non-zero flow across it, then by construction it's capacity, which is equal to z_{ij} must be equal to 1. Finally, constraints (11) are satisfied since if more than $g(W; R_j)$ RNs are assigned to any MBN m_j, this would correspond to edge (j, t) having a greater flow than it's assigned capacity. ∎

The preceding lemma gives us the final piece of the puzzle needed in order to construct an efficient algorithm for the MFPA assignment subproblem. The algorithm is given below.

Algorithm 2 Fixed K MFPA assignment algorithm

initialize $W^* \leftarrow \infty$
for $k = 1$ to K **do**
  for $b = 1$ to N **do**
    set $W \leftarrow F(R_k, b)$
    if $W < W^*$ **then**
      construct flow graph $G = (V, E, C)$ as follows:
       set $V \leftarrow P \bigcup M \bigcup \{s, t\}$
       set $E \leftarrow E \bigcup \{(s, i)\}$, $c(s, i) \leftarrow 1, \forall i \in P$
       set $E \leftarrow E \bigcup \{(i, j)\}$, $c(i, j) \leftarrow z_{ij}, \forall i \in P, j \in M$
       set $E \leftarrow E \bigcup \{(j, t)\}$, $c(j, i) \leftarrow \lfloor g(W; R_j) \rfloor, \forall j \in M$
      solve $s - t$ Max-Flow on G. Let f_{ij} be the flows on each edge (i, j) and F_{max} the max-flow value.
      if $F_{max} = N$ **then**
        set $m(i) \leftarrow j$ if $f_{ij} = 1$, $\forall i \in P, j \in M$
        set $W^* \leftarrow W$
return $W^*, m(1), \ldots, m(N)$

We conclude the section by noting that the best Integer Max-Flow algorithm has running time $O(KN^2 \log N)$ [8]. Therefore, the algorithm depicted above has $O(K^2 N^3 \log N)$ complexity. The result is a *worst case* complexity $O(N^{3K+3} \log N)$ algorithm for the *fixed* K MFPA problem. As will be shown in section VIII, this algorithm can be applied to solve instances with relatively small K and N.

VI. MTPA Solution

It turns out the general MTPA problem as formulated in 4 is significantly more difficult to optimally solve than the MFPA problem. For example, consider the MTPA problem for $K = 1$ MBN (i.e. ignore the assignment subproblem). At first glance it would seem like the MTPA problem looks like the well known 1-median/Fermat-Weber problem (numerically solvable in polynomial time [1]), in which one seeks to place the MBN in the location that minimizes the sum of the distances to each RN. However, the general MTPA objective is actually to maximize the sum of *arbitrary decreasing functions* of each of the distances; the difference is quite substantial. For example, consider a very simple decreasing function $H(d_i) = 1/(d_i+\gamma)$, where d_i represents the distance from RN i to the placed MBN and γ some positive constant. Clearly minimizing $\sum_i d_i$ achieves a significantly different objective from maximizing $\sum_i 1/(d_i+\gamma)$ (for which to our knowledge no optimal algorithm exists).

Thus we consider a restriction on the general MTPA problem, in which we enforce the condition on the $H()$ function that all RNs within a cluster get the same throughput, which is a function of the cluster radius and size, i.e.,

$$TP'(i) = H\Big(R_{m(i)}, |P_{m(i)}|\Big), \forall i \in P \tag{13}$$

The reasoning behind this particular restriction is two-fold. First, the above expression yields a lower bound on the general MTPA objective, i.e. since $H(d(i,m(i)), |P_{m(i)}|) \geq H(R_{m(i)}, |P_{m(i)}|), \forall i \in P$. It is therefore still useful to optimize. Second, this approach allows us to heavily leverage the discussion we have evolved through this paper for the MFPA problem. To start, for $K = 1$ the 1-center algorithm optimally solves the restricted version of the MTPA problem.

For $K > 1$, we note that Observation 1 along with lemmas 1-2 all apply to the restricted MTPA problem. Therefore, the high-level framework in section V solves the placement portion of the problem. Additionally, for $K = 2$ the simple MFPA assignment algorithm in section V-A also solves the restricted MTPA assignment subproblem, as long as the appropriate (i.e. MTPA) objective function is used.

For $K > 2$, the brute force approach discussed in the beginning of section V-B applies to the restricted MTPA problem. However, the fixed K MFPA assignment algorithm does not solve the fixed K restricted MTPA assignment problem.

VII. Lower Complexity Heuristics

Although the algorithms developed so far in this paper find optimal solutions in polynomial time, their complexity is still prohibitively high unless both K and N are quite small. For example for $K = 3$, $N = 35$, the running time of the optimal MFPA algorithm was 3 hours on a Pentium 2.4GHz computer.

Thus in this section, our goal is to develop *suboptimal* approaches that have significantly less running time than the optimal approach, but still perform comparably well. We will discuss 2 such approaches: (i) An approximation algorithm that is based on cutting down the number of candidate MBN placements, and (ii) A simple and fast heuristic algorithm, but with no worst case performance guarantee. For the most part, the discussion applies to both the MFPA and restricted MTPA problems. For brevity, we will describe the algorithms in the context of the MFPA problem, noting any key issues specific to the restricted MTPA when appropriate.

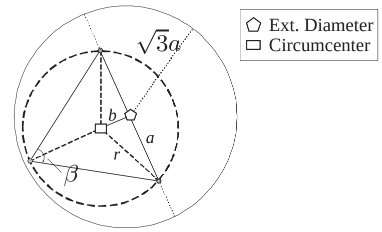

Fig. 6. Extended Diameter-type vs. Circumcenter-type placement.

A. Extended Diameter Algorithm (EDA)

As discussed in section V, the complexity of the optimal MFPA algorithm is dominated by the number of (optimality-preserving) possible placements, $\binom{|Q_P|}{K} = O(N^{3K})$. Indeed, the set Q_P is of size $O(N^3)$ due to having to consider all possible locations/radii corresponding to circumcenters/circumradii of triplets of RNs (see Fact 1). If we did not consider such "circumcenter-type" locations, but instead only looked at locations defined by (i) pairs of RNs (i.e. "diameter-type") and (ii) single RNs (i.e. "singular type"), the number of possible placements would immediately reduce to $O(N^{2K})$. This is the main idea behind the approach in this section.

Recall that in the high-level framework, we only consider placements at locations $q \in Q_P$, and assignments such that if an RN i is assigned to MBN m_k located at $q \in Q_P$, then $d(i, m_k) \leq R_q$. We denote such solutions as *valid*. However, an issue that comes up when we remove circumcenter-type locations from Q_P is that a valid solution may not even exist (e.g. consider 3 RNs that form a equilateral triangle).

To compensate for this, we define *extended-diameter* type locations, shown in Fig. 6, whose locations are the same as the original diameter-type locations, but whose associated radii are $\sqrt{3}$ times larger. Let Q'_P denote the set of all extended-diameter and singular-type locations with respect to a set of RNs P. Note that a direct analog with lemma 2 applies, i.e. Q'_P contains all extended-diameter and singular-type locations (with associated radii) with respect to any subset of RNs $T \subseteq P$. The next lemma ensures that placements among locations in Q'_P are guaranteed to contain a valid MFPA solution.

Lemma 5: For a set of RNs P, there exists a valid solution to the MFPA problem with placements at locations in Q'_P.

Proof: To prove the lemma, we need to show that for every circumcenter-type location/radii tuple in Q_P, there exists an extended-diameter-type location/radii tuple in Q'_P that covers the same set of RNs. To this end, consider some circumcenter-type placement, and the extended-diameter location corresponding to the midpoint of the *longest* side (of length $2a$) of the acute triangle formed by the circumcenters' defining RNs. The situation is depicted in Fig. 6. Let b be the distance between the extended-diameter and circumcenter locations. Next, let r denote the circumradius. By the triangle inequality, we know that the distance between the extended diameter location and any RN covered by the circumcenter

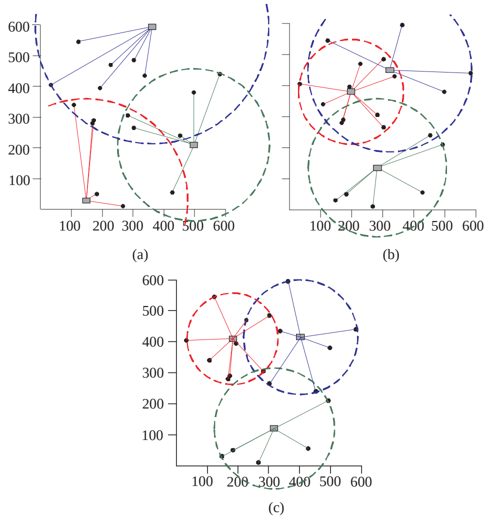

Fig. 7. Example of solutions found for a $K = 3$, $N = 20$ instance of the MFPA problem with the Slotted Aloha throughput function. (a) Unoptimized Farthest Point Heuristic (FPH) (b) Unoptimized Extended Diameter Algorithm (c) Optimal MFPA algorithm.

placement is at most $b + r$. Therefore, we have that,

$$\begin{aligned} b+r &= r+\sqrt{r^2-a^2} = \frac{a}{sin\beta}+\sqrt{\frac{a^2}{sin^2\beta}-a^2} \\ &\leq \frac{2a}{\sqrt{3}}+\frac{a}{\sqrt{3}} = \sqrt{3}a \end{aligned} \tag{14}$$

where we have used a geometric property of circumcenters that $r = \frac{a}{sin\beta}$. Additionally, we have used the observation that since the defining triangle is acute and since the extended-diameter location under consideration is defined by the longest edge of the triangle, that $\pi/3 \leq \beta \leq \pi/2$. ■

Finally, define the *extended-diameter 1-center* of a set of RNs P as the location in Q'_P that minimizes the maximum distance from any RN in P. We now state the analog of lemma 1 applied to this context, whose proof follows from lemma 5.

Lemma 6: Let $P_1^*, P_2^*, \ldots, P_K^*$ represent the optimal assignments of RNs to MBNs $m_1, m_2, \ldots, m_K$ respectively. Then, there exists a solution to the overall MFPA problem in which MBNs are placed at the extended-diameter 1-centers of $P_1^*, P_2^*, \ldots, P_K^*$. Also, the objective value of this solution is at least $H(\sqrt{3}R^*, |P^*|)$, where R^* and $|P^*|$ represent the worst case cluster radius and size of the optimal solution.

We define the Extended-Diameter Algorithm (EDA) for the fixed K MFPA as well as the $K = 2$ MTPA problem, as basically the optimal algorithms described earlier, with Q'_P used in place of Q_P. The only difference is a final optimization step, in which after the suboptimal extended-diameter placement is decided, we move each of the MBNs to the actual 1-center location of their assigned RNs.

By the preceding discussion, the EDA algorithm is a $\frac{H(\sqrt{3}R^*,|P^*|)}{H(R^*,|P^*|)}$-approximation algorithm for the MFPA. For path loss exponent α=2 this ratio evaluates to $1/3$ for both the Slotted-Aloha and CDMA throughput functions. The worst case running time of the algorithm is $O(N^5)$ for K=2, and $O(N^{2K+3} \log N)$ for fixed K>2.

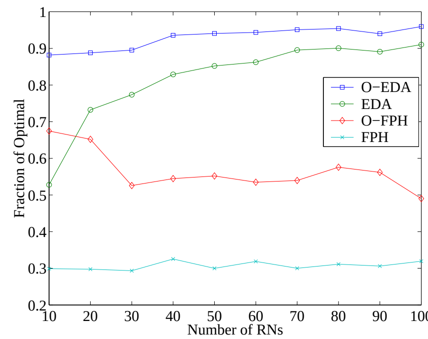

Fig. 8. Average case simulation for $K = 2$ for the MFPA problem with Slotted Aloha throughput function

B. Farthest Point Heuristic (FPH)

This next algorithm is simply an adaptation of Gonzalez's Farthest Point Heuristic (FPH) [9], with an additional optimization step tailored to our setting. The algorithm works as follows: Initialize the algorithm by placing an MBN on top an arbitrary RN, and assign all RNs to this MBN. Place the next MBN on top of the RN farthest from its assigned MBN, and re-assign RNs to their nearest MBN. Repeat the previous step until all K MBNs are placed. The above placement can be "optimized" by moving each MBN to the 1-center location of its assigned RNs. The running time of the unoptimized version of this algorithm is $O(N \log K)$, and using a practical 1-center algorithm [1], the optimized version takes $O(KN \log N)$ time.

VIII. Simulation Results

In this section we compare the performance of the various algorithms presented in this paper via simulation. To this end, we begin with an example of running the algorithms on a single $K = 3$ MBNs, $N = 20$ RNs, MFPA instance, shown in Fig. 7. We assume the RNs are randomly distributed in a 600×600 plane, and we use the Slotted-Aloha $H()$ throughput function given in (1), with $\alpha = 2$.

As can be seen, the optimal solution achieves the ideal balance between lightly loading clusters of large radii vs. heavily loading clusters of smaller radii. By contrast, the FPH solution potentially creates enormous radius clusters. Moreover, since nothing intelligent is done by the FPH regarding the assignment problem (i.e. just assign RNs to their closest MBN), the large radius clusters can also get heavily loaded. The EDA does better, in that even though its cluster radii are larger than optimal, it intelligently assigns RNs in a way that achieves optimal load balancing among the placed clusters.

Figs. 8 and 9 show an average case plot for varying numbers of RNs, and $K = 2$ and $K = 3$ MBNs. The parameters are the same as for the previous scenario, and we average each data point over 20 random instances. We present the average ratio of the throughput achieved by the suboptimal algorithms as compared to that of the optimal algorithms described in sections V. In both figures, we can notice that the optimization step significantly improves the performance of the heuristics.

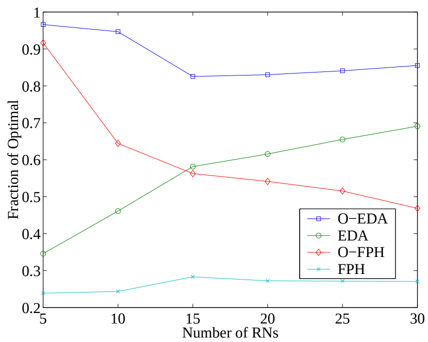

Fig. 9. Average case simulation for $K = 3$ for the MFPA problem with Slotted Aloha throughput function

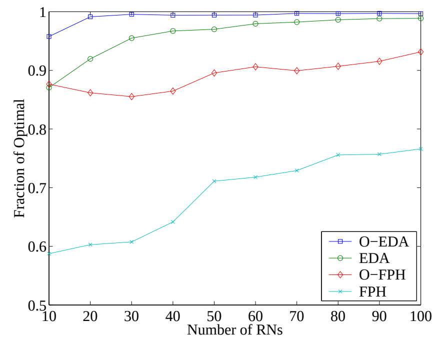

Fig. 10. Average case simulation for $K = 2$ for the MTPA problem with CDMA throughput function

However, as exhibited by the poor performance of both the optimized and regular FPH, the optimization step can only help insofar as lowering the cluster radius if possible; it cannot make up for already-made poor assignment decisions.

Finally, Fig. 10 shows an average case simulation for the $K = 2$ MTPA problem with the CDMA throughput objective function from (2). We set $\eta = 10^{-4}$ in order to normalize the SNR somewhat, and add 1 to the denominator so as to maintain $H() < \infty$ as mentioned in section III. Note that the O-EDA achieves aggregate throughput very close to optimal. In fact, all of the algorithms perform significantly better (relative to optimal) for the MTPA objective than for the MFPA objective, albeit with different $H()$ functions. Nevertheless, this would seem to indicate that the max-sum (i.e. MTPA) objective is less *sensitive* to suboptimal MBN placement/assignment than the max-min (i.e. MFPA) objective.

IX. Conclusion

The recently studied Mobile Backbone Network architecture can significantly improve the performance, lifetime and reliability of MANETs and WSNs. In this paper, we have focused on the key problem of how to jointly place the Mobile Backbone Nodes (MBNs), and assign every Regular Node to exactly one MBN. To this end, we have formulated two problems under a general communications model. The first is the Maximum Fair Placement and Assignment (MFPA) problem in which the objective is to maximize the throughput of the minimum throughput RN. The second is the Maximum Throughput Placement and Assignment (MTPA) problem, in which the objective is to maximize the aggregate throughput of the RNs. Our main result is a novel optimal polynomial time algorithm for the MFPA problem for fixed K. We have also provided an optimal solution for a restricted version of the MTPA problem for $K \leq 2$. We have developed two heuristic algorithms for both problems, including an approximation algorithm with bounded worst case performance loss. Finally, we have presented simulation results to evaluate the performance of the various algorithms developed in the paper.

To our knowledge the problems presented in this paper have not been considered before. Thus for this paper, our primary goal has been to provide a theoretical framework, as well as basic optimal solutions. Future work involves the development of more efficient, distributed and mobility-handling algorithms for both the MFPA and MTPA problems.

References

[1] P. Agarwal and M. Sharir, “Efficient Algorithms for Geometric Optimization,” *ACM Comput. Surveys*, 30, pp. 412-458, 1998.
[2] P. Agarwal, M. Pellegrini and M. Sharir, “Counting circular arc intersections,” *Siam J. Computing*, 22, pp. 778-793, 1993.
[3] R. Ahuja, T. Magnanti, and J. Orlin, “Network Flows,” Prentice Hall, New Jersey, 1993.
[4] E. Amaldi, A. Capone, and F. Malucelli, “Discrete models and algorithms for the capacitated location problems arising in UMTS network planning,” *ACM DIAL-M’01*, July 2001.
[5] D.J. Baker and A. Ephremides, “The Architectural Organization of a Mobile Radio Network via a Distributed Algorithm,” *IEEE Trans. Commun.*, 29, 11, pp. 1694-1701, 1981.
[6] S. Basagni, M. Mastrogiovanni, A. Panconesi, and C. Petrioli, “Localized protocols for ad hoc clustering and backbone formation: a performance comparison,” *IEEE Trans. on Parallel and Dist. Systems*, 17, 4, pp. 292-306, 2006.
[7] B. Das, R. Sivakumar, and V. Bharghavan, “Routing in Ad-Hoc Networks Using a Virtual Backbone,” *IEEE ICCCN’97*, 1997.
[8] H. Gabow, “Scaling algorithms for network problems,” *J. Comput. Syst. Sci.*, 31, pp. 148-168, 1985.
[9] T. Gonzalez, “Clustering to minimize the maximum intercluster distance,” *Theoretical Computer Science*, 38, pp. 293-306, 1985.
[10] S. Hanly, “An Algorithm for Combined Cell-Site Selection and Power Control to Maximize Cellular Spread Spectrum Capacity,” *IEEE J. Select. Areas Commun.*, 13, pp. 1332-1340, 1995.
[11] J. Luo and J.-P. Hubaux, “Joint Mobility and Routing for Lifetime Elongation in Wireless Sensor Networks,” *IEEE INFOCOM’05*, 2005.
[12] R. Mathar and T. Niessen, “Optimum positioning of base stations for cellular radio networks,” *Wireless Networks*, 6, pp. 421428, 2000.
[13] F. Preparata and M. Shamos, “Computational Geometry: An Introduction,” Spinger-Verlag, New York, 1985.
[14] R. Rao and G. Kesidis, “Purposeful Mobility for Relaying and Surveillance in Mobile Ad Hoc Sensor Networks,” *IEEE Tran. Mob. Comp.*, 3, 3, pp. 225-231, 2004.
[15] I. Rubin, A. Behzad, R. Zhang, H. Luo, and E. Caballero, “TBONE: a Mobile-Backbone Protocol for Ad Hoc Wireless Networks,” *Proc. IEEE Aerospace Conf.*, 6, 2002.
[16] A. Srinivas, G. Zussman, and E. Modiano, “Mobile Backbone Networks: Construction and Maintenance,” *ACM MOBIHOC’06*, May 2006.
[17] D. Stamatelos and A. Ephremides, “Spectral Efficiency and Optimal Base Placement for Indoor Wireless Networks,” *IEEE Journal on Selected Areas in Communications*, 14, 4, May 1996.
[18] K. Tutschku, “Demand-based Radio Network Planning of Cellular Mobile Communication Systems,” *IEEE INFOCOM’98*, 1998.
[19] K. Xu, X. Hong, and M. Gerla, “Landmark Routing in Ad Hoc Networks with Mobile Backbones,” *J. Parallel Distrib. Comput.*, 63, 2, pp. 110-122, 2003.