Approximating the Minimum Number of Maximum Power Users in Ad hoc Networks

Errol L. Lloyd[1], Rui Liu[1], and S. S. Ravi[2]

[1] University of Delaware, Newark, DE 19716.
Email: {elloyd, ruliu}@cis.udel.edu.
[2] University at Albany - SUNY, Albany, NY 12222.
Email: ravi@cs.albany.edu.

Abstract. Topology control is the problem of assigning transmission power values to the nodes of an ad hoc network so that the induced graph satisfies some specified property. The most fundamental such property is that the network/graph be connected. For connectivity, prior work on topology control gave a polynomial time algorithm for minimizing the maximum power assigned to any node (such that the induced graph is connected). In this paper we study the problem of minimizing the number of maximum power nodes. After establishing that this minimization problem is **NP**-Complete, we focus on approximation algorithms for graphs with symmetric power thresholds. We first show that the problem is reducible in an approximation preserving manner to the problem of assigning power values so that the sum of the powers is minimized. Using known results for that problem, this provides a family of approximation algorithms for the problem of minimizing the number of maximum power nodes with approximation ratios of $5/3 + \epsilon$ for every $\epsilon > 0$. Unfortunately, these algorithms, based on solving large linear programming problems are not practical. The main result of this paper is a practical algorithm having a 5/3 (exactly) approximation ratio. In addition, we present experimental results on randomly generated networks. Finally, based on the reduction to minimizing the total power problem, we outline some additional results for minimizing the number of maximum power users, both for graph properties other than connectivity and for graphs with asymmetric power thresholds.

1 Introduction

Considerable attention has been given to problems of **topology control** in ad hoc networks. Recall that an **ad hoc network** consists of a collection of transceivers for which all communication is based on radio propagation. For each ordered pair (u, v) of transceivers, there is a **transmission power threshold**, denoted by $p(u, v)$, where a signal transmitted by the transceiver u can be received by v only when the transmission power of u is at least $p(u, v)$. The transmission power threshold for a pair of transceivers depends on a number of factors including the distance between the transceivers, the direction of the antenna at the sender, interference, noise, etc. [RR00]. Further, due to those same factors, $p(u,v)$ and $p(v,u)$ need not be identical. When $p(u,v)$ and $p(v,u)$ are equal for all u and v, the power thresholds are **symmetric**. If they are unequal for some u and v, then the power thresholds are **asymmetric**. In this paper, unless otherwise specified, all of the problems considered utilize symmetric power thresholds.

Given the transmission powers of the transceivers, an ad hoc network can be represented by an undirected graph [KK+97]. In this graph, the nodes are in a one-to-one correspondence with the transceivers. There is an edge (u,v) if and only if the transmission powers of both u and v are at least the transmission power threshold $p(u,v)$. Note that every edge in the undirected graph corresponds to a two-way communication.

The main goal of **topology control** is to assign transmission powers to transceivers so that the resulting undirected graph satisfies a specified property, while minimizing some function of the transmission powers assigned to the transceivers. Limiting the maximum power used at any node, and using a minimum amount of power at each node to achieve a given task is likely to decrease the MAC layer interference between adjacent radios. The reader is referred to [LHB+01,RMM01,WL+01,RR00,RM99] for a thorough discussion of power control issues in ad hoc networks.

The most fundamental topology control problem [RR00] is to minimize the maximum power utilized by any node such that the resulting graph is connected. For this problem polynomial time algorithms are known [RR00,LLM02]. These algorithms are based on using binary search over all of the relevant power values (i.e. those power values corresponding to the transmission power thresholds). Since the goal is to minimize the maximum power assigned to any node, these algorithms may assign the computed minimum maximum power to *every* node of the network[3]. In practice however it is desirable not only to minimize the maximum power, but also to limit the number of nodes utilizing that maximum power.

In this paper we study the problem of assigning powers to nodes such that the induced graph is connected while minimizing the maximum power used by any node and minimizing the number of nodes that utilize that maximum power. In the next section we give formal definitions and some additional background. In Section 3 we show that the problem is **NP**-hard, and describe an approximation preserving reduction to the problem of minimizing the total power assigned to all of the nodes such that the resulting graph is connected. Using known results [ACMP04] for that problem, this provides a family of approximation algorithms for the problem of minimizing the number of maximum power nodes with approximation ratios of $5/3+\epsilon$ for every $\epsilon > 0$. Unfortunately, these algorithms, based on solving large linear programming problems are not practical. In Section 4 we give our main result, an algorithm with an approximation ratio of $5/3$ (exactly). Experimental results are given in Section 5. Finally, some further con-

[3] A method to minimalize the power at every node, once the maximum is found is discussed in [RR00], but was not implemented.

sequences of the reduction to the problem of minimizing the total power, along with open problems are discussed in Section 6.

2 Problem Formulation and Background

In this section we first give a formal definition of the problem studied in this paper, along with related terminology. The section also includes a brief review of some related prior work on topology control.

2.1 The Max-Power Users Problem

A formal statement of the decision version of the problem studied in this paper is as follows.

Max-power Users

Instance: A positive integer M, a positive number P (maximum allowable power value), a node set V, and a power threshold value $p(u, v)$ for each pair (u, v) of transceivers.

Question: Is there a power assignment where the power assigned to each node is at most P and the number of the nodes that are assigned power P is at most M, such that the resulting undirected graph G is connected?

Note that the above definition differs slightly from that described in the introduction, in that the formal statement assumes that the maximum allowable power value is given as an input. This is a reasonable assumption since for a given network, the problem of minimizing the maximum power such that the induced graph is connected can be solved in polynomial time [RR00,LLM02]. The problem considered here takes that maximum power value as an input and aims to minimize the number of nodes utilizing that power.

2.2 Some Prior Work

A notation was given in [LLM02] whereby a topology control problem is specified by a triple of the form $\langle \mathbb{M},\ \mathbb{P},\ \mathbb{O} \rangle$. In this notation, $\mathbb{M} \in \{$Undir, Dir$\}$ represents the graph model, $\mathbb{P}$ represents the desired graph property and $\mathbb{O}$ represents the minimization objective.

The form of topology control considered in this paper was proposed by Ramanathan and Rosales-Hain [RR00]. They presented efficient algorithms for two topology control problems, namely ⟨Undir, Connected, MaxP⟩ and ⟨Undir, 2-Node Connected, MaxP⟩. In addition, they presented efficient distributed heuristics for those problems.

Considerable work has been done over the past three years on a variety of topology control problems. For instance, several groups of researchers have studied the problems ⟨Undir, Connected, TotalP⟩, ⟨Undir, 2-Node Connected, TotalP⟩ and ⟨Undir, Diameter K, MaxP⟩ (see for example [LLM02,KL+03]). Likewise, work on ⟨Dir, Strongly Connected, TotalP⟩

may be found in [CH89,KK+97,CPS99,CPS00]. In most instances, the problems are shown to be **NP**-Hard and the focus is on the development of approximation algorithms having either $O(\log n)$ or constant approximation ratios.

The problem of minimizing the number of maximum power users was briefly addressed in [LLM02] where it was shown that minimizing the number of maximum power users is **NP**-Hard when the goal is to produce a connected graph with diameter at most 6.

3 Two Complexity Results

In this section we discuss two complexity results for the **Max-power Users** problem. The first result shows that the problem is **NP**-Complete. The second result gives an approximation preserving reduction to the ⟨Undir, Connected, TotalP⟩ problem.

3.1 Minimizing the Number of Max Power Nodes is NP-Complete

As noted above, it was shown in [LLM02] that minimizing the number of maximum power users is **NP**-Hard when the goal is produce a connected graph with diameter at most 6. In this section, we show that the problem is **NP**-Complete even when all that is desired is for the graph to be connected. The following result can be proven by a reduction from the minimum set cover problem [GJ79]. The proof is omitted due to space constraints.

Theorem 1. *The problem* **Max-power Users** *is* ***NP****-Complete.* □

3.2 A Reduction to Minimizing Total Power

In this section we give an approximation preserving reduction from the **Max-power Users** problem to ⟨Undir, Connected, TotalP⟩. Since several approximation results are known for ⟨Undir, Connected, TotalP⟩this will immediately provide identical results for **Max-power Users**.

Theorem 2. *There exists a polynomial-time reduction from* **Max-power Users** *to* ⟨Undir*,* Connected*,* TotalP⟩ *such that any α-approximation algorithm for* ⟨Undir*,* Connected*,* TotalP⟩ *is also an α-approximation algorithm for* **Max-power Users***.*

Proof: Consider an instance I of the **Max-power Users** problem. Let $p(x, y)$ denote the (symmetric) power threshold values for any pair of nodes x and y. Let P denote the smallest maximum power value for I (this can be computed efficiently using the algorithms of [RR00,LLM02]). Further, with that P, let P_0 = $\max\{p(x, y)$: $p(x, y) < P\}$. Since the goal is to minimize the number of nodes assigned the power value P, it is sufficient to consider solutions where all other nodes are assigned the power value P_0.

We now construct an instance I' of ⟨Undir, Connected, TotalP⟩ as follows. For the instance I', the power threshold value for each pair of nodes x and y, denoted by $p'(x, y)$, is chosen as follows:

$$
\begin{aligned}
p'(x,y) &= 1 \quad \text{if } P_0 < p(x,y) \leq P \\
&= 0 \quad \text{if } p(x,y) \leq P_0 \\
&= \infty \quad \text{otherwise}
\end{aligned}
$$

Now, for the instance I', the power value to be assigned to each node is either 0 or 1. If the power assigned to node x in I' is 1 (0), that corresponds to assigning the power value P (P_0) to the node x in I. Thus, the total assigned power value in I' is the number of nodes assigned the maximum power value in I. Further, it is clear that any α-approximation algorithm for ⟨UNDIR, CONNECTED, TOTALP⟩is also an α-approximation algorithm for **Max-power Users**. □

Combining the above theorem with the results of [ACMP04], the following are immediate:

- For any $\epsilon > 0$, there is a $(5/3+\epsilon)$-approximation algorithm for **Max-power Users**. Unfortunately, it is noted in [ACMP04] that "this algorithm is impractical", due to its reliance on solving large linear programming instances and other related issues.
- The minimum spanning tree based algorithm of [RR00] for ⟨UNDIR, CONNECTED, TOTALP⟩is a 2-approximation algorithm for **Max-power Users**. In contrast to the prior observation, the minimum spanning tree based algorithm of [RR00] has often been the method of choice for solving both ⟨UNDIR, CONNECTED, MAXP⟩and ⟨UNDIR, CONNECTED, TOTALP⟩. In the experimental studies described in this paper we include this algorithm as a baseline method for solving **Max-power Users**.

The NP-hardness results presented in this section apply to general instances of the **Max-power Users** problem rather than geometric instances (where the nodes are points in Euclidean space and the power threshold value for each pair of nodes is a function of the distance between the nodes). However, the approximation results presented in the subsequent sections of this paper are applicable to geometric instances as well.

4 A 5/3-approximation Algorithm

In this section we describe a 5/3-approximation algorithm for the **Max-power Users** problem, and show that the approximation ratio is tight. We begin by defining the following concept:

Definition 1. *Consider an undirected graph $G(V,E)$ and two subsets E' and H of E. For a node $u \in V$, a* **Maximal Connected Component Tree (MCCT)** *with root u for graph $G'(V,E')$ and edge subset H, is a subset T_u of H such that all of the following conditions hold:*

(a) The restriction of G to T_u is a tree. By abuse of notation, we let T_u denote that restriction of G.
(b) u is a node in T_u.

(c) Each node in T_u is in a different connected component of $G'(V, E')$.
(d) T_u is maximal with respect to properties (a) and (c) above (i.e. adding any edge $(x, y) \in H$ to T_u will either destroy property (a) or property (c)).

The number of edges in T_u is called the **size** *of T_u.*

Figure 1 illustrates the definition of MCCT.

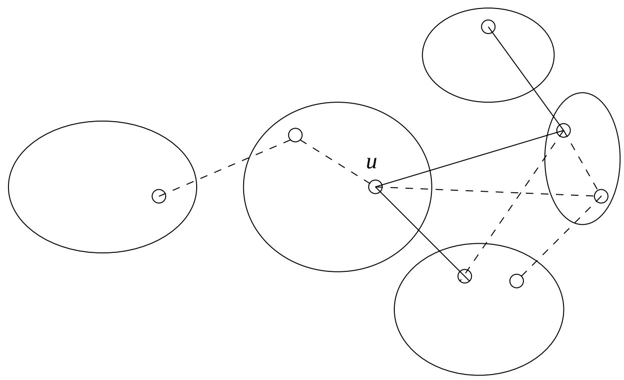

Fig. 1. MCCT of u: The ellipses illustrate the different connected components in G', H consists of all of the edges shown (solid and dashed), and T_u, a Maximal Connected Component Tree (MCCT) with root u for graph $G'(V, E')$ and edge subset H, is the set of all solid edges. The size of T_u is 3.

An MCCT based approximation algorithm for **Max-power Users** is given in Algorithm 1. For an understanding of the ideas behind this algorithm, we first consider the following method for finding a solution to **Max-power Users**: Using the power threshold graph, assign every node a power level equal to the greatest adjacent threshold that is less than p_{max} (the minimum maximum power). The graph induced by this power assignment will consist of several connected components. Then, one way to make this graph connected is to add maximum power edges (i.e. those having a power threshold of p_{max}) to this graph until the graph is connected. This is the essential idea underlying the minimum spanning tree based algorithm of [RR00] that was mentioned in the prior section. Clearly, in this approach if there are k connected components prior to the addition of any maximum power edges, then $k - 1$ maximum power edges must be added. The key issue is how many distinct nodes are adjacent to these edges. Clearly, k is a lower bound on that number of nodes. And, in [RR00] the number of distinct nodes might be as large as $2 \cdot (k - 1)$, hence the approximation ratio of 2.

How can we improve upon the ratio of 2? One approach would be to select maximum power edges that share adjacent nodes. Alternatively, we could tighten the lower bound of k on the number of maximum power nodes. Obviously, one or both approaches might apply for a particular instance of the **Max-power Users** problem. Algorithm 1 attempts to exploit the first approach, by focusing on MCCTs of particular sizes as it selects maximum power edges that result in a connected graph. In particular, the algorithm first finds MCCTs of size at least 3, then finds MCCTs of size 2, and finally includes single edges. Further, our proof on the approximation ratio associated with Algorithm 1 will tighten the lower bound as well.

Algorithm 1 A 5/3-Approximation Algorithm for **Max-power Users**

Input: A complete power threshold graph $G(V, E)$ and the minimum maximum power p_{max}.
Output: Power assignment A for each node u in V.

1: $E' \leftarrow \{(u, v) \in E | \; p(u, v) < p_{max}\}$
2: $H \leftarrow \{(u, v) \in E | \; p(u, v) = p_{max}\}$
3: **while** $\exists u \in V$, for which a MCCT T_u for $G'(V, E')$ and H has size greater than 2 **do**
4: $\quad E' \leftarrow E' \cup T_u$
5: **end while**
6: **while** $\exists u \in V$, for which a MCCT T_u for $G'(V, E')$ and H has size 2 **do**
7: $\quad E' \leftarrow E' \cup T_u$
8: **end while**
9: Compute the connected components $C_1, ..., C_L$ of $G'(V, E')$.
10: Add $L - 1$ edges in H to E' such that $G'(V, E')$ is connected.
11: **for** each u in V **do**
12: $\quad A(u) \leftarrow$ the weight of the largest edge in E' that is incident on u.
13: **end for**
14: Return A

It is easy to see from the algorithm's specification that it produces a power assignment for the nodes in V such that the induced graph is connected. In regard to the quality of that power assignment, we have:

Theorem 3. *Algorithm 1 is a 5/3-approximation algorithm.*

Proof: Suppose before the execution of the first *while* loop, that $G'(V, E')$ has K connected components. In iteration i of the *while* loop (lines 3–5), let s_i be the size of T_u. Note that $s_i \geq 3$, and that by adding T_u to E' the number of connected components in $G'(V, E')$ is reduced by s_i because T_u connects $s_i + 1$ different connected components.

We claim that T_u is node disjoint from any edges in H that were added to E' in the previous iterations. The reason is as follows. Suppose there exist edges (t, v) in T_u and (v', t) in T', a MCCT chosen in some previous iteration (i.e. these two edges have node t in common). Since (t, v) is in T_u, it follows that in iteration i, nodes t and v must be in different connected components of $G'(V, E')$. Also in iteration i all nodes in T' are in the same connected component. This component then includes node t, but not node v. But this means that in the prior iteration, $T' \cup (t, v)$ is a larger $MCCT$ than T'. This contradicts T' being a $MCCT$ (it was not maximal). Thus, the claim is proved.

As a result of the above, in both *while* loops, adding T_u to E' causes $s+1$ additional nodes to be assigned power p_{max} in A, where s is the size of T_u.

Now, suppose before the execution of the first *while* loop, that $G'(V, E')$ has K connected components, that before the execution of the second *while* loop, $G'(V, E')$ has M connected components, and that after the execution of the second *while* loop, $G'(V, E')$ has P connected components. Let s_i be the size of T_u in iteration i of the first *while* loop and let m_i be the size of T_u in iteration i of the second *while* loop. Let $N(A)$ be the number of nodes that are assigned power p_{max} in A. Then, since each $s_i \geq 3$ and each $m_i = 2$, we have:

$$
\begin{aligned}
N(A) &= \sum_i (s_i + 1) + \sum_i (m_i + 1) + 2(P-1) \\
&= \sum_i \left(\frac{s_i+1}{s_i} \cdot s_i\right) + \sum_i \left(\frac{m_i+1}{m_i} \cdot m_i\right) + 2(P-1) \\
&\leq \frac{4}{3}\sum_i s_i + \frac{3}{2}\sum_i m_i + 2(P-1) \\
&= \frac{4(K-M)}{3} + \frac{3(M-P)}{2} + 2(P-1)
\end{aligned}
$$

Consider an optimum solution for **Max-power Users**. It is obvious that:

Claim. 1 In any optimum solution, at least K nodes are assigned power p_{max}.

Let $G_{opt}(V, E_{opt})$ be the graph induced by an optimum solution. Consider the graph $G'(V, E')$ after the first *while* loop of Algorithm 1. Recalling that there are M connected components in G', we have:

Claim. 2 There exist node disjoint MCCTs for graph $G'(V, E')$ and edge set E_{opt} such that:

- Each MCCT[4] is of size 1 or 2.
- The sum of the sizes of these MCCTs is at least $M - 1$.

Proof of Claim 2: Let D_0 be the nodes in some connected component of $G'(V, E')$. Since $G_{opt}(V, E_{opt})$ is connected, for each cut of $G_{opt}(V, E_{opt})$ there must exist an edge in E_{opt} that crosses the cut. Let e_1 be an edge in E_{opt} that crosses the cut $(D_0, V - D_0)$. Note that e_1 must be a maximum power edge. Further, there must be a MCCT T_1 for $G'(V, E')$ and edge set E_{opt} that includes e_1. The size of that MCCT may be either 1 or 2 (it cannot be larger, since then it would have been discovered in the first *while* loop of Algorithm 1). Consider $G_1(V, E_1)$ where $E_1 = E' \cup T_1$. Now, let D_1 be the nodes in some connected

[4] Trivially, an MCCT of size 1 is a singleton edge. In the remainder of this section we will sometimes find it convenient to refer to edges that are not part of a larger MCCT in this fashion (i.e. as MCCTs of size 1).

component of G_1. As above, let e_2 be an edge (necessarily of maximum power) that crosses the cut $(D_1, V - D_1)$, and let T_2 be a MCCT for G_1 and edge set E_{opt} that includes e_2. Again, the size of that MCCT must be 1 or 2. Continuing in this fashion, we enumerate MCCTs $T_1, T_2, ..., T_g$, ending when the corresponding graph G_g is connected. Note that each of these MCCTs is of size 1 or 2, and that it follows from the definition of a MCCT that these MCCTs are node disjoint. It also follows that each of these MCCTs is present in graph G'. Finally, since G' had M connected components, the sum of the sizes of these MCCTs is $M-1$. □

It follows from the claim that if L of the MCCTs referenced there are of size 1, then there are at least $3/2(M-L-1)+2L = 3M/2+L/2-3/2$ nodes using power p_{max} in an optimum solution.

To complete the proof we need to relate the number of MCCTs of size 2 selected in Algorithm 1 to the values L and M. Specifically we have:

Claim. 3 In the second *while* loop of Algorithm 1 (lines 6–8), at least $(M-L)/4$ MCCTs are selected. Note that those MCCTs are all of size 2.

The proof of Claim 3is omitted.

Let $N(OPT)$ be the number of nodes that are assigned power p_{max} in an optimum solution. From Claims 1 and 2, we have that $N(OPT) \geq \text{Max}(K, 3M/2 + L/2 - 3/2)$. And, using Claim 3, $(M-P)/2 \geq (M-L)/4$, hence $P \leq M/2 + L/2$. We then have:

$$\frac{N(A)}{N(OPT)} \leq \frac{4/3(K-M)+3/2(M-P)+2(P-1)}{\text{Max}(K, 3M/2+L/2-3/2)}.$$

Using straightforward algebraic manipulations, it can be shown that the above ratio is bounded by 5/3. Thus, the approximation ratio of Algorithm 1 is 5/3. □

Due to space constraints, we state the following corollary without proof.

Corollary 1. *The* 5/3 *approximation ratio of Algorithm 1 is tight.*

An interesting question is what happens in Algorithm 1 if we do not give priority to MCCTs of size 3 or more, but rather consider equally acceptable any MCCT of size 2 or more. That is, in Algorithm 1, in line 3, replace "H has size greater than 2" with "H has size greater than 1", and we omit lines 6 to 8. We will call this *Algorithm 2*. Then we have the following corollary:

Corollary 2. *Algorithm 2 is a* 7/4*-approximation algorithm for* **Max Power Users***, and the bound is tight.*

Finally, we state without proof the running times of our algorithms:

Proposition 1. *The running times of Algorithms 1 and 2 are* $O(n^3)$ *and* $O(n^2)$ *respectively.*

5 Experimental Results

In this section we consider the experimental performance of Algorithms 1 and 2 along with the algorithm of [RR00].

The experimental environment used here is derived from the one described in [RR00]. The radio wave propagation model is the *Log-distance Path Loss Model*:

$$PL(d) \;=\; -10\log_{10}\left[\frac{G_t G_r \lambda^2}{(4\pi)^2 d_0^2}\right] + 10\eta\log_{10}\left[\frac{d}{d_0}\right]$$

where η is the path loss exponent, d_0 is the close-in reference distance, λ is the radio wavelength, G_t is the transmitter antenna gain, G_r is the receiver antenna gain, and d is the separation distance between transmitter and receiver (see [Ra96] for detailed descriptions of these parameters). All of the parameters are chosen to emulate a 2.4 GHz wireless radio, and if d is less than a certain threshold, the transmission power threshold is set to the minimum threshold of 1 dBm.

Experiments were conducted on networks with 200 nodes by varying the geographical distribution of the nodes and the power levels available to the nodes. Nodes were placed using a uniform random distribution, and with 200 nodes in each network, the node density was 12.5 nodes/sq mile in a 4 mile by 4 mile area. Three sets of power levels were used, with a varying ratio between the top two power levels. Those three sets of power levels were $(24, 8, 2, 1)$, $(24, 12, 6, 3)$ and $(18, 12, 9)$, corresponding to ratios between the top two power levels of 3, 2 and 1.5. Each data point represents the average taken over 19 trials.

In each experiment, after generating a placement of the nodes, three algorithms were run on the network consisting of those nodes. The three algorithms were our Algorithms 1 and 2, along with the minimum spanning tree algorithm from [RR00] that was briefly discussed in Section 3 (for brevity, we will refer to this as MST [RR00]). Each algorithm assigns powers to nodes such that the resulting network is connected. For each algorithm we measure the number of nodes that are assigned maximum power. In addition, we record the average power assigned, and the maximum and average degrees of the nodes in the resulting network.

The experimental results on nodes of maximum power are shown in Figure 2. There, in addition to the number of maximum power nodes produced by each algorithm, we provide a lower bound on the optimal number of maximum power nodes. This lower bound is calculated as in the proofs of the performance bounds in the prior section, and is based on the number of connected components, and the numbers of MCCTs of size 2 and size 3. Relative to the results in Figure 2 we make the following observations:

– The two MCCT based algorithms (Algorithms 1 and 2) outperform the MST based algorithm (MST [RR00]) in regard to the number of maximum power nodes. The reductions obtained by using the MCCT based algorithms range from 7% to 23%, with larger reductions occurring when there is a larger spread in the top two power levels.

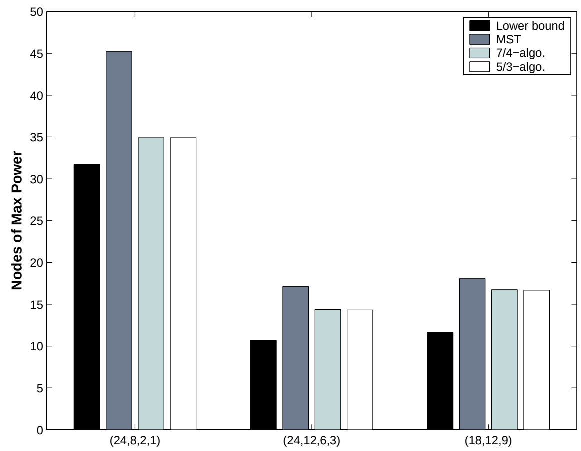

Fig. 2. Max Power Nodes

– There is virtually no difference in the performance of Algorithm 1 (the 5/3 approximation ratio) as compared with the performance of Algorithm 2 (the 7/4 approximation ratio). While they often obtain their MCCTs in different orders (in particular, the MCCTs of size 2 come last in Algorithm 1), it seems to be rare for the selection of one MCCT to eliminate another MCCT. Further, in most cases, the number of MCCTs found by either algorithm is fairly small. Usually there were fewer than 5 MCCTs, and in a surprisingly large number of cases there was only one, very large, MCCT. That single MCCT involved a node from the majority of the connected components that existed prior to MCCT selection in the *while* loops of Algorithms 1 and 2.
– On the average, Algorithms 1 and 2 use about 22% more maximum power nodes than the lower bound.

The results on average power are omitted due to space considerations, however we note that the improvements in average power due to Algorithms 1 and 2 are modest, ranging from under 1% to just over 10%.

Figure 3 shows the results on average degree. Here, some reductions in average degree occur when using the MCCT based algorithms. The results are most significant when the ratio between the top two power levels is high. Specifically, when the power levels are $(24, 8, 2, 1)$, the MCCT based algorithms reduce the average node degree by 17%.

6 Additional Observations and Future Research

Theorem 2 showed that the **Max-power Users** problem can be reduced in an approximation preserving manner to to $\langle$Undir, Connected, TotalP$\rangle$. Actually, that result is not restricted to power assignments producing connected

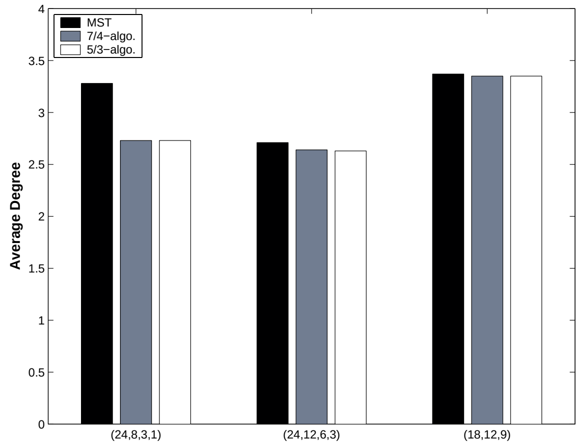

Fig. 3. Node Degree

graphs, but can be generalized to apply for any monotone graph property as well as for asymmetric power thresholds. As a consequence, approximation algorithms for several (generalized) **Max-power Users** problems can be obtained. These include: for symmetric power thresholds, constant factor approximations for the properties 2-connected and 2-edge-connected, as well as an $O(\log n)$ approximation for connectivity with asymmetric power thresholds. In addition, we can show using a reduction from minimum set cover that there is a matching $\Omega(\log n)$ lower bound for the connectivity problem under asymmetric thresholds.

The primary open question is whether there are algorithms with approximation ratios lower than 5/3 for the **Max-power Users** problem for connectivity. We conjecture that such algorithms exist. Specifically, if Algorithm 1 is extended in the obvious way by finding larger MCCTs before smaller ones (i.e. much as Algorithm 1 is an extension of Algorithm 2), it is likely that a smaller approximation ratio can be found. The difficulties in making this extension are: 1) Determining the relationship between the MCCTs of a given size found by the algorithm and the existence of MCCTs in optimal solutions; 2) Evaluating the resulting equations; and 3) The running time would seem to be dependent on the size of the MCCTs that are sought, hence it is likely that some kind of approximation scheme would result (though we suspect that even so, the achievable bound will not fall below 1.5).

References

[ACMP04] E. Althaus and G. Calinescu and I. Mandoiu and S. Prasad and N. Tchervenski and A. Zelikovksy. "Power Efficient Range Assignment for Symmetric Connectivity in Static Ad Hoc Wireless Networks", Private communication

[CH89] W. Chen and N. Huang. "The Strongly Connecting Problem on Multihop Packet Radio Networks", *IEEE Trans. Communication*, Vol. 37, No. 3, Mar. 1989, pp. 293-295.

[CPS99] A. E. F. Clementi, P. Penna and R. Silvestri. "Hardness Results for the Power Range Assignment Problem in Packet Radio Networks", *Proc. Third International Workshop on Randomization and Approximation in Computer Science* (APPROX 1999), Lecture Notes in Computer Science Vol. 1671, Springer-Verlag, July 1999, pp. 195–208.

[CPS00] A. E. F. Clementi, P. Penna and R. Silvestri. "The Power Range Assignment Problem in Packet Radio Networks in the Plane", *Proc. 17th Annual Symposium on Theoretical Aspects of Computer Science* (STACS 2000), Feb. 2000, pp. 651–660.

[GJ79] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*, W. H. Freeman and Co., San Francisco, CA, 1979.

[GK99] S. Guha and S. Khuller, "Improved Methods for Approximating Node Weighted Steiner Trees and Connected Dominating Sets", *Information and Computation*, Vol. 150, 1999, pp. 57–74.

[KK+97] L. M. Kirousis, E. Kranakis, D. Krizanc and A. Pelc. "Power Consumption in Packet Radio Networks", *Proc. 14th Annual Symposium on Theoretical Aspects of Computer Science* (STACS 97), Lecture Notes in Computer Science Vol. 1200, Springer-Verlag, Feb. 1997, pp. 363–374.

[KL+03] S. O. Krumke, R. Liu, E. L. Lloyd, M. V. Marathe, R. Ramanathan and S. S. Ravi, "Topology Control Problems Under Symmetric and Asymmetric Power Thresholds", *Proc. International Conference on Ad hoc and Wireless Networks* (ADHOC-NOW'03), Lecture Notes in CS, Vol. 2865, (Edited by S. Pierre, M. Barbeau and E. Kranakis), Montreal, Canada, Oct. 2003, pp. 187–198.

[LLM02] E. Lloyd, R. Liu, M. Marathe, R. Ramanathan, and S. S. Ravi, "Algorithmic Aspects of Topology Control Problems for Ad hoc Networks", *Proc. Third ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)* (2002), 123–134. (Complete version to appear in *Mobile Networks and Applications*.)

[LHB+01] L. Li, J. Y. Halpern, P. Bahl, Y. Wang and R. Wattenhofer, "Analysis of Cone-Based Distributed Topology Control Algorithm for Wireless Multihop Networks", *Proc. ACM Principles of Distributed Computing Conference* (PODC'01), Aug. 2001, pp. 264–273.

[Ra96] T. S. Rappaport. *Wireless Communications: Principles and Practice*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1996.

[RM99] V. Radoplu and T. H. Meng. "Minimum Energy Mobile Wireless Networks", *IEEE J. Selected Areas in Communications*, 17(8), Aug. 1999, pp. 1333-1344.

[RMM01] E. M. Royer, P. Melliar-Smith and L. Moser, "An Analysis of the Optimum Node Density for Ad hoc Mobile Networks", *Proc. IEEE Intl. Conf. on Communication* (ICC'01), Helsinki, Finland, June 2001, pp. 857–861.

[RR00] R. Ramanathan and R. Rosales-Hain. "Topology Control of Multihop Wireless Networks Using Transmit Power Adjustment", *Proc. IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000, pp. 404–413.

[WL+01] R. Wattenhofer, L. Li, P. Bahl and Y. Wang. "Distributed Topology Control for Power Efficient Operation in Multihop Wireless Ad Hoc Networks", *Proc. IEEE INFOCOM 2001*, Anchorage, Alaska, April 2001, pp. 1388–1397.