

Topology Control in Unreliable Ad hoc Networks: Obtaining Near-Optimal Power Efficiency and Low Interference

Maleq Khan V.S. Anil Kumar Madhav Marathe Gopal Pandurangan S.S. Ravi

Abstract

Topology control is a basic subroutine in many wireless networking problems and is, in general, a multi-criteria optimization problem involving (contradictory) objectives of connectivity, interference, and power minimization. Wireless devices are often unreliable, and motivated by mission critical networks (which require ensuring reliable system performance), we study topology control problems in ad hoc networks under node failures for *arbitrary* node distributions. We consider a simple and natural stochastic failure model, in which each node can fail independently with a given probability. The *Topology Control Problem Under Stochastic Failures* is to choose a power level for each node and a subset of edges such that the residual graph (i.e., the graph formed by the nodes which have not failed) is connected and can be scheduled efficiently, with high probability. We develop provably efficient bi-criteria approximation algorithms for this problem that simultaneously minimize power, reduce interference, and ensure that the surviving graph is connected with high probability. Our algorithms can be implemented efficiently in a distributed manner.

Keywords: Topology Control Distributed Algorithms; Approximation Algorithms; Minimum Spanning Tree; Energy-Efficient

I. INTRODUCTION

A. Overview

Topology control is a fundamental problem in multi-hop ad hoc wireless networks where it is undesirable for nodes to transmit at the maximum power levels — in addition to wasting energy, which is an important constraint, it leads to high interference. However, reducing power levels

M. Khan, V.S. Anil Kumar and Madhav Marathe are with the Network Dynamics and Simulation Science Laboratory, Virginia Bioinformatics Institute, Virginia Tech, Blacksburg, VA 24061, USA. e-mail: {maleq,akumar}@vbi.vt.edu.

G. Pandurangan is with the Department of Computer Science, Purdue University, West Lafayette, IN, 47907. Email: gopal@cs.purdue.edu.

S.S. Ravi is with the Dept. of Computer Science, University at Albany - State University of New York. Email: ravi@cs.albany.edu.

A shorter version of this paper has been accepted in the Proceedings of the 28th IEEE Conference on Computer Communications (INFOCOM), 2009.

can impact connectivity, and so topology control involves a trade-off between the (contradictory) objectives of connectivity, low interference, and energy minimization [28, 4]. The focus of this paper is on topology control in *unreliable* sensor networks, in which nodes can fail – this is especially important in mission critical applications involving wireless networks, such as disaster recovery, wildfire monitoring [20] and surveillance, in which the underlying connectivity needs to be ensured in spite of such failures.

There has been a lot of work on topologies which remain connected even when some nodes fail [11, 18, 5, 10]. Two classes of failure models have been studied: *adversarial* and *stochastic*. In the adversarial failures model, the topology is required to be well-connected no matter which k nodes fail; in other words, the topology should be a $(k + 1)$ -connected graph. The cost of such solutions (relative to the optimum) typically depends on k , and can become weak when k is large. The adversarial model has been studied in both random and arbitrary node distributions. In the stochastic failures model, nodes are assumed to fail independently with some probability. Under this model, the connectivity properties have been studied as a function of the transmission range in the case of random node distributions [29].

In this paper, we study topology control problems in the stochastic node failures model, for arbitrary node distributions. Formally, given a set of n nodes (transceivers) placed arbitrarily in the plane, wherein each node v fails independently with probability $1 - p(v) \in [0, 1]$, the *Topology Control Problem under Stochastic Failures* (TCPSF) is to choose a topology (i.e., a subset of edges) and power levels for all nodes, so that the random subset of surviving nodes form a “well connected” graph with low interference (formally defined later), with probability at least $1 - \epsilon$, where ϵ is a parameter. We consider a specific time horizon (which is of length polynomial in n), and each node can fail and can come back again at any time; the power levels for nodes are determined at the beginning, and when a nodes recovers, it uses the power level specified initially. In contrast to random node distributions, in which many connectivity properties can be mathematically analyzed very accurately (e.g., see [29]), stochastic failures in arbitrary node distributions leads to new challenges. For instance, we prove that even determining the probability that there is a component containing a μ -fraction of the surviving nodes for a given constant μ (or approximating it within a small constant factor) is #P-complete. To the best of our knowledge,

this is the first work that addresses the topology control problem under stochastic node failures for arbitrary node distributions.

Similar issues arise in the design of sleep scheduling protocols, which try to minimize energy consumption by cycling nodes between sleep and active modes, and scheduling transmissions to achieve convergecast [6, 23]. While many references present protocols that involve a careful choice of schedules for nodes to sleep, transmit, and receive [23], others such as Chiasserini et al. [6] study the performance of protocols in which nodes randomly go to sleep for some (randomly chosen) duration and then wake up. Solving the TCPSF problem will also yield topologies that would work well under such sleep scheduling models.

Early work on the theoretical foundations of topology control, e.g., [27, 28, 5, 21] involved optimization of transmission power levels so that the resulting topology is well connected. However, it has been observed that focusing on the transmission power alone is not very insightful [24, 4], and in recent years there has been interest in designing *low-interference* topologies, in addition to optimizing the transmission powers [4, 26, 25]. We focus on both these metrics, and develop algorithms to construct a topology formed by a set E' of edges that can be scheduled efficiently, while minimizing different functions of the transmission power levels needed to realize the set E' of links.

B. Preliminaries and Network Model

Let V be the set of n nodes on the plane, and let $\ell(u, v)$ denote the distance between u and v , i.e., the length of edge (u, v) , for all $u, v \in V$. We will assume throughout that $\min\{\ell(u, v) : u, v \in V\} = 1$ and $\max\{\ell(u, v) : u, v \in V\} = O(n^{O(1)})$. For node $v \in V$, let $B(v, r) = \{w \in V : \ell(v, w) \leq r\}$ be the set of nodes within distance r from v . Following [27, 21], we assume that transmission on a link (u, v) requires a power level $P(u) \geq \phi(u, v) = c \cdot \ell(u, v)^\gamma$, where c and γ are constants; we assume adaptive power control, meaning that to transmit on edge $e = (u, v)$, node u uses the minimum power $P(u) = \phi(u, v)$ required for transmission on this link. Let $G = (V, E)$ denote the graph formed by a set E of links. For a given set E of links, let $P(u, E) = \max\{\phi(u, v) : (u, v) \in E\}$ denote power level needed for node u to realize the edges incident on u , and $\vec{P}(E)$ denote the resulting vector of power levels. Designing a topology corresponds to constructing such

a set E of edges. For a given transmission power vector \vec{P} , let $G = (V, \vec{P})$ denote the graph formed by the set $E = E(G) = \{(u, v) \in V \times V : P(u) \geq \phi(u, v)\}$ of all possible edges realized by \vec{P} . If $P(u) = \phi$ for all nodes u , we also refer to the graph $G(V, \vec{P})$ by $G(V, \phi)$. Let $G(V)$ denote the complete weighted graph on V with each edge (u, v) having weight $\ell(u, v)$. We let $MST(V)$ denote a minimum spanning tree of $G(V)$.

We assume that each node v can fail with probability $1 - p(v)$; for simplicity, we assume a uniform probability $p(v) = p$, though some of the results in this paper extend to the more general case of non-uniform probabilities. Let $V(p)$ denote a random subset of V , containing each $v \in V$ with probability p . Let $G(V(p), \vec{P})$ denote a random subgraph of $G(V, \vec{P})$ induced by the random subset $V(p)$; similarly, let $G(V(p), E)$ denote a random subgraph of $G = (V, E)$.

The TCPSF problem involves constructing a set E of edges so that the subgraph $G(V(p), E)$ induced by the surviving nodes has some desired property \mathbb{P} , which could be, for instance, the property that $G(V(p), E)$ is strongly connected, has a giant component or has low diameter. In most of this paper, \mathbb{P} will be the property of strong connectivity. We are interested in two objectives related to E : the *interference complexity* (which quantifies how efficiently the edges in E can be scheduled), and the *energy cost*.

We use the *Distance-2* model for interference, which is an abstraction of the Tx-Rx model (see, e.g., [17, 2]); in this model, edges e and e' can be used simultaneously, provided they do not have a common end point, and there is no edge connecting their end points. We seek to construct edge sets E with low “scheduling complexity” or interference cost, i.e., they can be scheduled in the Distance-2 (or other similar models) efficiently. There have been several proposals for combinatorial measures for the interference cost, e.g., the “coverage” measure of [4] and the slightly different congestion measure of [17, 16]. Our interference cost is based on [17, 16], since it seems to be a better measure of the scheduling complexity if adaptive power levels are allowed (i.e., nodes can vary the transmission power level on different links). Our algorithms produce topologies with low interference cost. Given a set E of edges, for edge $e \in E$, define the interference set of $e = (u, v)$, denoted by $\mathbb{I}(e, E)$, as

$$\mathbb{I}(e, E) = \{e' = (u', v') : e' \in E, \ell(e) \leq \ell(e'), \text{ and } \min\{\ell(u, u'), \ell(u, v'), \ell(v, u'), \ell(v, v')\} \leq \ell(e')\},$$

and $\mathbb{I}(E) = \max_{e \in E} |\mathbb{I}(e, E)|$. When the set E is clear from the context, we will refer to the set

$\mathbb{I}(e, E)$ simply by $\mathbb{I}(e)$. Intuitively, $\mathbb{I}(e)$ consists of edges that are located “close” to e , and interfere with it, and are longer in length; this measure is motivated by the following result from [17, 16]:

Lemma 1.1: Given a set E of edges, it is possible to construct an interference-free schedule of length $O(\mathbb{I}(E))$ in the Distance-2 model.

It has also been shown in [17, 16] that the above lemma holds for many different interference models, and this interference metric is stronger than the one in [4] if adaptive power levels are used; the “ $\ell(e') \geq \ell(e)$ ” restriction in the definition is crucial in this case, as illustrated in Figure 1. The second measure we consider is the energy cost, $\text{cost}(\vec{P}(E))$, and we are interested in two objectives: the maximum power, $\max(\vec{P}) = \max_{u \in V} P(u)$ and the total power, $\text{sum}(\vec{P}) = \sum_{u \in V} P(u)$.




Fig. 1. An example illustrating the interference measure: Consider the set $E = \{e, e_1, e_2, e_3, e_4, e_5\}$ and let $\vec{P}(E)$ denote the power level vector. Then, $\mathbb{I}(e) = \{e_5\}$, $\mathbb{I}(e_4) = \{e, e_5\}$, and $\mathbb{I}(E) = 3$, so that all the edges shown can be scheduled in three time slots: e_1, \dots, e_4 can be scheduled simultaneously, while e and e_5 can be scheduled in two separate slots. In contrast, the coverage defined in [4] is higher, and leads to a longer schedule.

For a given parameter ϵ , let $E = E_{opt}$ denote a set of edges E so that $G(V(p), E)$ is strongly connected with high probability¹, and $\text{cost}(\vec{P}(E_{opt}))$ is minimized; among multiple such E_{opt} which may be possible, we consider one that minimizes $\mathbb{I}(E)$. The TCPSF problem is a bi-criteria optimization problem that involves constructing a set E of edges so that $G(V(p), E)$ is strongly connected with high probability such that $\mathbb{I}(E)$ and $\text{cost}(\vec{P}(E))$ are both small, relative to $\mathbb{I}(E_{opt})$ and $\text{cost}(\vec{P}(E_{opt}))$, respectively. The algorithms we study will be randomized bi-criteria approximations: we say our algorithm produces an (α_1, α_2) -approximation to the TCPSF problem if it produces a solution E such that $\mathbb{I}(E) \leq \alpha_1 \mathbb{I}(E_{opt})$ and $\text{cost}(\vec{P}(E)) \leq \alpha_2 \text{cost}(\vec{P}(E_{opt}))$, with high probability.

¹ Usually we take this to mean a probability of at least $1 - \frac{1}{n^3}$, but this can be any parameter

C. Summary of Results

1. Computational hardness of TCPSF problem. Various versions of this problem are computationally hard to solve exactly. We prove in Section III that even computing the probability that $G(V(p), \phi)$ has a component of size at least $\mu|V|$ is #P-complete, for a given constant μ - this is in contrast with results on stochastic failures in random node distributions [19, 29], where this quantity can be estimated analytically.

Therefore, we focus on approximation algorithms with provable approximation guarantees of the form described in Section I(B). In all our results, the interference complexity will be poly-logarithmic, and so we classify the results based on the energy objective. To avoid trivialities, we assume that $p > 1/n^{3/2}$ i.e., the failure probability is not very close to 1.

2. The Sum of Powers Objective. We present an algorithm that gives an $(O(\log^2 n), O(\log^{\gamma+2} n))$ -approximation² for the sum objective for energy; in other words, our algorithm chooses a set E of edges such that $\mathbb{I}(E) = O(\log^2 n)\mathbb{I}(E_{opt})$, and $sum(\vec{P}(E))/sum(\vec{P}(E_{opt})) = O(\log^{\gamma+2} n)$. Our algorithm is based on a local labeling algorithm, in which each node chooses a set E of edges to some number of higher labeled nodes. We prove that for *any* choice of distinct label numbers, the resulting set E of edges guarantees that $\Pr[G(V(p), E) \text{ is strongly connected}] \geq 1 - 1/n^3$, and $\mathbb{I}(E) = O(\log^2 n)\mathbb{I}(E_{opt})$. However, the energy cost of the solution, relative to the optimum, i.e., the ratio $sum(\vec{P}(E))/sum(\vec{P}(E_{opt}))$, can be bounded only by choosing the labels carefully. We show that this algorithm can be efficiently implemented in a distributed fashion. As mentioned earlier, the above bounds hold for any node distribution. If the nodes are distributed randomly on a plane, it can be shown that the labels can be chosen randomly by the nodes, with a little worse approximation guarantees.

3. The Maximum Power Objective. We develop an $(O(\log^2 n), O(\log n))$ -approximation for the max power objective. This algorithm is randomized and is based on Monte-Carlo sampling. However, a crucial aspect of our algorithm is that the sampling is done with probability not p , but a different value $q < p$. The overall running time of this algorithm is $O(n \log^3 n)$. This algorithm is a non-trivial extension of a result by Goemans and Vondrak [8], who develop an elegant technique for covering minimum spanning trees (MSTs) of random subgraphs of a given graph G . While their

² We use $\log^r n$ to mean $(\log n)^r$.

result implies that the solution computed by the algorithm is strongly connected with probability at least $1 - 1/n^3$, it does not give any bounds on the cost of the resulting solution. Our proof of this bound crucially uses the geometric structure.

Next, we show that at the expense of larger running time, the approximation factor of the above algorithm can be improved to $(O(\log^2 n), 1)$ for the max objective. This algorithm is a simpler Monte-Carlo algorithm than the one above, and has a running time of $O(n^3 \log n)$. For any “testable” monotone property \mathbb{P} (such as strong connectedness, existence of giant component or low diameter; this is defined formally in Section V(B)), this algorithm can be used to find a good power level vector \vec{P} so that $G(V(p), \vec{P})$ has property \mathbb{P} .

This algorithm works with the same guarantees, even in a setting where the failure probabilities are non-uniform. Both of these algorithms are based on MST computations, and can be efficiently implemented in a distributed manner following techniques given in [30].

4. Empirical Results. We analyze the empirical performance of our algorithms for both the objective functions. We observe a sharp threshold in the variation of the max power level as a function of ϵ , the probability that $G(V(p), \vec{P})$ is not strongly connected; a similar, though less sharp, threshold is observed for the sum objective also. We also implement our algorithm for the sum objective, and find that the approximation guarantee in practice is significantly better than what we are able to prove analytically in Section IV.

The main focus of this paper is a theoretical analysis of this class of problems. The approximation guarantees we prove here are *worst case* guarantees, which hold for *any* instance. In practice, the approximation guarantees are much better, as we find in our empirical results. We believe that the framework of stochastic failures can model failures in many mission critical applications, and the proof techniques we develop here are likely to be useful in more general settings. The rest of the paper is organized in the following manner. We discuss related work in Section II, and the challenges and computational hardness of this problem in Section III. We discuss our algorithms for the total power and max power objectives in Sections IV and V, respectively. We discuss the empirical results in Section VI and conclude in Section VII. Because of space limitations, we have omitted the proof of Lemma III.1, which gives the #P-hardness result and have given a sketch of the proof of Lemma V.1, since it is very similar to the proof of Lemma IV.1.

II. RELATED WORK

Early work on the theoretical foundations of topology control, e.g., [27, 28, 5, 21] involved optimization of transmission power levels so that the resulting topology is well connected. The maximum power objective can be minimized efficiently for any monotone graph property that can be tested efficiently [21]. Under the total power minimization objective, topology control problems for many graph properties (e.g. connectedness, bounded diameter) are known to be NP-hard and approximation algorithms for many such problems have been developed, e.g., [5, 14, 11].

There has been a lot of work on finding k -connected networks, to deal with deterministic failure models, e.g., [11, 18, 10]. These algorithms provide performance guarantees which are polynomial functions of k . Bredin et al. [3] consider the problem of adding a minimum number of sensor nodes to a given sensor network so that the augmented network is k -node-connected. Since this problem is NP-hard, they present an approximation algorithm with a performance guarantee which is a polynomial in k .

To our knowledge, under the stochastic failure model, the topology control problem has not received much attention. As mentioned earlier, some results when nodes are arranged on a grid, or placed randomly in the plane have been reported in the literature. For example, Li et al. [19] consider a set of n points placed randomly in the unit square and establish bounds on the transmission radius to be used for each node so that the resulting graph is k -connected with high probability. Shakkottai et al. [29] analyze a configuration of nodes on a grid with stochastic failures. They establish relationships between the failure probability and the transmission radius to be used for each node to ensure coverage, connectivity and low diameter. Kumar et al. [15] obtain similar results for other random distributions of nodes.

In recent years, some papers such as [24] have pointed out several issues arising out of focusing on transmission power alone. Therefore, there has been interest in designing *low-interference* topologies, in addition to optimizing the transmission powers [4, 26, 25]. These papers have developed interference measures, which are related to the complexity of scheduling the set of chosen edges, and have designed approximation algorithms for optimizing the interference. However, these papers do not consider unreliable ad hoc networks.

III. CHALLENGES IN ANALYZING STOCHASTIC FAILURE MODELS

Most algorithms for topology control either allow no failures, or consider worst case models for failures, in which any k nodes can fail, as mentioned earlier. We now discuss to what extent these algorithms can be used for the stochastic failure models.

Consider an example of n nodes v_1, \dots, v_n arranged on a line, with uniform spacing of ℓ , and a power level $\phi(v_i, v_{i+1}) = \phi = c \cdot \ell^\gamma$ for all i . Suppose each node fails with probability $1/2$. Then, it can be shown that the random surviving subgraph is strongly connected with high probability if and only if every node has a power level of $\phi' = \Omega(\phi \log^\gamma n)$. Therefore, in order to be robust to stochastic failures, the power levels have to be much higher than what is chosen in a “failure-free” setting. The right power levels depend crucially on the probability with which we want to ensure the random surviving subgraph be connected, making this a non-trivial problem.

A number of papers on topology control [18, 10, 11, 19] have considered a “worst-case” notion of failures for arbitrary node distributions - they give algorithms for power choice to ensure connectivity when *any* set of k nodes fail. These results do not directly lead to efficient solutions for the TCPSF problem. For the instance described above, with high probability, $\Theta(n)$ nodes will fail. If we run the algorithms from [18] or [10] with $k = \Theta(n)$, the resulting power level choice will indeed guarantee that the surviving subgraph will be connected, but would have a total cost of $\Theta(n^\gamma / \log^\gamma n)$ times the optimum described above. Thus, algorithms that ensure k -connectivity do not necessarily give good solutions for stochastic failures.

However, from a theoretical perspective, the stochastic failures add a different kind of complexity to the TCPSF problem. For instance, even determining the probability that $G(V(p), \vec{P})$ has a large connected component (which seems to be much simpler than the TCPSF problem, and would be needed to “verify” a solution) is hard, as shown in the following lemma; its proof is a simple extension of a result in [1] and is omitted.

Lemma III.1: Let $\mu > 1/2$ be a constant. For an arbitrary set V of nodes and power level ϕ , determining $\Pr[G(V(p), \phi)$ has a component with at least $\mu|V|$ nodes] is #P-complete.

IV. THE TOTAL POWER OBJECTIVE

We describe Algorithm MINSUM-TCPSF for approximating the TCPSF problem with the sum objective. The algorithm consists of two steps. A distinct label is chosen for each node in the first step. The second step uses these labels to choose the set E of edges, and it is interesting to note that *any* choice of distinct labels ensures that $\Pr[G(V(p), E) \text{ is strongly connected}] \geq 1 - 1/n^2$, and $\mathbb{I}(E) = O(\log n)$. However, the cost, $\text{sum}(\vec{P}(E))$ depends on how the labels are chosen. The simplest method would be to choose the labels randomly for each node - this, in fact, works well if the nodes are distributed uniformly at random in the unit square. However, for arbitrary distributions of nodes, choosing labels randomly could lead to a high power cost (though, in fact, the interference is small). At the end of this section, we present a distributed implementation of Algorithm MINSUM-TCPSF.

Algorithm MINSUM-TCPSF:

1. Run algorithm CHOOSELABELNUMBERS to choose distinct labels for all nodes.
2. Run algorithm CHOOSEEDGES to find a set E of edges, and the corresponding power levels $\vec{P}(E)$.

Algorithm CHOOSELABELNUMBERS:

1. Construct an Euclidean minimum spanning tree T on the set V of nodes.
2. Root the tree T at some arbitrary node r .
3. Traverse T in depth-first order starting at r . Let v_1, \dots, v_n be the nodes in this order, with $v_1 = r$.
4. For each node v_i , we define $L(v_i) = n + 1 - i$.

Algorithm CHOOSEEDGES:

1. Let $k = c_1 \log_{1/(1-p)} n$ for some constant c_1 . Let the nodes in V be ordered v_1, \dots, v_n such that $L(v_1) < L(v_2) < \dots < L(v_n)$.
2. For each node v_i , find the smallest radius $r(v_i)$ such that ball $B(v_i, r(v_i))$ contains $\min\{k, n - i\}$ nodes with labels larger than $L(v_i)$. For each such larger-labeled node v_j , add the edges (v_i, v_j) and (v_j, v_i) to set Q .
3. Set $r'(v_i) = \max_{e=(v_i, v_j) \in Q} \{\ell(e)\}$ and $P(v_i) = c \cdot r'(v_i)^\gamma$.

Recall the notation defined in Section I(B). We first describe the two steps of our algorithm at




Fig. 2. An illustration of Algorithm Choose-Label-Numbers and some of the terms used in Lemma IV.2 on an instance with 7 nodes v_1, \dots, v_7 . The dashed line shows the traversal order, leading to the sequence $\pi = v_1, v_2, v_3, v_2, v_4, v_2, v_1, v_5, v_6, v_5, v_7$ and the order of the nodes $v_1, v_2, v_3, v_4, v_5, v_6, v_7$. Thus the resulting labels of the nodes are $L(v_1) = 7, L(v_2) = 6, L(v_3) = 5, L(v_4) = 4, L(v_5) = 3, L(v_6) = 2,$ and $L(v_7) = 1$. The reverse sequence of π is $\pi^R = v_7, v_5, v_6, v_5, v_1, v_2, v_4, v_2, v_3, v_2, v_1$.

an intuitive level here. Let $L(v)$ denote the label chosen using Algorithm CHOOSELABELNUMBERS. Algorithm CHOOSEEDGES is really simple - for each node v , it chooses edges to certain number of closest nodes of higher labels; let $r(v)$ be the length of the longest edge chosen by v . The power level for node v is then $P(v) = c \cdot r(v)^\gamma$, following the model described in Section B. We show in Theorem IV.1 that for this set E , $G(V(p), E)$ is strongly connected with high probability. To construct a “good” labeling of the nodes in Algorithm CHOOSELABELNUMBERS, we first build a minimum spanning tree (MST) and order the nodes based on a depth-first traversal order on this MST. This type of labeling ensures that for each node, the radius $r(v)$ chosen for it is small on average.

The two steps of Algorithm MINSUM-TCPSF are illustrated in Figures 2 and 3. The running time of this algorithm is $O(n^2 \log n)$ if implemented sequentially. The following theorem shows that the power levels chosen by this algorithm ensure that the surviving subgraph is strongly connected, with high probability.

Theorem IV.1: Let Q be the set of edges chosen by Algorithm CHOOSEEDGES. Then, the probability that $G(V(p), Q)$ is connected is at least $1 - 1/n^3$.

Proof: Let v_1, v_2, \dots, v_n be the nodes in increasing order of their label numbers, i.e., $L(v_i) = i$. For each node v_i , let $H(v_i) = \{v_j \in B(v_i, r(v_i)) : L(v_j) > L(v_i)\}$. Let $S = \{v_1, \dots, v_{n-k}\}$ and $T = V \setminus S$. Consider any node $v_i \in S$. By construction, we have $|H(v_i)| = k$, and so




Fig. 3. An illustration of Algorithm Choose-Edges with $k = 1$. The number next to each node v_i indicates its label number $L(v_i)$. The directed edge (v, w) implies that w is the closest node with label larger than $L(v)$. In this example, we have radius $r(v_2) = \ell(v_2, v_1)$, $r(v_3) = \ell(v_3, v_2)$, $r(v_4) = \ell(v_4, v_1)$ and $r(v_5) = \ell(v_5, v_4)$.

$\Pr[\text{all nodes in } H(v_i) \text{ fail to be in } V(p)] = (1-p)^k = \frac{1}{n^{c_1}}$. By the union bound, the probability that there is a node $v_i \in S$ such that all nodes in $H(v_i)$ fail is $\frac{1}{n^{c_1-1}}$, and so with probability $1 - \frac{1}{n^{c_1-1}}$, corresponding to each node $v_i \in S$, some node $u \in H(v_i)$ survives in a random subset $V(p)$.

By a similar argument, since $|T| = k$, at least one node from T survives in $V(p)$ with probability at least $1 - \frac{1}{n^{c_1}}$. Let w be the node of the largest label that survives from set T in the random set $V(p)$. By construction, we have $H(v_i) = \{v_{i+1}, \dots, v_n\}$ for each $v_i \in T$. Therefore, with probability $1 - \frac{1}{n^{c_1-1}} - \frac{1}{n^{c_1}}$, for each node $v \in V(p)$, $v \neq w$, some node $u \in H(v)$ survives. By construction of the topology in Algorithm CHOOSEEDGES, we have the edges (v, u) and (u, v) in $G(V(p), Q)$, which implies that $G(V(p), Q)$ is connected, since there is a path from any surviving node $v \in V(p)$ to node w (and back), passing through nodes of increasing label numbers. By choosing c_1 to be greater than 4, we have the result stated in the theorem. ■

We now bound the interference of the set Q chosen by the algorithm.

Lemma IV.1: $\mathbb{I}(Q) = O(\log^2 n)\mathbb{I}(E_{opt})$.

Proof: For any edge $e = (u, v)$, we show that $\mathbb{I}(e) = O(\log^3 n)$. Let $Q' = \{e' = (u', v') \in \mathbb{I}(e) : L(u') < L(v')\}$, where $L()$ denotes the labels chosen by Algorithm CHOOSELABELNUMBERS; then $|Q'| \geq \mathbb{I}(e)/2$. Let $A_1 = \{e' = (u', v') \in Q' : u' \in B(u, \ell(e)) \cup B(v, \ell(e))\}$. Let $A_2(i) = \{e' \in Q' - A_1 : \ell(e') \in [2^i, 2^{i+1}]\}$, for $i \leq c \log n$ for some constant c . Now we show that each of the sets A_1 and $A_2(i)$, for all i , has size at most $O(\log^2 n)$.

First, consider the set A_1 . Let V_1 denote the set of end points of edges in A_1 that lie in

$B(u, \ell(e)) \cup B(v, \ell(e))$. Note that if $|V_1| \geq c_2 \log n$, for some constant c_2 , by a simple packing argument, it follows that there exists some node $u_1 \in V_1$ such that $B(u_1, \ell(e)/2)$ contains at least k nodes of higher label than u_1 . This implies that Algorithm CHOOSEEDGES would choose $r(u_1) < \ell(e)$, which contradicts the fact that some edge $e_1 = (u_1, v_1) \in Q'$; therefore, it follows that $|V_1| = O(\log n)$. By construction, each node has an outdegree of $O(\log n)$ in the set Q chosen by Algorithm CHOOSEEDGES, and therefore, $|A_1| = O(\log^2 n)$.

Next, consider the set $A_2(i)$ and let $V_2(i)$ denote the set of all end points u' such that $(u', v') \in A_2(i)$ for some node v' . As before, we can show that if $|V_2(i)| > c_3 \log n$ for some constant c_3 , then there exists a node $u_2 \in A_2(i)$ such that $B(u_2, 2^{i-1})$ would contain k nodes of higher label than u_2 , which contradicts the assumption that Algorithm CHOOSEEDGES added some incident edge (u_2, v_2) to Q' . Therefore, $|V_2(i)| = O(\log n)$. Again, since the outdegree of each node in Q' is $O(\log n)$, we have $|A_2(i)| = O(\log^2 n)$. Therefore, $\mathbb{I}(e) = O(\log^3 n)$.

Finally, observe that $\mathbb{I}(E_{opt}) = \Omega(\log n)$ since the optimum solution E_{opt} must have at least one node v with degree $\Omega(\log n)$, in order for E_{opt} to be robust to failures. It then follows that for the shortest edge e incident on v , $|\mathbb{I}(e)| = \Omega(\log n)$. ■

Next, we bound the cost $sum(\vec{P})$ of the solution produced by Algorithm CHOOSEEDGES, relative to the optimum in the following lemma. The specific labeling assigned by Algorithm CHOOSELABELNUMBERS turns out to be crucial.

Lemma IV.2: Let Q be the set of edges computed in Algorithm CHOOSEEDGES. We have $\sum_{e \in Q} \ell(e)^\gamma \leq (2k)^{\gamma+2} \sum_{e \in T} \ell(e)^\gamma$, where T is the Euclidean minimum spanning tree on the set V of nodes.

Proof: Let π be the in-order traversal on T , which contains the exact sequence of nodes visited - this includes nodes repeated during the traversal (see Figure 2). Let π^R denote the reverse sequence of π .

Consider any node v_i and the minimal subsequence $\pi_j^R, \pi_{j+1}^R, \dots, \pi_{j'}^R$ such that: (i) $\pi_j^R = v_i$ and $\pi_{j'}^R$ is the last occurrence of v_i in π^R (ii) there are $k' = \min\{k, n - i\}$ distinct nodes in the sequence $\pi_{j+1}^R, \dots, \pi_{j'}^R$. For simplicity of notation, define $w_1 = \pi_{j+1}^R, w_2 = \pi_{j+2}^R, \dots, w_{j'-j} = \pi_{j'}^R$. As the degree of any node in an Euclidean MST can be at most 5, some node may occur as much as 5 times in this subsequence. However, a subsequence of size $2k'$ is sufficient to have k' distinct

nodes; because each edge is traversed at most twice, and each traversed edge introduce a new node to the sequence. Thus $j' - j \leq 2k' \leq 2k$. Let A denote the set of distinct nodes that occur in this subsequence. By construction, all these nodes have label numbers larger than $L(v_i)$, since they appear before the first occurrence of v_i in sequence π . Consider the radius $r(v_i)$ computed in Algorithm CHOOSEEDGES. Radius $r(v_i)$ denotes the smallest radius such that ball $B(v_i, r(v_i))$ contains k' nodes with label numbers larger than $L(v_i)$. That is, $r(v_i) \leq \max\{\ell(v, w) : w \in A\}$.

Next, observe that

$$\left(\sum_{s=1}^{j'-j-1} \ell(w_s, w_{s+1}) \right)^\gamma \leq (2k)^\gamma \sum_{s=1}^{j'-j-1} \ell(w_s, w_{s+1})^\gamma$$

because of the fact that $j' - j \leq 2k$. By the triangle inequality, we have $\max\{\ell(v, w) : w \in A\} \leq \sum_{s=1}^{j'-j-1} \ell(w_s, w_{s+1})$, which implies $r(v_i)^\gamma \leq (2k)^\gamma \sum_{s=1}^{j'-j-1} \ell(w_s, w_{s+1})^\gamma$. Let $H(v_i)$ denote the set of nodes with label numbers larger than $L(v_i)$ in the set $B(v_i, r(v_i))$. By definition, for each $w \in H(v_i)$, we have $\ell(v_i, w) \leq r(v_i)$, which implies $\sum_{w \in H(v_i)} \ell(v_i, w)^\gamma \leq k(2k)^\gamma \sum_{s=1}^{j'-j-1} \ell(w_s, w_{s+1})^\gamma$. In this case, we say that node v_i places a charge of $k(2k)^\gamma$ on each of the edges (w_s, w_{s+1}) along this subsequence.

By construction, we have $Q = \cup_{v_i} \{(v_i, w) : w \in H(v_i)\}$, and therefore, $\sum_{e \in Q} \ell(e)^\gamma$ can be expressed in terms of the costs of the edges in the subsequence π^R . The only problem is that edges in this subsequence could appear in the summations of a number of nodes, and we need to bound this charge. When we consider a node v_i , we only consider the subsequence of π^R starting at v_i of length at most $2k$. Therefore, an edge (w_s, w_{s+1}) could get charged by at most $2k$ such nodes v_i . This implies

$$\begin{aligned} \sum_{e \in Q} \ell(e)^\gamma &= \sum_{v_i} \sum_{w \in H(v_i)} \ell(v_i, w)^\gamma \\ &\leq \sum_{v_i} k(2k)^\gamma \sum_{s=1}^{j'-j-1} \ell(w_s(v_i), w_{s+1}(v_i))^\gamma \\ &\leq 2k^2(2k)^\gamma \sum_{s \geq 1} \ell(\pi_s^R, \pi_{s+1}^R)^\gamma \\ &\leq 2^2 k^2 (2k)^\gamma \sum_{e \in T} \ell(e)^\gamma \end{aligned}$$

■

Theorem IV.2: Algorithm MinSum-TCPSF is an $(O(\log^2 n), O(\log^{\gamma+2} n))$ -approximation algorithm for the sum objective.

Proof: Since any edge $e \in Q$ can contribute to the radii of at most two nodes, the two end points of e , we have $\text{sum}(\vec{P}(Q)) \leq 2 \sum_{e \in Q} c\ell(e)^\gamma$.

It is easy to see that if we construct an MST using the weights $c\ell(e)^\gamma$ for all edge e , it would be the same MST using the weights $\ell(e)$. Along with this fact, using Theorem 3.2 (Claim 1) in [13], we have, for any ϵ , $\sum_{e \in \text{MST}(V)} c\ell(e)^\gamma \leq \text{sum}(P_\epsilon^{\vec{opt}})$. Thus using Lemma IV.2, we have $\text{sum}(\vec{P}) \leq 2(2k)^{\gamma+2} \sum_{e \in \text{MST}} c\ell(e)^\gamma \leq 2(2k)^{\gamma+2} \text{sum}(P_\epsilon^{\vec{opt}})$. With $k = O(\log n)$ the result now follows from Lemma IV.1. ■

A. Distributed Implementation

Algorithms CHOOSEEDGES and CHOOSELABELNUMBERS can be implemented in a distributed manner, as described below.

1. Constructing MST and choosing root. Construct an MST using the distributed algorithm due to Gallager, Humblet and Spira (GHS) [7]. This algorithm takes $O(n \log n)$ time and $O(|E| + n \log n)$ messages. (Reference [12] discusses how this algorithm can be adapted to run in a wireless network setting.) The GHS algorithm also elects a leader, which serves as a root of the MST.

2. Node counting. The root broadcasts a *count* message using the tree edges to all other nodes. A leaf-node after receiving the *count* message, immediately sends back a *count-reply* message to its parent with count = 1. Any intermediate node waits until it receives *count-reply* from all of its children, then aggregates the count and sends a *count-reply* message to its parent with count equal to the aggregated count. Thus, the root can determine the total number of nodes n . Each intermediate node (and the root) also stores the counts received from all of its children; thus it knows the number of nodes in the subtree rooted at each of its children.

3. Label number selection. The root picks the label number n and divides the range $[a..b] = [1..n - 1]$ as follows: let the root have t children and the node counts received from its children are C_1, C_2, \dots, C_t ; the order of the children is determined by the reverse depth-first order. Then the root sends the range $[L_i, L_i + C_i - 1]$ to its i^{th} child, where $L_i = a + \sum_{j=1}^{i-1} C_j$. Similarly, an intermediate node, after receiving the range $[a..b]$ from its parent, picks the label number b and distributes the range $[a..b - 1]$ to its descendants.

Example. In Figure 2, the root $r = v_1$ has total count 7. In reverse depth-first order, count C_1

is the number nodes in the subtree rooted at v_5 and C_2 the number of nodes in the subtree rooted at v_2 . Root v_1 picks the label number 7 and sends the range $[1 .. 3]$ to v_5 and $[4 .. 6]$ to v_2 . Node v_5 selects 3 and sends $[1 .. 1]$ to v_7 and $[2 .. 2]$ to v_6 , and so on.

4. Finding the k the Nearest Nodes with Larger Labels. Each node executes the following algorithm to find the closest nodes with larger label numbers. Let d be the largest possible distance between any two nodes.

1. Initialize R to be the distance to the closest neighbor.
2. Each node v repeats the following until $R/2 \geq d$ or k nodes with larger label numbers are found.
 - (a) Set transmission radius r to R and broadcast a message containing $L(v)$.
 - (b) Any node u , on receiving the message from v , sends back $L(u)$ to v iff $L(u) > L(v)$.
 - (c) Set $R \leftarrow 2R$.

V. THE MAXIMUM POWER LEVEL

We now discuss algorithms for the max objective. In Section A, we describe a randomized algorithm that gives an $(O(\log^2 n), O(\log n))$ -approximation for the max objective, with a running time of $O(n \log^3 n)$. In Section V(B), we describe another algorithm which gives an improved approximation of $(O(\log^2 n), 1)$, but with a higher running time of $O(n^4 \log n)$, thus illustrating the tradeoff between the quality of approximation and running time. Both algorithms are simple, and involve Monte Carlo sampling. However, the analysis, is non-trivial.

A. An $(O(\log^2 n), O(\log n))$ -approximation algorithm

Our algorithm MINMAX-TCPSF builds on an interesting result by [8], who show that a small subset of $O(n \log_b n)$ edges can “cover” the MST of the surviving subgraph, with high probability. However, we need a non-trivial extension of their analysis to bound the approximation guarantee for the algorithm, which crucially uses the geometric properties. Recall the notation from Section I(B). We assume that the power threshold $\phi(u, v)$ for any pair of nodes $u, v \in V$ is given by $\phi(u, v) = c \cdot \ell(u, v)^\gamma$ for some constants c and γ .

Algorithm MINMAX-TCPSF:

- I. Let $b = 1/(1 - p)$ and $k = \lceil 5 \log_b n \rceil + 1$.

II. For $i = 1$ to $32ek^2 \ln n$ do

1. Generate a random subset $V_i = V(q)$ by choosing each vertex independently from V with probability $q = 1/k$.
2. Find a Euclidean minimum spanning tree T_i of the complete graph on V_i , with the length of each edge (u, v) equal to $\ell(u, v)$.
3. For each edge e , include e in Q if it appears in at least $16 \ln n$ different T_i 's.

III. Set $r = \max_{e \in Q} \ell(e)$ for all nodes $v \in V$.

We need the following result due to [8], which shows that the set Q constructed in the above algorithm contains an MST of the random subgraph $V(q)$, with high probability.

Theorem V.1 ([8]) Let Q be the set computed in the above algorithm. Then, we have $|Q| \leq 10en \log_b n + O(n)$, and $\Pr\{MST(V(p)) \subseteq Q\} > 1 - \frac{1}{n^3}$.

We first show that the interference complexity of Q is low. We only sketch its proof, since it is very similar to that of Lemma IV.1.

Lemma V.1: $\mathbb{I}(Q) = O(\log^2 n) \mathbb{I}(E_{opt})$.

Proof: **(Sketch)** For any edge $e = (u, v) \in Q$, we prove that $\mathbb{I}(e) = O(\log^3 n)$. Let $A_i(e) = \{e' \in \mathbb{I}(e) : \ell(e') \in [2^i, 2^{i+1})\}$, for each i , and let $V_i(e) = \{u' : \exists v' \text{ such that } (u', v') \in \mathbb{I}(e), \text{ and } v' \in B(u, \ell(e)) \cup B(v, \ell(e))\}$ be the set of end points of edges in $A_i(e)$ that do not lie in $B(u, \ell(e)) \cup B(v, \ell(e))$. If $|V_i(e)| > c_4 \log^2 n$ for some constant c_4 , there would exist a node $u_4 \in V_i(e)$ such that $|B(u_4, \ell(e)/2) \cap V_i(e)| \geq c_5 \log n$. A direct application of Chernoff's bound implies that $|B(u_4, \ell(e)/2) \cap V_i(e) \cap V(q)| = \Omega(\log n)$ with high probability, and therefore, many nodes in $B(u_4, \ell(e)/2) \cap V_i(e)$ will survive in step II(1) of Algorithm MINMAX-TCPSF. This implies that the Euclidean MST in Step II(2) will not end up choosing many edges between nodes in $B(u_4, \ell(e)/2) \cap V_i(e) \cap V(q)$ and nodes in $B(u, \ell(e)) \cup B(v, \ell(e))$, which contradicts the assumption that each node in $V_i(e)$ has an edge to some node in $B(u, \ell(e)) \cup B(v, \ell(e))$. Since there are $O(\log n)$ values of i , we have $\mathbb{I}(e) = O(\log^3 n)$. By the same argument as in the proof of Lemma IV.1, this lemma now follows. ■

Theorem V.2: The solution Q produced by Algorithm MINMAX-TCPSF is an $(O(\log^2 n), O(\log n))$ -approximation for the max power objective.

Proof: By Theorem V.1, it follows that $\Pr[G(V(p), Q) \text{ is strongly connected}] \geq 1 - \frac{1}{n^3}$. Let r_{opt}

be the transmission radius corresponding to the optimum solution. We show that with probability at least $1 - \frac{1}{n^3}$, every edge e chosen in Q satisfies $\ell(e) \leq c_1 r_{opt} \log_{b'} n$, where c_1 is a constant to be specified later, and $b' = 1/(1-q)$. To do so, we first bound the probability that an edge $e = (u, v)$ with $\ell(e) > c_1 r_{opt} \log_{b'} n$ is added to the set Q in the algorithm. Recall our assumption in Section B that $p^2 > 1/n^3$. We claim that $G(V, r_{opt})$ contains a path connecting u and v - if there is no such path connecting u and v , $G(V(p), r_{opt})$ would be disconnected whenever both u and v survive, which happens with probability $p^2 > 1/n^3$, implying that $\Pr[G(V(p), r_{opt}) \text{ is connected}] < 1 - 1/n^3$, which contradicts the definition of r_{opt} . Let $P = \langle w_1, \dots, w_k \rangle$ be any such path between $w_1 = u$ and $w_k = v$ in $G(V, r_{opt})$. By definition of $G(V, r_{opt})$, every edge $e' = (w_i, w_{i+1}) \in P$ must have $\ell(e') \leq r_{opt} < \frac{\ell(e)}{c_1 \log_{b'} n}$. By the triangle inequality, we have $|P| > c_1 \log_{b'} n$. Let c_2 be a constant such that $2c_2 < c_1$. We partition P into $k/k' = \lceil c_1/c_2 \rceil$ blocks $B_1, \dots, B_{k/k'}$, each of size roughly $k' = c_2 \log_{b'} n$. We do the analysis assuming that k' is an integer; if this is not true, the analysis can be modified easily to deal with the slight non-uniformity in sizes. Therefore, $B_i = \{w_{(i-1)k'+1}, \dots, w_{ik'}\}$ as shown in Figure 4.

The probability that all the nodes in any block B_i fail in a random sample $V(q)$ is $(1-q)^{|B_i|} = (1-q)^{c_2 \log_{1/(1-q)} n} = \frac{1}{n^{c_2}}$. By the union bound, the probability that there is a block B_i in which all nodes fail in $V(q)$ is therefore at most $\frac{c_1}{c_2 n^{c_2}}$. Thus, with probability at least $1 - \frac{c_1/c_2}{n^{c_2}}$, at least one node survives in each block B_i in $V(q)$. Consider one such sample $V(q)$, and let $w_{g(i)}$ denote the node that survives in B_i , $i = 1, \dots, k/k'$. By the triangle inequality, we have

$$\begin{aligned} \ell(w_{g(i)}, w_{g(i+1)}) &\leq \sum_{j=g(i)}^{g(i+1)-1} \ell(w_j, w_{j+1}) \leq 2k' r_{opt} \\ &\leq 2k' \ell(e) / (c_1 \log n) \leq \frac{2c_2}{c_1} \ell(e). \end{aligned}$$

Since we have $2c_2 < c_1$, we would have $\ell(w_{g(i)}, w_{g(i+1)}) < \ell(e)$, for each $i = 1, \dots, k/k'$. This means we have an alternate path $u, w_{g(1)}, \dots, w_{g(k')}, v$ from u to v in $V(q)$, each of whose edges are shorter than e , with probability at least $1 - \frac{c_1/c_2}{n^{c_2}}$. We set $c_1 = 5$ and $c_2 = 2$. Then, in algorithm MINMAX-TCPSF, $\Pr[e \in T_i] \leq 1/n$, where T_i is the MST on the random subset $V(q)$. By a Chernoff bound, it follows that for any specific edge e with $\ell(e) > c_1 r_{opt} \log_{b'} n$, we have $\Pr[e \text{ appears in at least } 16 \ln n \text{ different } T_i\text{'s}] \leq \frac{1}{n^5}$. By the union bound, it follows that the probability Q contains any edge e' with $\ell(e') > c_1 r_{opt} \log_{b'} n$ is at most $\frac{n^2}{n^5} = \frac{1}{n^3}$. Therefore, the radius r

chosen by our algorithm satisfies $r \leq c_1 r_{opt} \log_{b'} n$ with probability at least $1 - 1/n^3$. By definition of r_{opt} and by Lemma V.1, the theorem follows. \blacksquare




Fig. 4. Edge (u, v) with $\ell(e) > c_1 r_{opt} \log_{b'} n$. The black circles denote the nodes w_1, \dots, w_k , and $B_1, \dots, B_{k'}$ are the blocks on path P , referred to in the proof. Node $w_{g(i)}$ is the survivor in block B_i .

B. A $(O(\log^2 n), 1)$ -approximation algorithm

In this section, we describe a simpler Monte-Carlo sampling based algorithm for the min objective, that works not just for connectivity, but any “testable” monotone property, and gives an $(O(\log^2 n), 1)$ -approximation, at the cost of a higher running time, and a lower probability of reliability. Let \mathbb{P} denote a monotone property that can be tested in polynomial time, and let $\mathcal{A}_{\mathbb{P}}$ denote such an algorithm to test for property \mathbb{P} ; examples of such properties are strong connectivity, existence of a giant component, low diameter, etc. For graph $G(V, \vec{P})$, let $\mathcal{A}_{\mathbb{P}}(G(V, \vec{P}))$ be 1 if $G(V, \vec{P})$ has property \mathbb{P} , and 0 otherwise. For a given power level ϕ , the following algorithm, Algorithm Test(ϕ), determines if a given candidate power level ϕ is adequate for all nodes. Using Algorithm Test(ϕ) as a query function, a binary search on the set of power thresholds for all edges, $\{\phi(u, v) \mid u, v \in V\}$, can find the required minimum ϕ by calling Test(ϕ) at most $O(\log n)$ times, since there are at most n^2 distinct power thresholds. Once the correct power threshold ϕ is found, we use the algorithm of [8] to find the set Q of edges.

Algorithm Test(ϕ):

Input: Set V of nodes on the plane, survival probability p for each node, and a common power level ϕ .

Output: YES, if $G(V(p), \phi)$ has property \mathbb{P} with probability at least $1 - 1/n$.

I. For $i = 1$ to n^2 do

1. Generate a random subset $V_i = V(p)$ of V by choosing each vertex from V with probability p .




Fig. 5. Connectivity of $G(V(p), r)$ using a structured topology with $n = 1749$ nodes.




Fig. 6. The variation of (an approximation to) the sum objective with ϵ . Note that the threshold is not as sharp as for the max objective.




Fig. 7. Approximation ratio $\sum(\vec{P}) / \sum(P_\epsilon^{opt})$ for random distributions of nodes computed by the Algorithm MinSum-TCPSF.

2. Compute $X_i = 1 - \mathcal{A}_P(G(V_i, \phi))$, which is 1 if $G(V_i, \phi)$ does not have property \mathbb{P} .

II. If $\sum_{i=1}^k X_i \leq 3n$, output YES.

Theorem V.3: If $\Pr[G(V(p), \phi) \text{ has property } \mathbb{P}] \geq 1 - \frac{1}{n}$, then $\Pr[\text{Test}(\phi) \text{ does not return YES}] \leq e^{-n}$. The solution Q computed by algorithm $Test()$ and the result of [8] is an $(O(\log^2 n), 1)$ -approximation to the max objective. It requires $O(n^2)$ invocations of algorithm \mathcal{A}_P ; for strong connectivity, this takes time $O(n^3 \log n)$.

VI. EMPIRICAL RESULTS

We present empirical results for the TCPSF problem. For most of our experiments, we use an instance consisting of a set of nodes distributed along the streets of the downtown Portland, OR, generated by the TRANSIMS mobility model [22]. We discuss the results for the max and sum objectives separately. In these experiments, we used $\gamma = 2$.

1. The max objective function. Figure 5 shows $\Pr[G(V(p), r) \text{ is connected}]$ versus radius r , for different values of the surviving probability p for the node distribution in Portland. Observe that there is a sharp threshold in the connectivity radius (or equivalently, the max power level needed) even in this setting. This threshold is known in the case of nodes distributed randomly on the plane [19, 9], or arranged on a grid [29], but it is surprising to observe this threshold even in this non-random setting.

2. The sum objective function. Figure 6 shows the variation of an approximation to the total power objective with ϵ for an instance of the Portland data. Observe that this has a similar threshold behavior as that for the max objective, though it seems to be a little less sharp.

For $\gamma = 2$, Theorem IV.2 gives a worst case bound of $O(\log^4 n)$ on the approximation factor of $sum(\vec{P})$ given by our algorithm MinSum-TCPSF to the optimum. In the average case, this bound can be much better. For randomly chosen 50-5000 nodes in a unit square, we computed this approximation factor for $\gamma = 2$. Figure 7 shows this experimental result. To understand its growth rate, we also plotted the function $\log n$ and $\log^2 n$. We observe that the growth rate of the approximation factor is even smaller than the growth rate of $\log^2 n$. In fact, the growth rate is pretty much close to that of $\log n$.

VII. CONCLUSION

In this paper, we study the problem of topology control for unreliable ad hoc networks, motivated by mission critical applications. We designed and analyzed algorithms that simultaneously achieve provably good approximations to multiple objectives, namely, connectivity, power efficiency and interference cost. Empirically, we find that our algorithm has significantly better performance guarantees in practice, than what we show analytically. Improving the performance guarantees, and extending our results to other topological properties (e.g., spanner) are interesting open problems. The specific algorithms and the analytical techniques we develop for the TCPSF problem would be useful in other applications requiring resilient topology design.

ACKNOWLEDGMENT

We would like to thank the referees for their valuable comments. M. Khan, V.S. Anil Kumar, and M. Marathe were partially supported by NSF Nets Award CNS-0626964, NSF Award CNS-0845700, NSF HSD Award SES-0729441, CDC Center of Excellence in Public Health Informatics Award 2506055-01, NIH-NIGMS MIDAS project 5 U01 GM070694-05, DTRA CNIMS Grant HDTRA1-07-C-0113 and NSF NETS CNS-0831633. G. Pandurangan was supported in part by NSF Award CCF-0830476.

REFERENCES

- [1] H. AboElFotouh and C. Colbourn. Computing 2-terminal reliability for radio-broadcast networks. *IEEE Transactions on Reliability*, 38(5):538–555, 1989.
- [2] H. Balakrishnan, C. Barrett, V.S. Anil Kumar, M. Marathe, and S. Thite. The Distance 2-Matching Problem and

- Its Relationship to the MAC Layer Capacity of Adhoc Wireless Networks. *special issue of IEEE J. Selected Areas in Communications*, 22(6):1069–1079, 2004.
- [3] J. Bredin, E. Demaine, M. Hajiaghayi, and D. Rus. Deploying sensor networks with guaranteed capacity and fault tolerance. *Proc. ACM MobiHoc*, pages 309–319, 2005.
- [4] M. Burkhart, P. Rickenbach, R. Wattenhofer, and A. Zollinger. Does topology control reduce interference? In *MobiHoc*, 2004.
- [5] G. Calinescu and P-J. Wan. Range assignment for high connectivity in wireless ad hoc networks. *Proc. International Conference on Ad hoc and Wireless Networks*, pages 235–246, 2003.
- [6] C. Chiasserini and M. Garetto. Modeling the performance of wireless sensor networks. In *INFOCOM*, 2004.
- [7] R. Gallager, P. Humblet, and P. Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Transactions on Programming Languages and Systems*, 5(1):66–77, January 1983.
- [8] M. Goemans and J. Vondrak. Covering minimum spanning trees of random subgraphs. *Random Structures and Algorithms*, 29:257–276, 2006.
- [9] P. Gupta and P. Kumar. Critical power for asymptotic connectivity in wireless networks. *Stochastic Analysis, Control, Optimization and Applications*, A Volume in Honor of W.H. Fleming. Eds. W. McEneaney, G. Yin, and Q. Zhang:547–566, 1998.
- [10] M. T. Hajiaghayi, N. Immorlica, and V. Mirrokni. Power optimization in fault-tolerant topology control algorithms for wireless multi-hop networks. *MobiCom'03*, pages 300–312, September 2003.
- [11] X. Jia, D. Kim, S. Makki, P. Wan, and C. Yi. Power assignment for k -connectivity in wireless ad hoc networks. *J. Combinatorial Optimization*, 9(2):213–222, March 2005.
- [12] M. Khan, G. Pandurangan, and V. S. Anil Kumar. Distributed algorithms for constructing approximate minimum spanning trees in wireless sensor networks. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 20(1):124–139, 2009.
- [13] L. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc. Power consumption in packet radio networks. *Theoretical Computer Science*, 243(1-2):289–305, July 2000.
- [14] 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*, 2865:187–198, October 2003.
- [15] S. Kumar, T. Lai, and J. Balogh. On k -coverage in a mostly sleeping sensor network. *MobiCom'04*, pages 144–158, September-October 2004.
- [16] V. S. Anil Kumar, Madhav V. Marathe, Srinivasan Parthasarathy, and Aravind Srinivasan. Algorithmic aspects of capacity in wireless networks. In *ACM SIGMETRICS*, pages 133–144, 2005.
- [17] V. S. Anil Kumar, Madhav V. Marathe, Srinivasan Parthasarathy, and Aravind Srinivasan. Provable algorithms for joint optimization of transport, routing and mac layers in wireless ad hoc networks. In *Proc. DialM-POMC Workshop on Foundations of Mobile Computing*, 2007.
- [18] N. Li and J. C. Hou. Flss: A fault-tolerant topology control algorithm for wireless networks. *MobiCom'04*, pages

275–286, September-October 2004.

- [19] X. Li, P.-J. Wan, Y. Wang, and C. Yi. Fault tolerant deployment and topology control in wireless networks. *Proc. ACM MobiHoc*, pages 117–128, 2003.
- [20] Yanjun Li, Zhi Wang, and Yeji Song. Wireless sensor network design for wildfire monitoring. In *Proc. of World Congress on Intelligent Control and Automation*, pages 109–113, 2006.
- [21] E. L. Lloyd, R. Liu, M. V. Marathe, R. Ramanathan, and S. S. Ravi. Algorithmic aspects of topology control problems for ad hoc networks. *Mobile Networks and Applications*, 10:19–34, 2005.
- [22] Los Alamos National Lab. *TRANSIMS: Transportation Analysis Simulation System*. <http://transims.tsasa.lanl.gov/>.
- [23] G. Lu, N. Sadagopan, B. Krishnamachari, and A. Goel. Delay efficient sleep scheduling in wireless sensor networks. In *INFOCOM*, 2005.
- [24] R. Min and A. Chandrakasan. Top five myths about the energy consumption of wireless communication. In *Mobile Computing and Communications Review*, 2003.
- [25] K. Moaveni-Nejad and X.-Y. Li. Low-interference topology control for wireless ad hoc networks. In *Proc. of IEEE SECON*, 2005.
- [26] T. Moscibroda and R. Wattenhofer. Minimizing interference in ad hoc and sensor networks. In *Proc. DIAL M-POMC*, pages 24–33, 2005.
- [27] R. Ramanathan and R. Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. *Proc. IEEE INFOCOM 2000*, pages 404–413, March 2000.
- [28] P. Santi. *Topology Control in Wireless Ad hoc and Sensor Networks*. John Wiley & Sons, Inc., New York, NY, 2005.
- [29] S. Shakkottai, R. Srikant, and N. Shroff. Unreliable sensor grids: Coverage, connectivity and diameter. *Ad Hoc Networks*, 3(6):702–716, 2005.
- [30] R. Wattenhofer, L. Li, P. Bahl, and Y. M. Wang. Distributed topology control for power efficient operation in multihop wireless ad hoc networks. In *Proceedings of IEEE INFOCOM*, 2001.