Divide et Impera is almost optimal for the bounded-hop MST problem on random Euclidean instances[0]

Andrea E. F. Clementi* Miriam Di Ianni* Massimo Lauria†
Angelo Monti† Gianluca Rossi*[1] Riccardo Silvestri†

August 30, 2005

Abstract

The d-Dim h-hops MST problem is defined as follows: Given a set S of points in the d-dimensional Euclidean space and $s \in S$, find a minimum-cost spanning tree for S rooted at s with height at most h. We investigate the problem for any constants h and $d > 0$. We prove the first non trivial lower bound on the solution cost *for almost all* Euclidean instances (i.e. the lower-bound holds with hight probability). Then we introduce an easy-to-implement, very fast divide et impera heuristic and we prove that its solution cost matches the lower bound.

1 Introduction

Given a positive integer h, an *h-tree* T is a rooted tree such that the number of hops (edges) in the path from the root to any other node is not greater than h. The *cost* of T, denoted as $\mathsf{cost}(T)$, is the sum of its edge weights. The *Minimum h-hops Spanning Tree* problem (h-hops MST) is defined as follows: Given a graph $G(V, E)$ with nonnegative edge weights and a node $s \in V$, find a minimum-cost h-tree rooted at s and spanning G. The h-hops MST problem and the related problem in which the constraint is on the tree diameter find applications in several areas: networks [4], distributed system design [22, 7], bit-compression for information retrieval [6].

The efficient construction of a (minimum) spanning tree of a communication network yields good protocols for *broadcast* and *anti-broadcast*[2] operations. The hop restriction limits the maximum number of links or connections in the communication paths between source and destination nodes: It is thus closely related to restricting the maximum delay transmission

*Dipartimento di Matematica, Università degli Studi di Roma "Tor Vergata". E-mail: {clementi,diianni,rossig}@mat.uniroma2.it.

†Dipartimento di Informatica, Università degli Studi di Roma "La Sapienza". Email: {massimo.lauria@mclink.net}, {monti,silver}@di.uniroma1.it.

[0]Research partially funded by the Information Society Technologies programm of the European Commission, Future and Emerging Technologies, under the IST-2001-33135 project CRESCCO

[1]Supported by the European Union under the IST FET Project CRESCCO.

[2]The anti-broadcast operation is also known in literature as *Accumulation* or *All-to-One* operation.

time of such fundamental communication protocols. The hop restriction finds another relevant application in the context of *reliability*: Assume that, in a communication network, link faults happen with probability p and that all faults occur independently. Then, the probability that a multi-hop transmission fails *exponentially* increases with the number of hops. Summarizing, a *fixed* bound on the maximum number of hops is sometimes a necessary constraint in order to achieve *fast* and *reliable* communication protocols.
For further motivations in studying the h-HOPS MST problem see [5, 11, 15, 24].

The h-HOPS MST problem is NP-hard even when the edge weights of the input graphs form a *metric* and $h = 2$ [1]. Algorithmic research on this issue thus aims to design and analyse efficient approximation algorithms.
Several previous works [5, 11, 24] focused on the 2-dimensional geometric version of the problem (2-DIM h-HOPS MST), i.e., when nodes are points of the Euclidean 2-dimensional space, the graph is complete, and the edge weights are the Euclidean distances. As for the case $h = 2$, polynomial-time, constant-factor approximation algorithms are given in [23, 8, 17, 14, 18] and a PTAS is provided in [3] for 2-DIM h-HOPS MST. We remark that all such approximation algorithms are not fast and/or easy-to-implement and, for $h \geq 2$, neither hardness results nor polynomial-time (exact) algorithms are known for the 2-DIM h-HOPS MST problem. Even more, for $h \geq 3$, no polynomial-time, constant-factor approximation algorithms are known.

Another series of papers have been devoted to evaluate and compare solutions for the 2-DIM h-HOPS MST problem returned by some heuristics on *random* planar instances by performing computer experiments [9, 11, 12, 21, 24]. Almost all such works adopt the *uniform input random model*, i.e., points are chosen independently and uniformly at random from a fixed square of the plane. The motivation on this input model is twofold: On one hand, the uniform distribution is the most suitable choice when *nothing* is known about the *real* input distribution or when the goal is to perform a preliminary study of the heuristic on arbitrary instances. On the other hand, uniform distribution well models important applications in the area of ad-hoc wireless and sensor networks: In such scenarios, once *base* stations are efficiently located, a large set of small wireless (mobile or not) devices are *well-spread* over a geographical region. Clearly, in these networks, efficient and reliable protocols for broadcast and accumulation is a primary goal [10].
We emphasize that no theoretical analysis is currently available on the expected performance of any efficient algorithm for the 2-DIM h-HOPS MST problem.

Our first result is a lower bound on the cost of any h-tree spanning a random set of points, i.e., a finite set of points chosen independently and uniformly at random from a fixed d-dimensional hypercube (d-cube).

Theorem 1 *Let $h, d \geq 1$ be constants. Let S be a random set of n points in a d-cube of side length L and let T be any h-tree spanning S. Then, it holds that*

$$\mathsf{cost}(T) = \begin{cases} \Omega\left(L \cdot n^{\frac{1}{h}}\right) & \textit{if } d = 1 \\ \Omega\left(L \cdot n^{1-\frac{1}{d}+\frac{d-1}{d^{h+1}-d}}\right) & \textit{otherwise} \end{cases} \quad \textit{with high probability.}$$

Here and in the sequel the term *with high probability* (in short, *w.h.p.*) means that the event holds with probability at least $1 - e^{-c \cdot n}$, for some constant $c > 0$. So, according to our input model, claiming that a given bound holds w.h.p. is equivalent to claiming that it holds *for almost all* inputs.

We then introduce a simple *Divide et Impera* heuristic h-PARTY. It makes a partition of the smallest d-cube containing S into *cells*. In each non-empty cell, it selects an *arbitrary* sub-root s' and connects s' to the root s; finally, it solves the non-empty cell sub-instances of the problem with $h-1$ hops, recursively. Choosing the size of the cells is the critical technical issue: This is solved by means of the lower bound in Theorem 1.

Theorem 2 *Let $h, d \geq 1$ be constants. Let S be a set of n points in a d-cube of side length L and let $s \in S$. For any h-tree T returned by h-PARTY on input (S, s), it holds that*

$$\mathsf{cost}(T) = \begin{cases} \mathrm{O}\left(L \cdot n^{\frac{1}{h}}\right) & \text{if } d = 1 \\ \mathrm{O}\left(L \cdot n^{1-\frac{1}{d}+\frac{d-1}{d^{h+1}-d}}\right) & \text{otherwise.} \end{cases}$$

Theorems 1 and 2 imply that, for any fixed h, h-PARTY returns a solution which is, with high probability, a constant factor approximation of the optimum. So, even though this fast algorithm provides no provably-good approximation in the *worst-case*, it works well on *almost-all* Euclidean instances.

h-PARTY is the first heuristic that works in $O(n)$ time and it can be thus efficiently applied to very large instances. In fact, the heuristic has been implemented and tested on instances of hundreds of thousands points [9].

Notice that, differently from Theorem 1, the bound in Theorem 2 holds for *any* Euclidean instance. It thus follows that random instances are those having the largest cost.

1.1 Related works

The 2-DIM 2-HOPS MST problem can be easily reduced to the classic Facility Location Problem on the plane. Indeed, the distance of the root from vertex i can be seen as the cost of opening a facility at vertex i. It thus follows that all the approximation algorithms for the latter problem apply to the 2-DIM 2-HOPS MST as well. In particular, the best result is the PTAS given by Arora et al in [3]. The algorithm works also in higher dimensions; however, it is based on a complex dynamic programming technique that makes any implementation very far to be practical. Several polynomial-time approximation algorithms for the METRIC 2-HOPS MST problem have been presented in the literature. Notice that, in [1] Alfandari and Paschos proved that METRIC 2-HOPS MST is MAX SNP-hard and, hence, PTAS cannot be found for this problem unless P = NP. The first constant factor approximation algorithm was given by Shmoys et al in [23], they presented a 3.16 approximation algorithm. After this, a series of constant factor approximation algorithms was published, see [8, 17, 14]. Currently, the best factor is 1.52 due to Mahdian et al [18]. All such algorithms make use of Linear Programming relaxations that yield not practically efficient implementations.

As for the general h-HOPS MST problem, Gouveia [12] and, successively, Gouveia and Requejo [13] provided and experimentally tested exact super-polynomial-time algorithms, based on the branch and bound technique. In [2, 16] a polynomial-time $O(\log n)$-approximation algorithm is given, but its time complexity is $n^{\mathrm{O}(h)}$. Voss in [24] presented a tabu-search heuristic for the h-HOPS MST problem, but the time complexity is very high when the graph is dense. In [21] heuristics based on Prim algorithm and on Evolutionary techniques have been experimentally tested. Finally, in [9] experimental tests have been performed on greedy heuristics and on h-PARTY.

2 Preliminaries

In the proof of our results we make use of the well-known Hölder inequality. We thus present it in the following convenient forms. Let x_i, $i = 1, \ldots, k$ be a set of k non negative reals and let $p, q \in \mathcal{R}$ such that $p \geq 1$ and $q \leq 1$. Then, it holds that:

$$\sum_{i=1}^{k} x_i^p \geq k \left(\frac{\sum_{i=1}^{k} x_i}{k} \right)^p; \tag{1}$$

$$\sum_{i=1}^{k} x_i^q \leq k \left(\frac{\sum_{i=1}^{k} x_i}{k} \right)^q. \tag{2}$$

3 The lower bound

Next lemma is the first lower bound on the cost of h-trees for general Euclidean instances.

Lemma 1 *Let $h, d \geq 1$ be constants. Let S be a set of points in a d-dimensional Euclidean space. Consider a partition of the space in d-cubes with the side length of each d-cube being l and let n_l be the number of the d-cubes containing points of S. For any h-tree T spanning S it holds that*

$$\mathsf{cost}(T) = \begin{cases} \Omega\left(l \cdot n_l^{1+\frac{1}{h}} \right) & \textit{if } d = 1 \\ \Omega\left(l \cdot n_l^{1+\frac{d-1}{d^{h+1}-d}} \right) & \textit{otherwise.} \end{cases}$$

Proof. We equivalently show that $\mathsf{cost}(T) = \Omega\left(L \cdot n^{1+\frac{1}{g(h)}} \right)$ where

$$g(h) = \begin{cases} d & \text{if } h = 1 \\ d \cdot g(h-1) + d & \text{otherwise.} \end{cases}$$

Notice that $1/g(h) = \frac{d-1}{d^{h+1}-d}$ if $d > 1$ and $1/g(h) = \frac{1}{h}$ if $d = 1$.
Let s be the point root of the tree T and consider a d-sphere centered at s and of radius $r = \Theta(l \cdot (n_l)^{\frac{1}{d}})$ such that the number n'_l of non-empty d-cubes outside the sphere is at least $\frac{n_l}{2}$. Finally let B be the set of points in these n'_l d-cubes.
The proof is by induction on the height h of the tree T. If $h = 1$, for each of the n'_l d-cubes, there is an edge in T of length at least r. This implies that

$$\mathsf{cost}(T) \geq r \cdot n'_l = \Omega\left(l \cdot n_l^{1+\frac{1}{d}} \right) = \Omega\left(l \cdot n_l^{1+\frac{1}{g(1)}} \right).$$

Let $h \geq 2$. Let $A = \{a_1, a_2, \ldots, a_{|A|}\}$ be the set of points whose father is at distance at least $\frac{r}{h}$ and let $\beta = 1 - \frac{1}{d} + \frac{1}{g(h)}$. Two cases may arise.
- *Case* $|A| \geq n_l^{\beta}$. Since there are at least $|A|$ edges of length $\frac{r}{h}$, it holds that

$$\mathsf{cost}(T) \geq \frac{r}{h} \cdot |A| = \Omega(l \cdot n_l^{\beta+\frac{1}{d}}) = \Omega\left(l \cdot n_l^{1+\frac{1}{g(h)}} \right).$$

- *Case* $|A| < n_l^{\beta}$. For every point x in B there is a path from x to the root s with at most h hops. Since the distance from x to s is at least r in the path there is at least one edge of length at least $\frac{r}{h}$. Hence we can partition the points in $A \cup B$ into $|A|$ subsets $A_1, A_2, \dots A_{|A|}$ where a point y is in A_i if a_i is the first point in A in the path from y to s. Notice that the points in the subsets A_i, $1 \le i \le |A|$, belong to disjoint subtrees $T_1, T_2, \dots, T_{|A|}$ of T where T_i is an $(h-1)$-tree rooted at a_i. Let $n_{l,i}$ be the number of d-cubes containing the points of T_i, $1 \le i \le |A|$. It holds that

$$
\begin{array}{rcll}
\mathsf{cost}(T) & \ge & \sum_{i=1}^{|A|} \mathsf{cost}(T_i) & \\
 & = & \Omega\left(\sum_{i=1}^{|A|} l \cdot {n_{l,i}}^{1+\frac{1}{g(h-1)}}\right) & \text{by inductive hypothesis} \\
 & = & \Omega\left(l \cdot |A| \cdot \left(\frac{\sum_{i=1}^{|A|} n_{l,i}}{|A|}\right)^{1+\frac{1}{g(h-1)}}\right) & \text{by the Hölder inequality} \\
 & = & \Omega\left(l \cdot |A|^{-\frac{1}{g(h-1)}} \cdot {n_l}^{1+\frac{1}{g(h-1)}}\right) & \text{since } \sum_{i=1}^{|A|} n_{l,i} \ge n_l' \ge \frac{n_l}{2} \\
 & = & \Omega\left(l \cdot {n_l}^{-\frac{\beta}{g(h-1)}+1+\frac{1}{g(h-1)}}\right) & \text{since } |A| < n_l^{\beta} \\
 & = & \Omega\left(l \cdot {n_l}^{1+\frac{g(h)-d}{d \cdot g(h-1) \cdot g(h)}}\right) & \\
 & = & \Omega\left(l \cdot {n_l}^{1+\frac{d \cdot g(h-1)}{d \cdot g(h-1) \cdot g(h)}}\right) & \text{since } g(h) = d \cdot g(h-1) + d. \\
 & = & \Omega\left(l \cdot {n_l}^{1+\frac{1}{g(h)}}\right) &
\end{array}
$$

The thesis follows. □

By applying the probabilistic method of *bounded differences* [19], we derive the following lower bound

Theorem 1 *Let $h, d \ge 1$. Let S be a random set of n points in a d-cube of side length L and let T be any h-tree spanning S. Then, it holds that*

$$
\mathsf{cost}(T) = \begin{cases} \Omega\left(L \cdot n^{\frac{1}{h}}\right) & \textit{if } d = 1 \\ \Omega\left(L \cdot n^{1-\frac{1}{d}+\frac{d-1}{d^{h+1}-d}}\right) & \textit{otherwise} \end{cases} \quad \textit{with high probability.}
$$

Proof. Let us partition the d-cube into n d-cubes, each of them with side length $l = \frac{L}{n^{\frac{1}{d}}}$. Let n_l be the number of non-empty d-cubes. Lemma 1 implies that

$$
\mathsf{cost}(T) = \begin{cases} \Omega\left(L \cdot n^{-1} \cdot {n_l}^{1+\frac{1}{h}}\right) & \text{if } d = 1 \\ \Omega\left(L \cdot n^{-\frac{1}{d}} \cdot {n_l}^{1+\frac{d-1}{d^{h+1}-d}}\right) & \text{otherwise} \end{cases}
$$

The theorem follows by noticing that, by applying the method of *bounded differences* [19], we have that $n_l \ge n/4$, with high probability. □

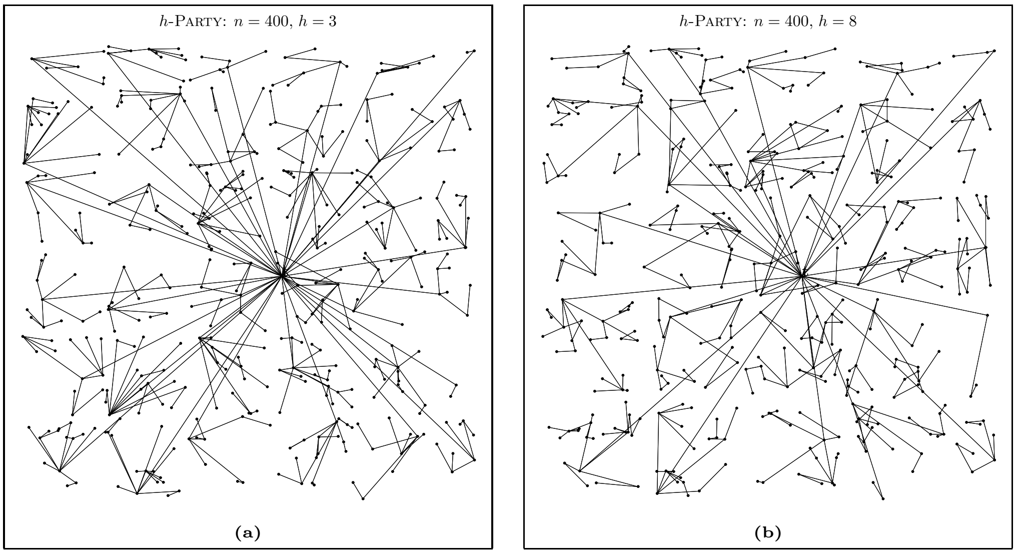

Figure 1: The trees yielded by the h-PARTY heuristics on the same random instance with 400 points and $h = 3, 8$.

procedure h-PARTY(S, s)
    if $h = 1$ **then** $T \leftarrow \{\{x, s\} | x \in S - \{s\}\}$;
    else begin
        $T \leftarrow \emptyset$;
        if $d = 1$ **then** $k \leftarrow \left\lfloor |S|^{\frac{1}{h}} \right\rfloor$;
        else $k \leftarrow \left\lfloor |S|^{1-\frac{1}{d}+\frac{d-1}{d^{h+1}-d}} \right\rfloor$;
        Let L be the side length of the smallest d-cube containing all points in S;
        Partition the d-cube into d-cubes of side length $\frac{L}{\left\lfloor k^{\frac{1}{d}} \right\rfloor}$;
        Let k' be the number of d-cubes and let S_i be the points of S in the i-th d-cube, $1 \leq i \leq k'$;
        for $i \leftarrow 1$ **to** k' **do**
            if $|S_i| \geq 1$ **then begin**
                choose a point s' in S_i;
                $T \leftarrow T \cup \{\{s', s\}\}$;
                if $|S_i| > 1$ **then** $T \leftarrow T \cup$ $(h-1)$-PARTY(S_i, s');
            end;
    end;
    output T

Figure 2: Algorithm h-PARTY.

4 The Divide-et-Impera heuristic

The heuristic h-PARTY is described in Figure 2 while its solution cost is proved in the following

Theorem 2 *Let $h, d \geq 1$ be constants. Let S be a set of n points in a d-cube of side length L and let $s \in S$. For any h-tree T returned by h-PARTY on input (S, s), it holds that*

$$\mathsf{cost}(T) = \begin{cases} \mathrm{O}\left(L \cdot n^{\frac{1}{h}}\right) & \text{if } d = 1 \\ \mathrm{O}\left(L \cdot n^{1-\frac{1}{d}+\frac{d-1}{d^{h+1}-d}}\right) & \text{otherwise.} \end{cases}$$

Proof. We equivalently show that $\mathsf{cost}(T) = O\left(L \cdot n^{1-\frac{1}{d}+\frac{1}{g(h)}}\right)$ where

$$g(h) = \begin{cases} d & \text{if } h = 1 \\ d \cdot g(h-1) + d & \text{otherwise} \end{cases}$$

Notice that, as in Lemma 1, $1/g(h) = \frac{d-1}{d^{h+1}-d}$ if $d > 1$ and $1/g(h) = \frac{1}{h}$ if $d = 1$. The proof is by induction on h. If $h = 1$ it is clear that $\mathsf{cost}(T) = O(L \cdot n)$.
For $h \geq 2$, let t be the number of non-empty d-cubes in the d-cube of size length L and $\{q_1, q_2, \ldots q_t\}$ be the set of points selected by the procedure in the t non-empty d-cubes, let T_i be the $(h-1)$-tree rooted in q_i and S_i be the set of points spanned by T_i, $1 \leq i \leq t$. By inductive hypothesis, we get $\mathsf{cost}(T_i) = O\left(\frac{L}{k^{\frac{1}{d}}} \cdot |S_i|^{1-\frac{1}{d}+\frac{1}{g(h-1)}}\right)$. We thus have that

$$\begin{aligned}
\mathsf{cost}(T) &= \textstyle\sum_{i=1}^{t} d(q_i, s) + \sum_{i=1}^{t} \mathsf{cost}(T_i) \\
&\leq L \cdot t + \textstyle\sum_{i=1}^{t} \mathsf{cost}(T_i) && \text{since } d(a_i, s) \leq L, 1 \leq i \leq t \\
&= O\left(L \cdot t + \textstyle\sum_{i=1}^{t} \frac{L}{k^{\frac{1}{d}}} \cdot |S_i|^{1-\frac{1}{d}+\frac{1}{g(h-1)}}\right) && \text{by inductive hypothesis} \\
&= O\left(L \cdot t + \frac{L}{k^{\frac{1}{d}}} \cdot t \cdot \left(\frac{\sum_{i=1}^{t} |S_i|}{t}\right)^{1-\frac{1}{d}+\frac{1}{g(h-1)}}\right) && \text{by the Hölder inequality} \\
&= O\left(L \cdot t + \frac{L}{k^{\frac{1}{d}}} \cdot t^{\frac{1}{d}-\frac{1}{g(h-1)}} \cdot n^{1-\frac{1}{d}+\frac{1}{g(h-1)}}\right) && \text{since } \textstyle\sum_{i=1}^{t} |S_i| = n \\
&= O\left(L \cdot k + L \cdot k^{-\frac{1}{g(h-1)}} \cdot n^{1-\frac{1}{d}+\frac{1}{g(h-1)}}\right) && \text{since } t \leq k \\
&= O\left(L \cdot n^{1-\frac{1}{d}+\frac{1}{g(h)}} + L \cdot n^{-\frac{1}{g(h-1)}(1-\frac{1}{d}+\frac{1}{g(h)})+1-\frac{1}{d}+\frac{1}{g(h-1)}}\right) \\
&= O\left(L \cdot n^{1-\frac{1}{d}+\frac{1}{g(h)}} + L \cdot n^{1-\frac{1}{d}+\frac{g(h)-d}{d \cdot g(h-1) \cdot g(h)}}\right) \\
&= O\left(L \cdot n^{1-\frac{1}{d}+\frac{1}{g(h)}}\right).
\end{aligned}$$

where the last step follows since $\frac{g(h)-d}{d \cdot g(h-1) \cdot g(h)} = \frac{d \cdot g(h-1)}{d \cdot g(h-1) \cdot g(h)} = \frac{1}{g(h)}$. □

Finally, it is not hard to verify that, for any $h > 0$, the worst-case time complexity is $O(n)$.

5 Open problems

The main open question is to refine the asymptotical analysis in order to obtain bounds on the constant factors. It would be interesting to understand how the constant factors depend on h: As suggested by the experimental results [9], they can even depend exponentially on h. Finally, it would be interesting to extend our asymptotical analysis to non constant h (e.g. $h = \Omega(\log n)$).

References

[1] L. Alfandari and V.T. Paschos. Approximating minimum spanning tree of depth 2. *Intl. Trans. In Op. Res.* 6:607-622, 1999.

[2] E. Althaus, S. Funke, S. Har-Peled, J. Koenemann, E. A. Ramos, M. Skutella, Approximation k-hop minimum-spanning trees, *Operations Research Letters*, 33, 115-120, 2005.

[3] S. Arora, P. Raghavan, and S. Rao. Approximation schemes for Euclidean k-medians and related problems. *Proc. 30-th ACM Symposium on Theory of Computing* 106-113, 1998.

[4] K. Bala, K. Petropoulos, and T.E. Stern. Multicasting in a Linear Lightwave Network. *Proc. of INFOCOM*, 1350–1358, 1993.

[5] A. Balakrishnan and K. Altinkemer. Using a hop-constrained model to generate alternative communication network design. *ORSA Journal of Computing* 4: 147-159, 1992.

[6] A. Bookstein and S.T. Klein. Compression of Correlated Bit-Vectors. *Information Systems*, 16(4):110–118, 1996.

[7] R. Chou and T. Johnson. *Distributed Operating Systems and Algorithms.* Addison-Wesley. Reading, MA 1997

[8] F.A. Chudak. Improved approximation algorithms for uncapacitated facility location problem. *Proc. of the 6-th Conference on Integer Programming and Combinatorial Optimization* 1998.

[9] A.E.F. Clementi, M. Di Ianni, A. Monti, G. Rossi, and R. Silvestri, Experimental Analysis of Practically Efficient Algorithms for Bounded-Hop Accumulation in Ad-Hoc Wireless Networks, In *Proc. of the IEEE IPDPS-WMAN*, 2005, to appear.

[10] A. Ephremides, G.D. Nguyen, and J.E. Wieselthier. On the Construction of Energy-Efficient Broadcast and Multicast Trees in Wireless Networks. In *Proc. of the 19th INFOCOM*, 585–594, 2000.

[11] L. Gouveia. Using the Miller-Tucker-Zemlin constraints to formulate a minimal spanning tree problem with hop constraints. *Computers and Operations Research* 22: 959-970, 1995.

[12] L. Gouveia. Multicommodity flow models for spanning trees with hop constraints. *European Journal of Operational Research* 95:178-190, 1996.

[13] L. Gouveia and C. Requejo. A new relaxation approach for the hop-constrain minimum spanning tree problem. *European Journal of Operational Research* 132:539-552, 2001.

[14] S. Guha and S. Khuller. Greedy strikes back: Improved facility location algorithms. *Journal of Algorithms* 31:228-248, 1999.

[15] M. Haenggi. Twelve Reasons not to Route over Many Short Hops, in *Proc. IEEE Vehicular Technology Conference (VTC'04 Fall)*, (Los Angeles, CA), Sept. 2004.

[16] G. Kortsarz and D. Peleg. Approximating the weight of shallow Steiner trees. *Discrete Applied Mathematics* 93:265-285, 1999.

[17] M.R. Korupolu, C.G. Plaxton and R. Rajaraman. Analysis of a Local Search Heuristic for Facility Location Problems. *Proc. of the 9th annual ACM-SIAM Symposium on Discrete algorithms* 1-10, 1998.

[18] M. Mahdian, Y. Ye and J. Zhang, A 1.52-approximation algorithm for the uncapacitated facility location problem. *Proc. of APPROX 2002* LNCS 2462: 229-242, 2002.

[19] C.J.H. McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics: Invited Papers at the 12th British Combinatorial Conference.* J. Siemons Ed. London Mathematical Society Lecture Notes Series, 141:148-188, 1989.

[20] P. Raghavan and R. Motwani, *Randomized Algorithms.* Cambridge Univ. Press, 1995.

[21] G. R. Raidl and B. A. Julstrom, Greedy Heuristics and an Evolutionary Algorithm for the Bounded-Diameter Minimum Spanning Tree Problem. *Proc. of the 2003 ACM Symposium on Applied Computing* 747-752, 2003.

[22] K. Raymond. A Tree-Based Algorithm for Distributed Mutual Exclusion. *ACM Transactions on Computer Systems*, 7(1):61–77, 1989.

[23] D.B. Shmoys, E. Tardos and K. Aardal, Approximation algorithms for facility location problems. *Proc. of the 29-th Annual ACM Symposium on Theory of Computing* 265-274, 1997.

[24] S. Voss, The steiner tree problem with hop constraint, *Annals of Operations Research* 86:321-345, 1999.