Minimizing Interference in Ad Hoc and Sensor Networks

Thomas Moscibroda
Computer Engineering and
Networks Laboratory, ETH Zurich
8092 Zurich, Switzerland
moscitho@tik.ee.ethz.ch

Roger Wattenhofer
Computer Engineering and
Networks Laboratory, ETH Zurich
8092 Zurich, Switzerland
wattenhofer@tik.ee.ethz.ch

ABSTRACT

Reducing interference is one of the main challenges in wireless communication, and particularly in ad hoc networks. The amount of interference experienced by a node v corresponds to the number of other nodes whose transmission range covers v. At the cost of communication links being dropped, interference can be reduced by decreasing the node's transmission power. In this paper, we study the problem of minimizing the average interference while still maintaining desired network properties, such as connectivity, point-to-point connections, or multicast trees. In particular, we present a greedy algorithm that computes an $O(\log n)$ approximation to the interference problem with connectivity requirement, where n is the number of nodes in the network. We then show the algorithm to be asymptotically optimal by proving a corresponding $\Omega(\log n)$ lower bound that holds even in a more restricted interference model. Finally, we show how the algorithm can be generalized towards solving the interference problem for network properties that can be formulated as a 0-1 proper function.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*computations on discrete structures*;
G.2.2 [**Discrete Mathematics**]: Graph Theory—*graph algorithms*;
G.2.2 [**Discrete Mathematics**]: Graph Theory—*network problems*

General Terms

Algorithms, Theory

Keywords

ad hoc networks, sensor networks, interference, approximation, primal-dual algorithms

1. INTRODUCTION AND RELATED WORK

The advent of wireless communication in almost all aspects of daily life has not only spawned a multi-billion dollar market, but has also changed people's lifestyle. While heterogenous *cellular networks* have become omnipresent, a novel, complementing network paradigm has been rapidly emerging in recent years. *Ad hoc and sensor networks* consist of autonomous devices that can communicate via short-range radio, without any a-priori built-in infrastructure. If two nodes are not within mutual transmission range, they communicate through intermediate nodes relaying their message. In view of the great potential of ad hoc and sensor networks in a variety of application scenarios such as disaster relief, community mesh networks, monitoring and surveillance, or data gathering, it is not surprising that there has been a flurry of research activity in the field.

One of the main challenges in both models of wireless communication is *interference*. Informally speaking, a node a may interfere with another node b if a's interference range unintentionally covers b. Consequently, the *amount of interference* that b experiences is the number of such nodes a. In frequency division multiplexing cellular network, reducing the amount of interference results in fewer channels which, in turn, can be exploited to increase the bandwidth per frequency channel. In systems using code division multiplexing, small interference helps in reducing coding overhead. In the context of ad hoc and sensor networks, there is an additional motivation for keeping interference low. In these networks consisting of battery-driven devices, energy is typically scarce and the frugal usage of it is critical in order to prolong system operability and network lifetime. In addition to enhancing throughput, minimizing interference may help in lowering node energy dissipation by reducing the number of collisions (or the amount of energy spent in an effort of avoiding them) and consequently retransmissions on the media access layer.

Interference can be reduced by having nodes send with less transmission power. The area covered by the smaller transmission range will contain fewer nodes, yielding less interference. On the other hand, reducing the transmission range has the consequence of communication links being dropped. However, there is surely a limit to how much the transmission power can be decreased. In ad hoc networks, if the node's transmission ranges become too small and too many links are abandoned, the network may become disconnected. Hence, transmission ranges must be assigned to nodes in such a way that the desired global network properties are maintained.

The most basic network property that we seek to maintain is *connectivity*. That is, we want the transmission ranges to be large enough such that there exists a communication path between all pairs of nodes in the network. However, more elaborate network properties such as multicast connectivity, spanning subgraphs with limited stretch factor, or simultaneous point-to-point connections between pairs of nodes are not only conceivable, but practically relevant. Obviously, there exists a trade-off between interference reduction on the one hand and the maintenance of these properties on the other hand. This allows us to formulate a clear-cut combinatorial optimization problem for each property of the form: assign transmission ranges to nodes such that the resulting communication graph satisfies the network property, and the induced interference is minimized.

In the ad hoc network community, much effort has been invested in finding network structures having desirable properties such as sparseness or low node-degree, e.g. [17, 14, 12, 13, 9]. Most protocols adopt structures from the field of computational geometry, focusing on preserving energy efficient paths or computing planar subgraphs in order to facilitate routing [1]. Yet, in spite of its practical importance, there has been virtually no algorithmic work dealing with *interference* reduction *explicitly*. Instead, it has been assumed that the above mentioned traditional network structures (which guarantee sparseness and low node-degree) also help in confining interference. By disproving this long-standing conjecture, a recent paper by Burkhart et al [2] has changed the focus in the wireless networking community towards an *explicit* study of interference in ad hoc radio networks [7, 16, 18]. The work of [2] is of practical nature and verifies the efficiency of the proposed solution by means of simulation. A recent algorithmic work is [18] in which the problem of minimizing the *maximum interference* is studied. For the special case when the nodes are located on the Euclidean line, [18] presents an $O(\sqrt[4]{n})$ approximation algorithm. For a different, less robust link-based definition of interference, the work of [16] presents optimal or approximate solutions of interference reduction. As for the node-based definition studied in this paper, [16] proposes a heuristic, but without providing analytic approximation guarantees.

The reason why interference problems are difficult to track in practice and challenging from a theoretical point of view is their somewhat counter-intuitive behavior. Consider for instance the simple network shown in Figure 1 consisting of nodes $a_1, \ldots, a_n$ on a line [18].
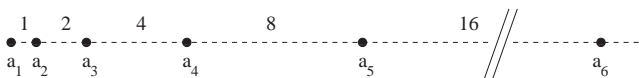

Figure 1: Example with nodes a_i being located at position $pos(a_i) = 2^i$, $i = 1, \ldots, n$.

Surprisingly, connecting the network in the obvious and intuitive way (each node connects to its left and right neighbor) produces a catastrophic result. Specifically, connecting the network linearly leads to linear interference, meaning that on average every node is disturbed by $n/2$ other nodes. In contrast, our approximation algorithm proposed in Section 3 yields an asymptotically optimal solution in which the average interference experienced by nodes is in $\Theta(\sqrt{n})$.

Moreover, it is intriguing that all commonly used structures such as the minimum spanning tree, shortest path trees, or relative neighborhood graphs produce the same linear list solution. In fact, any solution that includes the *nearest neighbor forest* is asymptotically worst possible. The downright failure of all these orthodox geometric/graph theoretical structures highlights the need for a different structuring of the communication graph in order to reduce interference explicitly. Ideally, such a structure for reducing interference combines the conflicting characteristics of both *clustering* (only few nodes have long transmission range) and traditional approaches seeking *sparseness* and *low node-degrees*. Informally, the algorithm we present in Section 3 constructs such a hybrid structure, reconciling the two conflicting aims.

In the literature on ad hoc networks, it is sometimes assumed that the coverage area of omnidirectional antennas form a clear-cut disk. Unfortunately, practical measurements show that modeling the coverage area by disks is far from realistic. In addition to inaccuracies of the antenna itself, the presence of obstacles to the propagation of electromagnetic signals can result in coverage areas that hardly resemble disks. In this paper, we therefore study the interference problem in a more general model.

We start out with the *general interference problem* which we model as a network property Φ and a point-set $S = A \cup P$ consisting of *active* nodes A and *passive* nodes P. Modeling the scenario exhibited in cellular networks or heterogenous two-tier ad hoc networks, nodes in set A need to maintain property Φ in the network, while interference is exhibited at nodes in P. For each transmission range chosen by a node $a \in A$, there is a (monotonously increasing) set of nodes R that it covers. The interference of a node $p \in P$ is the number of nodes $a \in A$ whose transmission range includes p. The *average interference* of a solution is the sum over all these interferences, divided by the number of passive nodes.

While modeling signal propagation by clear-cut disks is too idealistic, resorting to an abstract graph model in its full generality may be overly pessimistic. An intermediate model that we choose to study is therefore the *metric* case in which the transmission range corresponds to a distance, and a node interferes with all nodes being inside its transmission range. Besides this distinction between the non-metric and metric model, we additionally study another, orthogonal variant of the problem, the so-called *ad hoc model*. In these models, there is only one kind of node which experiences interference *and* is both responsible for maintaining network properties.

Surprisingly, we prove in this paper that from an approximation point of view, the general graph model and the metric model are equally hard. Specifically, for the basic connectivity property, we prove that a greedy algorithm achieves an approximation ratio of $O(\log n)$ in the *general interference problem* formulation. We then establish that this bound is asymptotically tight by presenting an $\Omega(\log n)$ approximation lower bound that holds even in the restricted *ad hoc metric* case. Finally, we also show how various other network properties can be solved. These include properties that can be modelled as 0-1 proper functions [8].

In fact, the upper bounds of some of the interference problems studied in this paper have been obtained earlier in an entirely different context. Particularly, the authors of [3, 4] study the problem of reducing the *total amount of energy* spent by the nodes in wireless ad hoc networks, while keep-

ing connectivity properties. By transforming the network instances of [3, 4] properly, their greedy algorithms can also be used to obtain $O(\log n)$ approximation ratios for the basic connectivity interference problem [3], as well as some of its generalizations [4]. Therefore, these nice and elegant algorithms for energy reduction were the first asymptotically optimal solutions for the interference problem. In this paper, however, we formulate and analyze the algorithms directly in the *interference model* which yields a proof that is different from [3, 4]. Additionally, our result on 0-1 proper functions is strictly more general than the connectivity requirements studied in [3, 4] and hence, is a true generalization of their result.

The remainder of the paper is organized as follows. In Section 2, we formally introduce the different models and problem formulations. Section 3 presents the greedy algorithm for the connectivity case. Tightness of this upper bound is shown in subsequent Section 4 in which we present the corresponding approximation hardness result. Section 5 generalizes the greedy algorithm to other network properties before Section 6 concludes the paper.

2. MODELS AND DEFINITIONS

Interference problems as studied in this paper model the situation in which a set of entities (such as nodes in a wireless multi-hop network) need to satisfy a certain global property Φ. We call these entities *active*. On the other hand, there are *passive* entities that experience interference caused by active nodes trying to fulfill property Φ. Passive entities may or may not correspond directly to active entities. An algorithm for the interference problem needs to assign large enough transmission ranges to active nodes such that Φ is fulfilled, while keeping the interference at passive nodes as small as possible.
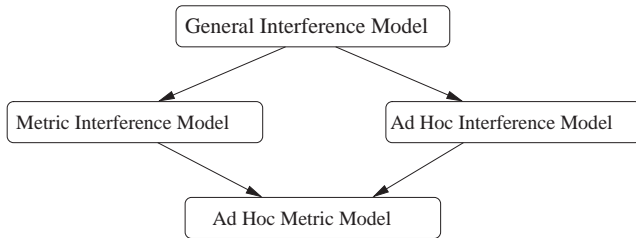

Figure 2: Relationship between interference models.

We can state the general interference problem as a tuple (S, Φ), where $S = A \cup P$ is a set of nodes representing the network entities. A corresponds to the set of *active nodes* and P contains the *passive nodes* in the network graph. The different transmission ranges of an active node $a \in A$ are denoted by a set $\mathcal{R}^a$ consisting of a monotonous sequence of inclusion sets, i.e., $\mathcal{R}^a = \{R_0^a, R_1^a, \ldots, R_{k_a}^a\}$ for $k_a \in \mathbb{N}_0^+$, such that for every i and for every $s \in S$, it holds that $R_i^s \subseteq S$ and $R_{i-1}^s \subset R_i^s$. For every $a \in A$, we define R_0^a to be the empty set. Informally, R_i^a is the set of all (passive and active) nodes that are disturbed/reached if a sends at transmission level i. A *power assignment* r is a mapping $r : A \to \mathbb{N}^+$ where $r(a) \leq k_a$ for all active nodes $a \in A$.

Given a mapping r of transmission ranges to active nodes, we can construct the *resulting graph* $RG(A, r)$. In this graph,

the vertex set corresponds to the points in A; there is an edge between two vertices $a_1, a_2 \in A$ if and only if these nodes are in *mutual* transmission range. Formally, we define $RG(A, r) = (V_{RG}, E_{RG})$ as $V_{RG} = A$ and

$$E_{RG} = \{(a_1, a_2) \mid a_1 \in R_{r(a_2)}^{a_2} \wedge a_2 \in R_{r(a_1)}^{a_1}\}.$$

Note that an edge between a_1 and a_2 exists in the resulting graph if and only if *both* nodes set their transmission range to a sufficiently high value. We call a resulting graph $RG(A, r)$ *valid* for property Φ, if and only if it satisfies property Φ.

In addition to the active nodes to which a is capable of connecting, the set R_i^a contains all passive nodes that experience interference *if* $r(a) = i$. The *interference* $\chi(p)$ *experienced* by a passive node $p \in P$ is defined as

$$\chi(p) = |\{a \in A \mid p \in R_{r(a)}^a\}|.$$

On the other hand, we say that an active node a transmitting at power level i *causes* an interference of t_i^a, where $t_i^a = |R_i^a \cap P|$. The *interference problem* consists of having the *active nodes* fulfill property Φ, while minimizing the average interference experienced at *passive nodes*.

MODEL 1. *In the general interference problem, we are given a tuple (S, Φ), where $S = A \cup P$ is a finite set of active and passive nodes and Φ is the property the resulting graph is required to fulfil. We need to find a power assignment r such that the resulting graph $RG(A, r)$ is valid for Φ and the average interference $\frac{1}{|P|} \sum_{p \in P} \chi(p)$ is minimized.*

Depending on the specific application at hand, various sub-classes of the average interference problem may be of interest. The ad hoc interference model captures the scenario in which there is only one kind of entity in the network, i.e., the set of active and passive nodes is equal.

MODEL 2. *In the ad hoc interference problem, the set of active and passive nodes is equal, i.e., $A = P$.*

In the metric version of the problem, we assume points $s \in S$ to be located in a metric space (S, μ). We slightly change the problem definition in the way that the power assignment r maps every active node to a real value, $r : A \to \mathbb{R}^+$. As in [3, 4], this real value represents a node's transmission range. In other words, a node $a \in A$ interferes with all nodes $b \in B$ whose distance to a is at most $r(a)$. It is clear that at the cost of a less intuitive definition, metric interference problems could also be stated using the notion of inclusions sets in the same way as in the general case.

MODEL 3. *In the metric interference problem, we consider the triple (M, Φ), where M is a finite metric space $M = (A \cup P, \mu)$ and Φ is the property the resulting graph is required to fulfill. A power assignment is a function $r : A \to \mathbb{R}^+$ of active nodes to a real value. The sets R_d^a are defined as $R_d^a = \{s' \mid s' \in S, \mu(a, s') \leq d\}$, where $d = r(a)$ denotes the transmission range assigned to a.*

MODEL 4. *The ad hoc metric interference problem is a special metric interference problem in which the set of active and passive nodes is equal, i.e., $A = P$.*

The relationship between the different models is shown in Figure 2. We conclude the section with a simple Lemma that is implicitly used throughout the paper. It says that the total amount of interference *experienced at passive nodes* $p \in P$ can equivalently be measured as the amount of interference *caused by active nodes* $a \in A$.

LEMMA 2.1. *Consider an instance of the interference problem. For any given power assignment r, it holds that*

$$\frac{1}{|P|}\sum_{p\in P}\chi(p) \;=\; \frac{1}{|P|}\sum_{a\in A} t^a_{r(a)}.$$

PROOF. The proof follows from

$$\sum_{p\in P}\chi(p) \;=\; \sum_{p\in P}|\{a\in A|p\in R^a_{r(a)}\}| = \sum_{a\in A}|R^a_{r(a)}|.$$

$\square$

3. APPROXIMATION ALGORITHM

In this section, we present a greedy algorithm that achieves an $O(\log n)$ approximation ratio, the same guarantee that can be achieved by the greedy algorithms given in [3, 4]. Our solution differs in the sense that it is directly formulated and analyzed in the interference model.

Many of the algorithmic challenges faced in designing approximation algorithms for interference problems stem from the fact that an edge in the resulting graph exists only if *both* nodes set their transmission range to a sufficiently large value. This leads to situations in which an edge (a_1, a_2) may be cheap (interference-wise) at node a_1, but costly at node a_2. Another peculiar aspect is that establishing an edge between two nodes (a_1, a_2) may potentially be free of charge if both nodes have an edge that requires higher transmission range (this is because we demand the interference sets R^a_i to be inclusion sets). Similarly, setting up an edge can be free of charge at one node, but expensive at the other. The combination of these characteristics render interference problems difficult to tackle in practice and appealing for theoretical study.

In this section, we consider the most basic graph property, i.e., we want the resulting graph to satisfy *connectivity*. Specifically, we propose a polynomial time approximation algorithm for the *general interference problem* that we show to be asymptotically optimal in Section 4.

3.1 Algorithm

Starting from a minimal, invalid power assignment, Algorithm 1 proceeds greedily and connects in every step the most cost-efficient star of nodes in the network by appropriately increasing the transmission ranges of the nodes in the star. The algorithm thus increases transmission ranges of stars until there remains only a single connected component in the resulting graph.

Let a *star* B consist of a center node and one or more leaf nodes. $\mathcal{B}$ denotes the set of all stars $B \subseteq A$. Given a power assignment r_i, a star is *active* if all its nodes (including the center) are mutually disconnected in $RG(A, r_i)$, i.e., for all pairs of nodes $a_1, a_2 \in B$, a_1 and a_2 are in different connected components in $RG(A, r_i)$. Further, we need the following two definitions. Let $f_a(a')$ denote the smallest k such that $a' \in R^a_k$, and let $g_a(\ell)$ be the smallest k such that $t^a_k = |R^a_k| \geq \ell$. That is, $f_a(a')$ denotes the minimal transmission range which enables a to reach a', and $g_a(\ell)$ is the smallest transmission range such that a interferes with at least ℓ other nodes. If for two nodes $a, a' \in A$, there does not exist any k such that $a' \in R^a_k$, then $f_a(a') = \infty$ and $t^a_{f_a(a')} = \infty$. For a star B, the *induced power assignment* $\hat{r}_B$ denotes the *minimal* power assignment for which all nodes $a' \in B$ have an edge to B's center a in $RG(A, \hat{r}_B)$. Formally,

for a star B with center a, the induced power assignment for node a' is defined as

$$\hat{r}_B(a') = \begin{cases} f_{a'}(a) & , a' \in B \setminus \{a\} \\ \max_{a^* \in B \setminus \{a\}} f_a(a^*) & , a' = a. \\ 0 & , a' \notin B. \end{cases}$$

Given a power assignment r_i, we define the *cost* of a star B as $cost(B) = \sum_{a \in B} cost(a)$ where

$$cost(a) = \begin{cases} t^a_{\hat{r}_B(a)} & , t^a_{\hat{r}_B(a)} > t^a_{r_i(a)} \\ 0 & , t^a_{\hat{r}_B(a)} \leq t^a_{r_i(a)}. \end{cases}$$

Intuitively, the cost of a star measures the amount of *additional* interference created by connecting the nodes of the star by edges, given the current power assignment r_i. Particularly, $cost(a')$ of a leaf node a' is 0, if $r_i(a')$ is already large enough such that the center node a_c is in $R^{a'}_{r(a')}$, i.e., if the transmission range of a' does not need to be increased in order to reach a_c. If, on the other hand, the transmission range of a node a' needs to be increased in order to connect a star, $cost(a')$ is the full interference created by $t^{a'}_{\hat{r}_B(a')}$ (not merely the incremental cost that comes from increasing $r_i(a')$ to $\hat{r}_B(a')$). Algorithm 1 chooses in every round the active star with minimal cost efficiency, and (minimally) increases the transmission ranges of the nodes such that all leaf nodes are connected to the star's center by an edge.

Algorithm 1 Greedy Algorithm

1: $i = 0$;
2: **forall** $a \in A$ **do** $r_0(a) := 0$;
3: **while** $RG(A, r_i)$ is not valid **do**
4:    $\mathcal{B}_i := \{B \in \mathcal{B} \mid B$ is active$\}$;
5:    $\tilde{B}_i := \mathrm{argmin}_{B \in \mathcal{B}_i}(cost(B)/|B|)$;
6:    **forall** $a \in \tilde{B}_i$ **do** $r_{i+1}(a) := \max\{\hat{r}_{\tilde{B}_i}(a), r_i(a)\}$;
7:    **forall** $a \notin \tilde{B}_i$ **do** $r_{i+1}(a) := r_i(a)$;
8:    $i := i + 1$;
9: **end while**

THEOREM 3.1. *Algorithm 1 runs in polynomial time and produces a valid power assignment r_{ALG}.*

PROOF. Correctness follows directly from the stop criterion of the while loop. As for runtime, observe that one greedy-step can be implemented in polynomial time in spite of there being an exponential number of stars $B \in \mathcal{B}_i$. For a given center-node $a \in A$, we compute $cost(a')$ for every $a' \in A$ as the cost of connecting a' to a. Let $a'_1, a'_2, \ldots a'_k$ be an ordering of $a' \in A \setminus \{a\}$ such that $cost(a'_1) \leq cost(a'_2) \leq \ldots \leq cost(a'_k)$. The set of leaf nodes of the most cost efficient star B (i.e., the star minimizing $cost(B)/|B|$) centered at a is of the form $\{a'_1, a'_2, \ldots a'_i\}$, for some $1 \leq i \leq k$. For a given center-node, it is therefore possible to choose its minimal cost star in polynomial time. Consequently, we can find the global minimum cost star $\tilde{B}_i$ by repeating the process for all $a \in A$. Finally, note that the number of iterations is at most $|A|$. $\square$

Before moving on to proving the algorithm's approximation ratio, we would like to add a remark on some practical issues related to the above algorithm. In physical reality, a node's interference range exceeds its transmission range.

Therefore, a node may experience interference from a transmitting node even though it is located outside this node's transmission range. Our algorithm can easily be extended to handle this case. All that needs to be done is to adjust the definition of $g_a(\ell)$. Specifically, it denotes the smallest transmission range such that at least ℓ other nodes are located within a's *interference range*. The remainder of the algorithm as well as its analysis remains the same.

3.2 Analysis

The analysis makes use of an integer linear program (ILP) formulation of the interference problem. Let $\mathcal{T}$ denote the set of all power assignments r, and let $\mathcal{T}_{inv} \subseteq \mathcal{T}$ be the set of all *invalid* assignments. A pair (a, k_a) denotes a specific power level at node a. We define the *border* of a given power assignment r as $\delta(r) = \{(a,k) \mid a \in A, k \in \mathcal{R}^a, k > r(a)\}$. That is, the border of r contains all power assignment pairs (a,k) that would denote an increase of the transmission range at a as compared to assignment r. Using these definitions, the interference problem can be captured as the following covering integer linear program.

$$\min \sum_{a \in A, k \in \mathcal{R}^a} x_{a,k} t_k^a$$

$$\text{s.t.} \sum_{\substack{a \in A, k \in \mathcal{R}^a \\ (a,k) \in \delta(r)}} x_{a,k} \geq 1 \qquad , \forall r \in \mathcal{T}_{inv} \qquad (1)$$

$$x_{a,k} \in \{0,1\} , \forall a \in A, k \in \mathcal{R}^a.$$

For each power level (a,k), the variable $x_{a,k}$ is set to 1 if a's transmission range is set to k, and to 0 otherwise. By Lemma 2.1, minimizing the objective function will result in a minimal total interference. The optimal *average interference* is then obtained by dividing the optimal value by $|P|$. The inequality constraint (1) guarantees that for every invalid power assignment $r \in \mathcal{T}_{inv}$, at least one node in its border must increase its power level. It can be shown that condition (1) is both sufficient and necessary to identify a valid power assignment r that yields a connected resulting graph $RG(A, r)$. In the relaxed version of the above program's dual, there is a variable α_r for each invalid power assignment $r \in \mathcal{T}_{inv}$.

$$\max \sum_{r \in \mathcal{T}_{inv}} \alpha_r$$

$$\text{s.t.} \sum_{r:(a,k) \in \delta(r)} \alpha_r \leq t_k^a , \forall a \in A, k \in \mathcal{R}^a$$

$$\alpha_r \geq 0 , \forall r \in \mathcal{T}_{inv}.$$

The following key lemma relates the cost of the star picked in each iteration of the greedy procedure to $\chi(r_{OPT})$, the cost incurred by the optimal valid power assignment.

LEMMA 3.2. *Let r_{OPT} be the optimal valid power assignment and let $\chi(r_{OPT})$ the interference incurred by r_{OPT}. Further, let λ_i denote the number of connected components in $RG(A, r_i)$. For every iteration i of Algorithm 1, it holds that*

$$\frac{cost(\tilde{B}_i) \cdot \lambda_{i-1}}{|\tilde{B}_i|} \leq \chi(r_{OPT}).$$

PROOF. We use the dual linear program to prove the claim. Specifically, we show that in every iteration i, we can distribute the cost of $cost(\tilde{B}_i)/|\tilde{B}_i|$ to dual variables α_r

corresponding to each of the λ_{i-1} connected components in $RG(A, r_{i-1})$. Specifically, we can do so in a way that the assignment constitutes a feasible dual solution. By the laws of LP duality, the objective value of a dual feasible solution is a lower bound on $\chi(r_{OPT})$, a fact which will allow us to derive the claimed approximation ratio.

Let C be one of the λ_{i-1} connected components that exist in $RG(A, r_{i-1})$. For each C, we define a series $r_1^c, \ldots, r_m^c$ of $m = \lceil cost(\tilde{B}_i)/|\tilde{B}_i| \rceil$ different invalid power assignments $r_j^c \in \mathcal{T}_{inv}$ as

$$r_j^c(a) := \begin{cases} g_a(j) - 1 & , a \in C \\ \min_{a' \in C} f_a(a') - 1 & , a \in \Gamma_j(C) \\ \infty & , \text{else} \end{cases}$$

where $\Gamma_j(C)$ denotes the set of nodes $a \in A \setminus C$ for which there exists an $a' \in C$ such that $j > t_{f_{a'}(a)}^{a'}$. The dual variable $\alpha_{r_j^c}$ assigned to power assignment r_j^c is defined as

$$\alpha_{r_j^c} := \begin{cases} 1 & , j < m \\ \frac{cost(\tilde{B}_i)}{|\tilde{B}_i|} - (m-1) & , j = m. \end{cases}$$

Note that with this definition of $\alpha_{r_j^c}$, the sum of dual variables at each connected component C is exactly

$$\sum_{j=1}^{m} \alpha_{r_j^c} = \frac{cost(\tilde{B}_i)}{|\tilde{B}_i|}. \qquad (2)$$

The total amount of dual variables assigned to invalid power assignments $r_j^c \in \mathcal{T}_{inv}$ in iteration i is therefore given by $\lambda_{i-1} \cdot cost(\tilde{B}_i)/|\tilde{B}_i|$.

We claim that the dual solution thus assigned in every iteration of Algorithm 1 is feasible. Assume for contradiction that there exists a power level (a, k) with $a \in A$ and $k \in R^a$ for which the dual condition is violated, i.e., $\sum_{r:(a,k) \in \delta(r)} \alpha_r > t_k^a$. Assume further that a is in connected component C. We distinguish two cases, depending on whether t_k^a is larger or smaller than $cost(\tilde{B}_i)/|\tilde{B}_i|$.

In the first case, $t_k^a < cost(\tilde{B}_i)/|\tilde{B}_i|$. We first show that the sum of the dual variables assigned to power assignments r_j^c (i.e., assignments corresponding to component C) is at most t_k^a. By the definition of r_j^c and its border $\delta(r_j^c)$, it holds that $(a,k) \in \delta(r_j^c)$ if and only if $g_a(j) - 1 < k$. Because $g_a(j)$ is the smallest k' such that $j \leq t_{k'}^a$, there are exactly t_k^a power assignments r_j^c for which $(a,k) \in \delta(r_j^c)$. Because for each of these t_k^a assignments r_j^c, the dual variable $\alpha_{r_j^c} \leq 1$, it follows that the sum of dual variables assigned to power assignments r_j^c is at most t_k^a.

To conclude the argument in the first case, we now show that power assignment (a,k) is not in the border of any $r_j^{c'}$, for any $C' \neq C$. Assume for contradiction that there is a connected component C' and some j for which $(a,k) \in \delta(r_j^{c'})$. By definition of $\Gamma_j(C)$, this is the case only if there exists a node $a' \in C'$ for which $j > t_{f_{a'}(a)}^{a'}$ and thus $m - 1 \geq t_{f_{a'}(a)}^{a'}$ holds. In words, there must be a node $a' \notin C$ for which the number of nodes it interferes with at transmission level $f_{a'}(a)$ is less than j. Now, consider the star B^*

consisting of nodes a and a'. The cost of this star can be upper bounded by

$$
\begin{aligned}
cost(B^*) &\leq t_k^a + t_{f_{a'(a)}}^{a'} < \frac{cost(\tilde{B}_i)}{|\tilde{B}_i|} + t_{f_{a'(a)}}^{a'} \\
&\leq \frac{cost(\tilde{B}_i)}{|\tilde{B}_i|} + m - 1 \leq \frac{2cost(\tilde{B}_i)}{|\tilde{B}_i|}.
\end{aligned}
$$

The cost efficiency of star B^* is less than that of $\tilde{B}_i$ and therefore, Algorithm 1 would have picked star B^* instead of $\tilde{B}_i$, yielding the contradiction which proves that $(a, k) \notin \delta(r_j^{c'})$ for all $C' \neq C$.

For the second case, assume that $t_k^a \geq cost(\tilde{B}_i)/|\tilde{B}_i|$. Unlike in the first case, (a, k) may also be in the border of power assignments $r_j^{c'}$ for some connected components $C' \neq C$. Let $\mathcal{C}$ be the set of all components C' for which there exists a power assignments $r_j^{c'}$ such that $(a, k) \in \delta(r_j^{c'})$. Further, for a component C', let $\hat{a}$ denote the node $a' \in C'$ that minimizes $t_{f_{a'(a)}}^{a'}$, i.e., $\hat{a} \in C'$ is the node that can reach a with minimal interference. By the definition of $r_j^{c'}$, (a, k) is only in $r_j^{c'}$ if $j > t_{f_{\hat{a}(a)}}^{\hat{a}}$. Consequently, for $j = 1, \ldots, t_{f_{\hat{a}(a)}}^{\hat{a}}$, $(a, k) \notin r_j^{c'}$. Because the total amount of dual variable assigned to power assignments associated with component C' is exactly $cost(\tilde{B}_i)/|\tilde{B}_i|$ by (2), the sum of the dual variables $\alpha_{r_j^{c'}}$ for which $a \in \delta(r_j^{c'})$ is

$$
\sum_{r_j^{c'}:(a,k)\in\delta(r_j^{c'})} \alpha_{r_j^{c'}} = \frac{cost(\tilde{B})}{|\tilde{B}|} - t_{f_{\hat{a}(a)}}^{\hat{a}}.
$$

Summing up over all $C' \in \mathcal{C}$ yields

$$
\sum_{r:(a,k)\in\delta(r)} \alpha_r = \frac{|\mathcal{C}| \cdot cost(\tilde{B})}{|\tilde{B}|} - \sum_{C'\in\mathcal{C}} t_{f_{\hat{a}(a)}}^{\hat{a}}.
$$

By assumption, the left hand side is strictly larger than t_k^a which means that

$$
\frac{\sum_{C'\in\mathcal{C}} t_{f_{\hat{a}(a)}}^{\hat{a}} + t_k^a}{|\mathcal{C}|} < \frac{cost(\tilde{B})}{|\tilde{B}|}. \tag{3}
$$

Consider the star B^* consisting of a and all nodes $\hat{a}$ for $C' \in \mathcal{C}$. The cost efficiency of $cost(B^*)/|B^*|$ is upper bounded by the left hand side of the Inequality (3). This contradicts Algorithm 1 choosing star $\tilde{B}$ in iteration i. This concludes the proof of Lemma 3.2.

$\square$

Given Lemma 3.2, it is now straightforward to derive the main theorem using a proof technique used for instance in [11] and subsequently in [10].

THEOREM 3.3. *Let r_{ALG} and r_{OPT} be the power assignments obtained by Algorithm 1 and by an optimal algorithm, respectively. Further, $\chi(r_{ALG})$ and $\chi(r_{OPT})$ denote their respective average interferences. It holds that $\chi(r_{ALG}) \leq O(\log n) \cdot \chi(r_{OPT})$.*

PROOF. By the definition of Algorithm 1, it follows that the number of connected components λ_i is at most

$$
\begin{aligned}
\lambda_i &= \lambda_{i-1} - |\tilde{B}_i| + 1 \\
&\underset{\text{Lemma 3.2}}{\leq} \lambda_{i-1} - \frac{cost(\tilde{B}_i) \cdot \lambda_{i-1}}{\chi(r_{OPT})} + 1 \\
&\leq \lambda_{i-1} \cdot \left(1 - \frac{cost(\tilde{B}_i)}{2\chi(r_{OPT})}\right).
\end{aligned}
$$

Assume that Algorithm 1 requires w iterations. Because $\lambda_0 = |A|$ and $\lambda_{w-1} \geq 2$, it follows that

$$
2 \leq \lambda_{w-1} \leq |A| \prod_{\ell=1}^{w-1} \left(1 - \frac{cost(\tilde{B}_\ell)}{2\chi(r_{OPT})}\right)
$$

By taking the logarithm on both sides, this simplifies to

$$
\begin{aligned}
\ln 2 &\leq \ln|A| + \sum_{\ell=1}^{w-1} \left(1 - \frac{cost(\tilde{B}_\ell)}{2\chi(r_{OPT})}\right) \\
&\leq \ln|A| - \sum_{\ell=1}^{w-1} \frac{cost(\tilde{B}_\ell)}{2\chi(r_{OPT})}.
\end{aligned}
$$

By rearranging the above expression, we therefore obtain the following upper bound,

$$
\sum_{\ell=1}^{w-1} cost(\tilde{B}_\ell) \leq 2\chi(r_{OPT})(\ln|A| - \ln 2).
$$

Finally, observe that the total interference created by Algorithm 1, $\chi(r_{ALG})$ is at most $\sum_{\ell=1}^{w-1} cost(\tilde{B}_\ell)$, because whenever the transmission range of a node $a \in A$ is increased, $cost(\tilde{B}_\ell)$ accounts for (at least) the entire cost of this increase. The cost of the last iteration, $\tilde{B}_w$, is at most $\chi(r_{OPT})$ and the theorem thus follows. $\square$

In Section 5, we will show how a similar algorithmic approach can also help in solving the interference problem for various other network properties Φ.

4. LOWER BOUND

In this section, we establish a tight $\Omega(\log n)$ approximation lower bound for the average interference problem even in its most restricted model, the *ad hoc metric* model. Corresponding lower bounds for the more general interference models follow as a straightforward corollary. This implies that from an approximation point of view, all models introduced in Section 2 are asymptotically equally hard.

THEOREM 4.1. *Consider the ad hoc metric interference problem with property Φ being the connectivity of nodes S, $n = |S|$. There is no approximation algorithm having an approximation ratio better than $\Omega(\log n)$, unless it holds that $NP \in DTIME(n^{\log\log n})$.*

PROOF. The proof makes use of a reduction to the *set cover problem* that preserves the approximation. Consider a set cover instance I_s consisting of a collection $\mathcal{S}$ of subsets of a finite set E of elements. Let n and m denote the number of elements and sets, respectively. A solution to I_s is a subset $\mathcal{S}' \subseteq \mathcal{S}$ such that every element $e \in E$ is at least in one set $S \in \mathcal{S}'$. Let OPT_s be an optimal solution to I_s, i.e., one that covers all elements with as few sets as possible.
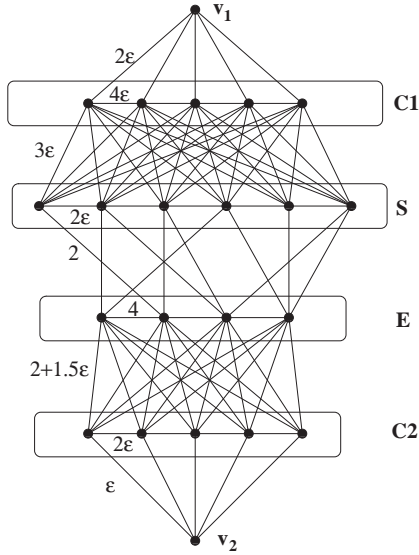
Figure 3: The lower bound instance I_{int} for the set cover instance $S_1 = \{b\}, S_2 = \{a, c\}, S_3 = \{b, c\}, S_4 = \{a, d\}, S_5 = \{d\}$, **and** $S_6 = \{c, d\}$ **over the universe** $E = \{a, b, c, d\}$.

Given I_s, we define an instance I_{int} of the interference problem consisting of $n + m + k_1 + k_2 + 2$ nodes, where k_1 is to be defined later and $k_2 = k_1 + m + n + 1$. The node-set A (which is equivalent to P) of I_{int} is partitioned into six clusters, two of which contain only a single node (see Figure 3 for an example). Clusters C_1 and C_2 consist of k_1 and k_2 nodes, respectively. For every set $S \in \mathcal{S}$ of I_s, we add one node to cluster C_S, and for every element $e \in E$, there is a node in cluster C_E. For a node $a_1 \in C_S$, let $S(a_1)$ denote the set $S \in \mathcal{S}$ that a_1 represents. Similarly, $e(a_2)$ is the element $e \in E$ corresponding to node a_2 in I_{int}. Finally, we denote the two remaining nodes by v_1 and v_2.

Considering the *ad hoc metric* version of the interference problem, we need to define the distance-function d that, along with the set of nodes S, defines the metric space. The distance between two nodes $a_1, a_2 \in A$ is defined as

$$d(a_1, a_2) = \begin{cases} 2\epsilon & , a_1 = v_1, a_2 \in C_1 \\ 4\epsilon & , a_1, a_2 \in C_1 \\ 3\epsilon & , a_1 \in C_1, a_2 \in C_S \\ 2\epsilon & , a_1, a_2 \in C_S \\ 2 & , a_1 \in C_S, a_2 \in C_E, \\ & \quad e(a_2) \in S(a_2) \\ 4 & , a_1, a_2 \in C_E \\ 2 + \frac{3\epsilon}{2} & , a_1 \in C_E, a_2 \in C_2 \\ 2\epsilon & , a_1, a_2 \in C_2 \\ \epsilon & , a_1 = v_2, a_2 \in C_2 \end{cases}$$

for an arbitrarily small ϵ. All other distances follow from symmetry or are induced by the shortest path metric. Specifically, notice that the distance between two nodes $a_1 \in C_E$ and $a_2 \in C_S$ for which $e(a_1) \notin S(a_2)$ is $2 + 2\epsilon$. Finally, let OPT_{int} be the optimal solution to the instance I_{int} of the interference problem, and let $\chi(OPT_{int})$ denote the average interference incurred by OPT_{int}.

It has been shown in [15, 6] that set cover cannot be approximated better than $\ln n$, unless $NP \in DTIME(n^{\log \log n})$.

Hence, we merely need to show that the following two implications hold

$$|OPT_s| \leq \alpha \implies \chi(OPT_{int}) \leq \beta$$
$$|OPT_s| > \ln n \cdot \alpha \implies \chi(OPT_{int}) > O(\log n) \cdot \beta$$

in order to prove the $\Omega(\log n)$ lower bound. Moreover, inspecting the proof of Feige's hardness result on set cover, it can be observed that the bound holds even for instances having $|S| \leq |E|$, e.g. [5, 6]. This allows us to use inequality $\ln(m + n) \leq \ln n + c$, for a small constant c.

We begin by showing that any solution Q_{int} to the interference problem in I_{int} can be transformed into a solution to the underlying set-cover instance I_s. Consider such a solution Q_{int}. Every element-node $a_i \in C_E$ must be connected to the rest of the graph. In the following lemma, we claim that we can transform ALG_{int} into a solution ALG'_{int} in which the ranges of all but one element-node is set to distance 2 without increasing interference. That is, each such element-node reaches all set-elements to which it belongs, but no other nodes.

LEMMA 4.2. *Given a solution Q_{int}, we can obtain in polynomial time a transformed solution Q'_{int} in which all but one $a_i \in C_E$ is assigned transmission range $r(a_i) < 2 + 3\epsilon/2$, and $\chi(Q'_{int}) \leq \chi(Q_{int})$.*

PROOF. Consider a solution Q_{int} in which more than one node $a_i \in C_E$ has transmission range at least $2 + 3\epsilon/2$. Every such node covers all nodes $a_j \in C_2$, causing an interference of k_2 at these nodes. We transform Q_{int} into a solution Q'_{int} as follows: For all but one $a_e \in C_E$ with $r(a_e) > 2 + 3\epsilon/2$, choose an $a_s \in C_S$ such that $e(a_e) \in S(a_s)$ and set $r(a_s) := 2$ and $r(a_e) = 2$. By doing so, the interference caused by each a_e is reduced by k_2. On the other hand, setting $r(a_s) := 2$ increases the interference by at most $1 + k_1 + n + |S(a_s)| \leq k_2$ by the definition of k_2. Clearly, Q'_{int} remains connected and hence $\chi(Q'_{int}) \leq \chi(Q_{int})$. $\square$

In the following, we given an upper bound on the optimal interference $\chi(OPT_{int})$ by relating it to the optimal solution OPT_s of the original set cover instance.

LEMMA 4.3. *Given a set cover instance I_s having an optimal solution OPT_s of size α_{OPT}. The optimal average interference $\chi(OPT_{int})$ in the corresponding interference instance I_{int} is bounded by*

$$\chi(OPT_{int}) \leq 5n + 5m + 2nm + m^2 + 2 + k_1(7 + \alpha_{OPT}).$$

PROOF. For the proof, we construct a power assignment r that yields a valid resulting graph $RG(A, r)$. The interference of r is an upper bound on $\chi(OPT_{int})$. The following power assignment r yields a valid resulting graph for I_{int}:

$$r(v_1) = 2\epsilon, \quad r(v_2) = \epsilon,$$
$$r(\hat{a}_1) = 3\epsilon, \quad \forall a_1 \in C_1 \setminus \{\hat{a}_1\} : r(a_1) = 2\epsilon,$$
$$r(\hat{a}_2) = 2 + 3\epsilon/2, \quad r(\hat{a}_E) = 2 + 3\epsilon/2,$$
$$\forall a_2 \in C_2 \setminus \{\hat{a}_2\} : r(a_2) = \epsilon,$$
$$\forall a_E \in C_E \setminus \{\hat{a}_e\} : r(a_E) = 2,$$
$$\forall a_S \in C_S : \begin{cases} 2\epsilon & , S(a_S) \notin OPT_s \\ 2 & , S(a_S) \in OPT_s \end{cases}$$

where $\hat{a}_1$, $\hat{a}_2$, and $\hat{a}_E$ are a single node from each C_1, C_2, and C_E, respectively. To see that the resulting graph is

connected, observe that there is an edge between all nodes in C_1, C_2 and v_1, v_2, respectively. Furthermore, set-nodes are linked as a linked list. The longer transmission range of $\hat{a}_1$ enables this node to set up a link to any $a_S \in C_S$ which occurs in the optimal solution OPT_s, thus connecting clusters C_1 and C_S. Because all nodes a_S with $S(a_S) \in OPT_s$ set their transmission range to 2, $RG(A, r)$ contains an edge between each such set and the element it covers, i.e., for every a_s having $S(a_S) \in OPT_s$ and every a_e with $e(a_e) \in S(a_S)$, $(a_s, a_e) \in E_{RG}$. Particularly, this means that all elements are connected. Finally, the *connector nodes* $\hat{a}_2$, and $\hat{a}_E$ ensure that cluster C_2 is connected to one element-node, and consequently to the rest of $RG(A, r)$.

Let ξ_C denote the total interference caused by nodes in set C. Because OPT_{int} is the power assignment resulting in the least possible interference, the interference of assignment r provides an upper bound on OPT_{int}. Thus, counting the interferences exhibited with assignment r, we can upper bound OPT_{int} as

$$
\begin{aligned}
OPT_{int} \quad \le \quad & \xi_{v_1} + \xi_{v_2} + \xi_{C_1} + \xi_{C_2} + \xi_{C_S} + \xi_{C_E} \\
\le \quad & k_1 + k_2 + (2k_1 - 1 + m) + \\
& + (2k_2 - 1 + n) + \xi_{C_S} + \xi_{C_E} \\
\le \quad & 3k_1 + 3k_2 - 2 + n + m + \\
& + \xi_{C_S} + k_2 + nm \\
\le \quad & 3k_1 + 4k_2 - 2 + n + m + nm + \\
& + m^2 + k_1 \alpha_{OPT} + n\alpha_{OPT} \\
< \quad & 5n + 5m + 2nm + m^2 + 2 + \\
& + k_1(7 + \alpha_{OPT}).
\end{aligned}
$$

The third inequality follows from $\xi_{\{\hat{a}_E\}} \le m + k_2$ as well as from $\xi_{C_E \setminus \{\hat{a}_E\}} \le (n-1)m$, because every element can clearly be in at most m sets. Since every set-nodes (regardless of whether it is in OPT_s) covers all other set-nodes, and because all nodes a_S with $S(a_S) \in OPT_s$ additionally interfere with all nodes in $C_1 \cup \{v_1\}$ and possibly as many as n elements, the fourth inequality holds. Because OPT_{int} is the power assignment resulting in the least possible interference, the claim follows. $\square$

Having obtained a bound for the first part of the reduction, we will now consider the implication $|OPT_s| > \ln n \cdot \alpha \Rightarrow \chi(OPT_{int}) > O(\log n) \cdot \beta$.

LEMMA 4.4. *Given a set cover instance I_s having an optimal solution OPT_s of size at least $\ln n \cdot \alpha_{OPT}$. The optimal average interference $\chi(OPT_{int})$ in the corresponding interference instance I_{int} is lower bounded by $\chi(OPT_{int}) \ge k_1 \ln n \cdot \alpha_{OPT}$.*

PROOF. Consider the optimal solution OPT_{int} to instance I_{int}. Applying the transformation of Lemma 4.2 to OPT_{int}, we obtain an optimal solution OPT'_{int} in which all but one $a_e \in C_E$ are assigned range $r(a_e) < 2 + 3\epsilon/2$. Observe that an element-node a_e can only be connected to set-nodes a_s for which $e(a_e) \in S(a_s)$, and only if $r(a_s) \ge 2$. On the other hand, every set-node with $r(a_s) \ge 2$ interferes with all k_1 nodes in C_1. Because by assumption at least $\ln n \cdot \alpha_{OPT}$ such sets are needed to cover all elements, $\chi(OPT'_{int})$ is at least $k_1 \ln n \cdot \alpha_{OPT}$. $\square$

Having established both upper (Lemma 4.4) and lower (Lemma 4.3) bounds, we can now conclude the proof of Theorem 4.1. Specifically, when defining $k_1 := 5m + 5n + 2nm +$

$m^2 + 2$, the inapproximability factor δ boils down to the expression

$$
\begin{aligned}
\delta \quad \ge \quad & \frac{k_1 \ln n \cdot \alpha_{OPT}}{5m + 5n + 2mn + m^2 + 2 + k_1(7 + \alpha_{OPT})} \\
\ge \quad & \frac{\ln n \cdot \alpha_{OPT}}{8 + \alpha_{OPT}} \ge \frac{\ln n}{8}. \quad (4)
\end{aligned}
$$

To conclude the proof, it remains to establish a relationship between n and $|A|$,

$$
\begin{aligned}
|A| \quad = \quad & n + m + k_1 + k_2 + 2 \\
= \quad & 12m + 12n + 4nm + 2m^2 + 7 \\
< \quad & 16(m + n)^2 + 7
\end{aligned}
$$

and hence, $\ln |A| < 2\ln(n+m) + c < 2\ln(n) + c'$, for constants c and c'. Plugging this into Inequality (4) concludes the proof of Theorem 4.1.

Having proved the lower bound for the *metric ad hoc* model, i.e., most restricted interference model, the following corollary is straightforward.

COROLLARY 4.5. *There is no algorithm that approximates the general interference problem better than $\Omega(\log n)$, unless $NP \in DTIME(n^{\log \log n})$.*

5. GENERALIZATION

The approach presented in Section 3 is not limited to solving the interference problem with regard to the basic connectivity problem. The algorithm can be generalized to handle various other network-design problems. In particular, we consider problems that can be formulated as special cut-covering problems. A function $f : 2^A \to \{0, 1\}$ is called a *0-1 proper function* if it satisfies the following properties [8]:

- $f(A) = 0$

- $f(C) = f(A \setminus C)$ for all $C \subseteq A$

- if C and C' are disjoint, then
  $f(C) = f(C') = 0 \Rightarrow f(C \cup C') = 0$

As explained in [8], the class of network design problems which can be formulated by 0-1 proper functions is particularly rich, encompassing problems such as the shortest path problem, connectivity, point-to-point connection problems, generalized Steiner tree problems (which includes the practically important case of multicast connectivity), T-join problems, and many others. As mentioned in the introduction, the work of [4] presents an $O(\log n)$ approximation algorithm to a variety of generalizations of the basic interference problem. However, the class of problems modelled using point-to-point requirements [4] is a proper subset of the problems that can be formulated by 0-1 proper functions (i.e., the opposite is not true). Hence, our result is a true generalization of the approximation result obtained in [4].

In order to solve the interference problems for properties Φ that can be modelled as 0-1 proper functions, we adjust Algorithm 1 as follows. Consider the set of connected components $\mathcal{C}$ in an iteration i. A component $C \in \mathcal{C}$ is called *active* if $f(C) = 1$. At each iteration, the adjusted algorithm considers only active components and tries to connect them in the most cost-efficient way. Particularly, let a star B be a set of active components centered at a node a.

Consider a path $p = (a_0, a_1, \ldots, a_\ell)$ from $a = a_0$ to an active component C with $a_\ell \in C$. The *cost* of p is defined as $cost(p) = \sum_{j=1}^{\ell} cost(a_j)$, where for $j = 1, \ldots, \ell$, $cost(a_j)$ is given by

$$cost(a_j) = \begin{cases} \max\left\{ t^{a_j}_{f_{a_j}(a_{j-1})}, t^{a_j}_{f_{a_j}(a_{j+1})} \right\} & , \theta(a_j) = 1 \\ 0 & , \theta(a_j) = 0. \end{cases}$$

where the indicator variable $\theta(a_j)$ is defined as

$$\theta(a_j) = \begin{cases} 1 & , \max\left\{ t^{a_j}_{f_{a_j}(a_{j-1})}, t^{a_j}_{f_{a_j}(a_{j+1})} \right\} > t^{a_j}_{r_i(a_j)} \\ 0 & , \max\left\{ t^{a_j}_{f_{a_j}(a_{j-1})}, t^{a_j}_{f_{a_j}(a_{j+1})} \right\} \leq t^{a_j}_{r_i(a_j)}. \end{cases}$$

Note that this definition is a generalization of the definition of $cost(a)$ in Section 3, where each node $a' \in B$ is simply a path of length 1. The cost of an active connected component $cost(C, a)$ with regard to a node a is the cost of the lightest path p_C from a to a node $a' \in C$ (This path can be computed in polynomial time using an adaptation of a shortest path algorithm). If $a \in C$, then $cost(C, a) = 0$. Finally, the cost of B centered at a is

$$cost(B) = \frac{1}{|B|} \left(\sum_{C \in B} cost(C, a) + \max_{\substack{\forall C \in B: \\ a_1 \in p_C}} t^a_{f_a(a_1)} \right).$$

The last term denotes the cost of node a in order to connect to all first nodes a_1 of the different paths.

With these adjusted definitions, the greedy step again consists of choosing the star $B \in \mathcal{B}$ with minimum cost efficiency $cost(B)/|B|$ in each iteration. That is, all active components in the picked star are connected by raising the transmission powers of all nodes in the elected paths to the minimum required values. Let a *terminal* be a node $a \in A$ for which $f(\{a\}) = 1$. We can prove that the adjusted generalized algorithm exhibits the following properties.

THEOREM 5.1. *Consider a network property Φ that can be formulated as a 0-1 proper function. The adjusted version of Algorithm 1 runs in polynomial time and produces a valid power assignment for property Φ. Further, the algorithm computes an $O(\log k)$ approximation, where k is the number of terminals.*

PROOF. By the 0-1 proper function properties, it follows that there cannot be one single active connected component during the course of the algorithm's execution. The number of active components being reduced by at least one in every iteration, the algorithm terminates leaving only inactive components. The proof of the approximation ratio follows exactly along the lines of the proof of Theorem 3.3. $\square$

6. CONCLUSION

We believe that interference in wireless networks (and particularly so in ad hoc and sensor networks) will be a major issue in the future that ought to be of interest to both practitioners and theoreticians alike. There are many paths for future research. The study of more complex network properties that need to be maintained is clearly one of them. Of most notable practical relevance appear to be fault-tolerant structures such as k-edge/node connected subgraphs (as pointed out in [4]) and, in hindsight of routing applications, subgraphs having bounded (hop or energy) stretch. It will also be interesting to see how well the interference problem can be approximated in metric spaces with restricted geometry, such as doubling metrics. The lower bound metric in Section 4 is not restricted, leaving open the possibility of sublogarithmic or even constant approximation algorithms in these cases.

8. REFERENCES

[1] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with Guaranteed Delivery in Ad Hoc Wireless Networks. In *Proceedings of the 3^{rd} International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM)*, pages 48–55, 1999.

[2] M. Burkhart, P. von Rickenbach, R. Wattenhofer, and A. Zollinger. Does Topology Control Reduce Interference. In *Proc. of the 5^{th} ACM Int. Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2004.

[3] G. Calinescu, S. Kapoor, A. Olshevsky, and A. Zelikovsky. Network Lifetime and Power Assignment in Ad Hoc Wireless Networks. In *Proceedings of 11^{th} European Symposium on Algorithms (ESA)*, pages 114–126, 2003.

[4] I. Caragiannis, C. Kaklamanis, and P. Kanellopolous. Energy-Efficient Wireless Network Design. *To appear in Theory of Computing Systems*, 2005.

[5] M. Chlebík and J. Chlebíková. Approximation Hardness of Dominating Set Problems. In *Proceedings of 12^{th} European Symposium on Algorithms (ESA)*, pages 192–203, 2004.

[6] U. Feige. A Threshold of ln n for Approximating Set Cover. *Journal of the ACM*, 45(4):634–652, 1998.

[7] M. Fussen, R. Wattenhofer, and A. Zollinger. Interference Arises at the Receiver. In *Proceedings of Int. Conference on Wireless Networks, Communications, and Mobile Computing (WIRELESSCOM)*, 2005.

[8] M. X. Goemans and D. P. Williamson. A General Approximation Technique for Constrained Forest Problems. *SIAM Journal of Computing*, 24(2):296–317, 1995.

[9] L. Jia, R. Rajaraman, and C. Scheideler. On Local Algorithms for Topology Control and Routing in Ad Hoc Networks. In *Proc. 15^{th} Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 220–229, 2003.

[10] P. Klein and R. Ravi. A Nearly Best-Possible Approximation Algorithm for Node-Weighted Steiner Trees. *Journal of Algorithms*, 19(1):104–115, 1995.

[11] F. T. Leighton and S. Rao. An Approximate Max-Flow Min-Cut Theorem for Uniform Multicommodity Flow Problems with Applications to Approximation Algorithms. In *Proceedings of the 29^{th} IEEE Symposium on the Foundations of Computer Science (FOCS)*, pages 422–431, 1988.

[12] L. Li, J. Y. Halpern, P. Bahl, Y.-M. Wang, and R. Wattenhofer. Analysis of a Cone-Based Distributed Topology Control Algorithm for Wireless Multi-Hop

Networks. In *Proc. 20th Symp. on Principles of Distributed Computing (PODC)*, pages 264–273, 2001.

[13] N. Li, J. Hou, and L. Sha. Design and Analysis of an MST-based Topology Control Algorithm. In *Proceedings of INFOCOM*, pages 1702–1712, 2003.

[14] X.-Y. Li, G. Calinescu, and P.-J. Wan. Distributed Construction of Planar Spanner and Routing for Ad Hoc Networks. In *Proceedings of INFOCOM*, 2002.

[15] C. Lund and M. Yannakakis. On the Hardness of Approximating Minimization Problems. *Journal of the ACM*, 41(5).

[16] K. Moaveni-Nejad and X. Y. Li. Low-Interference Topology Control for Wireless Ad Hoc Networks. In *Proceedings of the 2nd IEEE Conference on Sensor and Ad Hoc Communications and Networks (SECON)*, 2005.

[17] R. Ramanathan and R. Hain. Topology Control of Multihop Wireless Networks using Transmit Power Adjustment. In *Proceedings of INFOCOM*, pages 404–413, 2000.

[18] P. von Rickenbach, S. Schmid, R. Wattenhofer, and A. Zollinger. A Robust Interference Model for Wireless Ad Hoc Networks. In *Proc. 5th IEEE International Workshop on Algorithms for Wireless, Mobile, Ad-Hoc and Sensor Networks (WMAN)*, 2005.